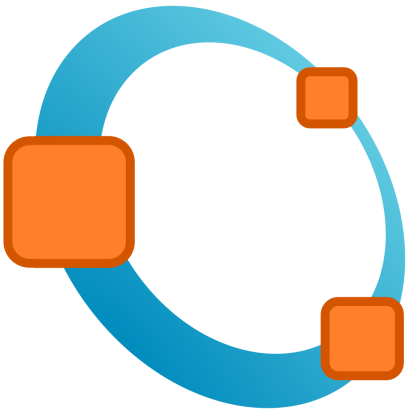
# GNU Octave

A high-level interactive language for numerical computations
Edition 11 for Octave version 11.3.0
August 2026



## Free Your Numbers

**John W. Eaton**
**David Bateman**
**Søren Hauberg**
**Rik Wehbring**

Copyright © 1996-2026 The Octave Project Developers

This is edition 11 of the Octave documentation, and is consistent with version 11.3.0 of Octave.

Permission is granted to make and distribute verbatim copies of this manual provided the copyright notice and this permission notice are preserved on all copies.

Permission is granted to copy and distribute modified versions of this manual under the conditions for verbatim copying, provided that the entire resulting derived work is distributed under the terms of a permission notice identical to this one.

Permission is granted to copy and distribute translations of this manual into another language, under the same conditions as for modified versions.

Portions of this document have been adapted from the `gawk`, `readline`, `gcc`, and C library manuals, published by the Free Software Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301–1307, USA.

# Table of Contents

**Preface** ..... 1
Acknowledgements ..... 1
Citing Octave in Publications ..... 5
How You Can Contribute to Octave ..... 6
Distribution ..... 6

**1 A Brief Introduction to Octave** ..... 7
1.1 Running Octave ..... 7
1.2 Simple Examples ..... 7
1.2.1 Elementary Calculations ..... 8
1.2.2 Creating a Matrix ..... 8
1.2.3 Matrix Arithmetic ..... 8
1.2.4 Solving Systems of Linear Equations ..... 9
1.2.5 Integrating Differential Equations ..... 9
1.2.6 Producing Graphical Output ..... 10
1.2.7 Help and Documentation ..... 10
1.2.8 Editing What You Have Typed ..... 11
1.3 Conventions ..... 11
1.3.1 Fonts ..... 11
1.3.2 Evaluation Notation ..... 11
1.3.3 Printing Notation ..... 12
1.3.4 Error Messages ..... 12
1.3.5 Format of Descriptions ..... 12
1.3.5.1 A Sample Function Description ..... 12
1.3.5.2 A Sample Command Description ..... 13

**2 Getting Started** ..... 15
2.1 Invoking Octave from the Command Line ..... 15
2.1.1 Command Line Options ..... 15
2.1.2 Startup Files ..... 19
2.2 Quitting Octave ..... 20
2.3 Commands for Getting Help ..... 21
2.4 Command Line Editing ..... 26
2.4.1 Cursor Motion ..... 26
2.4.2 Killing and Yanking ..... 27
2.4.3 Commands for Changing Text ..... 28
2.4.4 Commands for Completion ..... 28
2.4.5 Commands for Manipulating the History ..... 29
2.4.6 Customizing `readline` ..... 33
2.4.7 Customizing the Prompt ..... 33
2.4.8 Diary and Echo Commands ..... 35
2.5 How Octave Reports Errors ..... 36

2.6 Executable Octave Programs ... 37
2.6.1 Passing Arguments to Executable Scripts ... 38
2.6.2 Dual-Purpose Executable Scripts and Octave Functions ... 38
2.7 Comments in Octave Programs ... 39
2.7.1 Single Line Comments ... 39
2.7.2 Block Comments ... 39
2.7.3 Comments and the Help System ... 40

**3 Data Types ... 41**
3.1 Built-in Data Types ... 41
3.1.1 Numeric Objects ... 44
3.1.2 Missing Data ... 45
3.1.3 String Objects ... 46
3.1.4 Data Structure Objects ... 46
3.1.5 Cell Array Objects ... 46
3.2 User-defined Data Types ... 46
3.3 Object Sizes ... 46

**4 Numeric Data Types ... 51**
4.1 Matrices ... 52
4.1.1 Empty Matrices ... 56
4.2 Ranges ... 56
4.3 Single Precision Data Types ... 58
4.4 Integer Data Types ... 59
4.4.1 Hexadecimal and Binary Integer Constants ... 62
4.4.2 Integer Arithmetic ... 62
4.5 Bit Manipulations ... 63
4.6 Logical Values ... 66
4.7 Automatic Conversion of Data Types ... 67
4.8 Predicates for Numeric Objects ... 68

**5 Strings ... 75**
5.1 Escape Sequences in String Constants ... 75
5.2 Character Arrays ... 76
5.3 String Operations ... 78
5.3.1 Common String Operations ... 79
5.3.2 Concatenating Strings ... 81
5.3.3 Splitting and Joining Strings ... 85
5.3.4 Searching in Strings ... 88
5.3.5 Searching and Replacing in Strings ... 94
5.4 Converting Strings ... 99
5.4.1 String encoding ... 99
5.4.2 Numerical Data and Strings ... 100
5.4.3 JSON data encoding/decoding ... 109
5.5 Character Class Functions ... 113

**6 Data Containers ... 117**

6.1 Structures ... 117
6.1.1 Basic Usage and Examples ... 117
6.1.2 Structure Arrays ... 121
6.1.3 Creating Structures ... 122
6.1.4 Manipulating Structures ... 125
6.1.5 Processing Data in Structures ... 129
6.2 containers.Map ... 130
6.3 Cell Arrays ... 131
6.3.1 Basic Usage of Cell Arrays ... 131
6.3.2 Creating Cell Arrays ... 133
6.3.3 Indexing Cell Arrays ... 137
6.3.4 Cell Arrays of Strings ... 140
6.3.5 Processing Data in Cell Arrays ... 141
6.4 Comma-Separated Lists ... 141
6.4.1 Comma-Separated Lists Generated from Cell Arrays ... 142
6.4.2 Comma-Separated Lists Generated from Structure Arrays ... 143

**7 Variables ... 145**

7.1 Global Variables ... 147
7.2 Persistent Variables ... 148
7.3 Status of Variables ... 150

**8 Expressions ... 159**

8.1 Index Expressions ... 159
8.1.1 Advanced Indexing ... 162
8.2 Calling Functions ... 167
8.2.1 Call by Value ... 168
8.2.2 Recursion ... 169
8.2.3 Access via Handle ... 170
8.3 Arithmetic Operators ... 170
8.4 Comparison Operators ... 174
8.5 Boolean Expressions ... 176
8.5.1 Element-by-element Boolean Operators ... 176
8.5.2 Short-circuit Boolean Operators ... 177
8.6 Assignment Expressions ... 179
8.7 Increment Operators ... 183
8.8 Operator Precedence ... 184

**9 Evaluation ... 187**

9.1 Calling a Function by its Name ... 188
9.2 Evaluation in a Different Context ... 190

**10 Statements ... 193**

10.1 The if Statement ... 193
10.2 The switch Statement ... 195
10.2.1 Notes for the C Programmer ... 196
10.3 The while Statement ... 197
10.4 The do-until Statement ... 198
10.5 The for Statement ... 198
10.5.1 Looping Over Structure Elements ... 199
10.6 The break Statement ... 200
10.7 The continue Statement ... 201
10.8 The unwind_protect Statement ... 202
10.9 The try Statement ... 202
10.10 Continuation Lines ... 203

**11 Functions and Scripts ... 205**

11.1 Introduction to Function and Script Files ... 205
11.2 Defining Functions ... 205
11.3 Returning from a Function ... 208
11.4 Multiple Return Values ... 209
11.5 Variable-length Return Lists ... 212
11.6 Variable-length Argument Lists ... 213
11.7 Ignoring Arguments ... 214
11.8 Default Arguments ... 216
11.9 Validating Arguments ... 216
11.9.1 Validating the number of Arguments ... 216
11.9.2 Validating the type of Arguments ... 217
11.9.3 Parsing Arguments ... 222
11.10 Function Files ... 224
11.10.1 Manipulating the Load Path ... 227
11.10.2 Subfunctions ... 231
11.10.3 Private Functions ... 232
11.10.4 Nested Functions ... 232
11.10.5 Overloading and Autoloading ... 235
11.10.6 Function Locking ... 236
11.10.7 Function Precedence ... 237
11.11 Script Files ... 238
11.11.1 Publish Octave Script Files ... 239
11.11.2 Publishing Markup ... 242
11.11.2.1 Using Publishing Markup in Script Files ... 242
11.11.2.2 Text Formatting ... 243
11.11.2.3 Sections ... 243
11.11.2.4 Preformatted Code ... 244
11.11.2.5 Preformatted Text ... 244
11.11.2.6 Bulleted Lists ... 244
11.11.2.7 Numbered Lists ... 244
11.11.2.8 Including File Content ... 245
11.11.2.9 Including Graphics ... 245
11.11.2.10 Including URLs ... 245

11.11.2.11 Mathematical Equations ... 246
11.11.2.12 HTML Markup ... 246
11.11.2.13 LaTeX Markup ... 246
11.11.3 Jupyter Notebooks ... 246
11.12 Function Handles and Anonymous Functions ... 248
11.12.1 Function Handles ... 248
11.12.2 Anonymous Functions ... 250
11.13 Command Syntax and Function Syntax ... 250
11.14 Organization of Functions Distributed with Octave ... 252

**12 Errors and Warnings ... 255**
12.1 Handling Errors ... 255
12.1.1 Raising Errors ... 255
12.1.2 Catching Errors ... 258
12.1.3 Recovering From Errors ... 261
12.2 Handling Warnings ... 262
12.2.1 Issuing Warnings ... 262
12.2.2 Enabling and Disabling Warnings ... 271

**13 Debugging ... 273**
13.1 Entering Debug Mode ... 273
13.2 Leaving Debug Mode ... 274
13.3 Breakpoints ... 274
13.4 Debug Mode ... 278
13.5 Call Stack ... 279
13.6 Profiling ... 280
13.7 Profiler Example ... 282

**14 Input and Output ... 287**
14.1 Basic Input and Output ... 287
14.1.1 Terminal Output ... 287
14.1.1.1 Paging Screen Output ... 291
14.1.2 Terminal Input ... 293
14.1.3 Simple File I/O ... 294
14.1.3.1 Saving Data on Unexpected Exits ... 309
14.2 C-Style I/O Functions ... 311
14.2.1 Opening and Closing Files ... 311
14.2.2 Simple Output ... 313
14.2.3 Line-Oriented Input ... 314
14.2.4 Formatted Output ... 315
14.2.5 Output Conversion for Matrices ... 317
14.2.6 Output Conversion Syntax ... 317
14.2.7 Table of Output Conversions ... 318
14.2.8 Integer Conversions ... 319
14.2.9 Floating-Point Conversions ... 320
14.2.10 Other Output Conversions ... 320
14.2.11 Formatted Input ... 321

14.2.12 Input Conversion Syntax ... 323
14.2.13 Table of Input Conversions ... 323
14.2.14 Numeric Input Conversions ... 324
14.2.15 String Input Conversions ... 324
14.2.16 Binary I/O ... 325
14.2.17 Temporary Files ... 327
14.2.18 EOF (End of File) and Errors ... 329
14.2.19 File Positioning ... 330

**15 Plotting ... 333**

15.1 Introduction to Plotting ... 333
15.2 High-Level Plotting ... 333
15.2.1 Two-Dimensional Plots ... 333
15.2.1.1 Axis Configuration ... 366
15.2.1.2 Two-dimensional Function Plotting ... 378
15.2.1.3 Two-dimensional Geometric Shapes ... 381
15.2.2 Three-Dimensional Plots ... 382
15.2.2.1 Aspect Ratio ... 408
15.2.2.2 Three-dimensional Function Plotting ... 409
15.2.2.3 Three-dimensional Geometric Shapes ... 413
15.2.3 Plot Annotations ... 414
15.2.4 Multiple Plots on One Page ... 423
15.2.5 Multiple Plot Windows ... 425
15.2.6 Manipulation of Plot Objects ... 426
15.2.7 Manipulation of Plot Windows ... 427
15.2.8 Use of the "interpreter" Property ... 432
15.2.8.1 "none" interpreter ... 432
15.2.8.2 "tex" interpreter ... 432
15.2.8.3 "latex" interpreter ... 434
15.2.9 Printing and Saving Plots ... 435
15.2.10 Interacting with Plots ... 446
15.2.11 Test Plotting Functions ... 447
15.3 Graphics Data Structures ... 448
15.3.1 Introduction to Graphics Structures ... 448
15.3.2 Graphics Objects ... 449
15.3.2.1 Creating Graphics Objects ... 450
15.3.2.2 Handle Functions ... 453
15.3.3 Graphics Object Properties ... 460
15.3.3.1 Root Properties ... 460
15.3.3.2 Figure Properties ... 462
15.3.3.3 Axes Properties ... 470
15.3.3.4 Legend Properties ... 483
15.3.3.5 Line Properties ... 485
15.3.3.6 Text Properties ... 489
15.3.3.7 Image Properties ... 492
15.3.3.8 Patch Properties ... 495
15.3.3.9 Scatter Properties ... 501
15.3.3.10 Surface Properties ... 505

15.3.3.11 Light Properties ... 511
15.3.3.12 Uimenu Properties ... 513
15.3.3.13 Uibuttongroup Properties ... 516
15.3.3.14 Uicontextmenu Properties ... 519
15.3.3.15 Uipanel Properties ... 521
15.3.3.16 Uicontrol Properties ... 525
15.3.3.17 Uitable Properties ... 529
15.3.3.18 Uitoolbar Properties ... 534
15.3.3.19 Uipushtool Properties ... 536
15.3.3.20 Uitoggletool Properties ... 539
15.3.4 Searching Properties ... 542
15.3.5 Managing Default Properties ... 543
15.4 Advanced Plotting ... 545
15.4.1 Colors ... 545
15.4.2 Line Styles ... 545
15.4.3 Marker Styles ... 546
15.4.4 Callbacks ... 546
15.4.5 Application-defined Data ... 548
15.4.6 Object Groups ... 549
15.4.6.1 Data Sources in Object Groups ... 553
15.4.6.2 Area Series ... 554
15.4.6.3 Bar Series ... 554
15.4.6.4 Contour Groups ... 555
15.4.6.5 Error Bar Series ... 556
15.4.6.6 Line Series ... 557
15.4.6.7 Quiver Group ... 558
15.4.6.8 Stair Group ... 559
15.4.6.9 Stem Series ... 559
15.4.6.10 Surface Group ... 560
15.4.7 Transform Groups ... 561
15.4.8 Graphics Toolkits ... 561
15.4.8.1 Customizing Toolkit Behavior ... 562
15.4.8.2 Hardware vs Software Rendering ... 562
15.4.8.3 Precision issues ... 563

**16 Matrix Manipulation ... 565**

16.1 Finding Elements and Checking Conditions ... 565
16.2 Rearranging Matrices ... 569
16.3 Special Utility Matrices ... 580
16.4 Famous Matrices ... 593

**17 Arithmetic ... 603**

17.1 Exponents and Logarithms ... 603
17.2 Complex Arithmetic ... 605
17.3 Trigonometry ... 606
17.4 Sums and Products ... 611
17.5 Utility Functions ... 616

17.6 Special Functions . . . 628
17.7 Rational Approximations . . . 640
17.8 Coordinate Transformations . . . 641
17.9 Mathematical Constants . . . 643

**18 Linear Algebra . . . 647**
18.1 Techniques Used for Linear Algebra . . . 647
18.2 Basic Matrix Functions . . . 647
18.3 Matrix Factorizations . . . 657
18.4 Functions of a Matrix . . . 672
18.5 Specialized Solvers . . . 676

**19 Vectorization and Faster Code Execution . . . 689**
19.1 Basic Vectorization . . . 689
19.2 Broadcasting . . . 691
19.2.1 Broadcasting and Legacy Code . . . 694
19.3 Function Application . . . 695
19.4 Accumulation . . . 700
19.5 Memoization . . . 703
19.6 Miscellaneous Techniques . . . 704
19.7 Examples . . . 706

**20 Nonlinear Equations . . . 707**
20.1 Solvers . . . 707
20.2 Minimizers . . . 712

**21 Diagonal and Permutation Matrices . . . 717**
21.1 Creating and Manipulating Diagonal/Permutation Matrices . . . 717
21.1.1 Creating Diagonal Matrices . . . 718
21.1.2 Creating Permutation Matrices . . . 718
21.1.3 Explicit and Implicit Conversions . . . 719
21.2 Linear Algebra with Diagonal/Permutation Matrices . . . 720
21.2.1 Expressions Involving Diagonal Matrices . . . 720
21.2.2 Expressions Involving Permutation Matrices . . . 721
21.3 Functions That Are Aware of These Matrices . . . 722
21.3.1 Diagonal Matrix Functions . . . 722
21.3.2 Permutation Matrix Functions . . . 722
21.4 Examples of Usage . . . 722
21.5 Differences in Treatment of Zero Elements . . . 723

**22 Sparse Matrices ... 727**
22.1 Creation and Manipulation of Sparse Matrices ... 727
22.1.1 Storage of Sparse Matrices ... 727
22.1.2 Creating Sparse Matrices ... 728
22.1.3 Finding Information about Sparse Matrices ... 734
22.1.4 Basic Operators and Functions on Sparse Matrices ... 738
22.1.4.1 Sparse Functions ... 738
22.1.4.2 Return Types of Operators and Functions ... 739
22.1.4.3 Mathematical Considerations ... 740
22.2 Linear Algebra on Sparse Matrices ... 748
22.3 Iterative Techniques Applied to Sparse Matrices ... 758
22.4 Real Life Example using Sparse Matrices ... 766

**23 Numerical Integration ... 771**
23.1 Functions of One Variable ... 771
23.2 Orthogonal Collocation ... 780
23.3 Functions of Multiple Variables ... 781

**24 Differential Equations ... 791**
24.1 Ordinary Differential Equations ... 791
24.2 Differential-Algebraic Equations ... 794
24.3 Matlab-compatible solvers ... 802

**25 Optimization ... 813**
25.1 Linear Programming ... 813
25.2 Quadratic Programming ... 819
25.3 Nonlinear Programming ... 821
25.4 Linear Least Squares ... 823

**26 Statistics ... 829**
26.1 Descriptive Statistics ... 829
26.2 Statistics on Sliding Windows of Data ... 846
26.3 Basic Statistical Functions ... 857
26.4 Forecasting Metrics ... 863
26.5 Correlation and Regression Analysis ... 865
26.6 Distributions ... 869
26.7 Random Number Generation ... 869

**27 Sets ... 871**
27.1 Set Operations ... 873

**28 Polynomial Manipulations ... 877**
28.1 Evaluating Polynomials ... 877
28.2 Finding Roots ... 878
28.3 Products of Polynomials ... 879
28.4 Derivatives / Integrals / Transforms ... 882
28.5 Polynomial Interpolation ... 883
28.6 Miscellaneous Functions ... 893

**29 Interpolation ... 895**
29.1 One-dimensional Interpolation ... 895
29.2 Multi-dimensional Interpolation ... 899

**30 Geometry ... 905**
30.1 Delaunay Triangulation ... 905
30.1.1 Plotting the Triangulation ... 907
30.1.2 Identifying Points in Triangulation ... 910
30.2 Voronoi Diagrams ... 912
30.3 Convex Hull ... 916
30.4 Interpolation on Scattered Data ... 918
30.5 Vector Rotation Matrices ... 921

**31 Signal Processing ... 925**

**32 Image Processing ... 939**
32.1 Loading and Saving Images ... 939
32.2 Displaying Images ... 945
32.3 Representing Images ... 947
32.4 Plotting on top of Images ... 960
32.5 Color Conversion ... 960

**33 Audio Processing ... 963**
33.1 Audio File Utilities ... 963
33.2 Audio Device Information ... 964
33.3 Audio Player ... 965
33.3.1 Playback ... 966
33.3.2 Player Properties ... 967
33.4 Audio Recorder ... 967
33.4.1 Recording ... 968
33.4.2 Data Retrieval ... 969
33.4.3 Recorder Properties ... 969
33.5 Audio Data Processing ... 970

**34 Object Oriented Programming ... 973**
34.1 Creating a Class ... 973
34.2 Class Methods ... 975
34.3 Indexing Objects ... 978
34.3.1 Defining Indexing And Indexed Assignment ... 978
34.3.2 Indexed Assignment Optimization ... 982
34.4 Overloading Objects ... 983
34.4.1 Function Overloading ... 983
34.4.2 Operator Overloading ... 984
34.4.3 Precedence of Objects ... 985
34.5 Inheritance and Aggregation ... 987
34.6 `classdef` Classes ... 991
34.6.1 Creating a `classdef` Class ... 991
34.6.2 Properties ... 993
34.6.3 Methods ... 994
34.6.4 Inheritance ... 995
34.6.5 Value Classes vs. Handle Classes ... 996

**35 GUI Development ... 999**
35.1 I/O Dialogs ... 999
35.2 Progress Bar ... 1007
35.3 UI Elements ... 1008
35.4 GUI Utility Functions ... 1017
35.5 User-Defined Preferences ... 1021
35.6 Octave Workspace Windows ... 1023

**36 System Utilities ... 1025**
36.1 Timing Utilities ... 1025
36.2 Filesystem Utilities ... 1038
36.3 File Archiving Utilities ... 1048
36.4 Networking Utilities ... 1051
36.4.1 FTP Objects ... 1051
36.4.2 WWW Access ... 1053
36.4.3 Base64 and Binary Data Transmission ... 1058
36.5 Controlling Subprocesses ... 1058
36.6 Process, Group, and User IDs ... 1067
36.7 Environment Variables ... 1068
36.8 Current Working Directory ... 1069
36.9 Password Database Functions ... 1071
36.10 Group Database Functions ... 1072
36.11 System Information ... 1072
36.12 Hashing Functions ... 1081

**37 Packages ... 1083**
37.1 Background Information ... 1083
37.2 Installing and Removing Packages ... 1084
37.3 Using Packages ... 1089
37.4 Administrating Packages ... 1089
37.5 Creating Packages ... 1090
37.5.1 The DESCRIPTION File ... 1092
37.5.2 The INDEX File ... 1094
37.5.3 PKG_ADD and PKG_DEL Directives ... 1095
37.5.4 Missing Components ... 1095

**Appendix A External Code Interface ... 1097**
A.1 Oct-Files ... 1098
A.1.1 Getting Started with Oct-Files ... 1098
A.1.2 Matrices and Arrays in Oct-Files ... 1102
A.1.3 Character Strings in Oct-Files ... 1105
A.1.4 Cell Arrays in Oct-Files ... 1107
A.1.5 Structures in Oct-Files ... 1107
A.1.6 Sparse Matrices in Oct-Files ... 1109
A.1.6.1 Array and Sparse Class Differences ... 1109
A.1.6.2 Creating Sparse Matrices in Oct-Files ... 1110
A.1.6.3 Using Sparse Matrices in Oct-Files ... 1113
A.1.7 Accessing Global Variables in Oct-Files ... 1114
A.1.8 Calling Octave Functions from Oct-Files ... 1115
A.1.9 Calling External Code from Oct-Files ... 1116
A.1.10 Allocating Local Memory in Oct-Files ... 1119
A.1.11 Input Parameter Checking in Oct-Files ... 1119
A.1.12 Exception and Error Handling in Oct-Files ... 1120
A.1.13 Documentation and Testing of Oct-Files ... 1122
A.2 Mex-Files ... 1123
A.2.1 Getting Started with Mex-Files ... 1123
A.2.2 Working with Matrices and Arrays in Mex-Files ... 1125
A.2.3 Character Strings in Mex-Files ... 1127
A.2.4 Cell Arrays with Mex-Files ... 1128
A.2.5 Structures with Mex-Files ... 1129
A.2.6 Sparse Matrices with Mex-Files ... 1130
A.2.7 Calling Other Functions in Mex-Files ... 1134
A.3 Standalone Programs ... 1134
A.4 Java Interface ... 1138
A.4.1 Making Java Classes Available ... 1138
A.4.2 How to use Java from within Octave ... 1140
A.4.3 Set up the JVM ... 1142
A.4.4 Java Interface Functions ... 1143

**Appendix B Test and Demo Functions ... 1151**
B.1 Test Functions ... 1151
B.2 Demonstration Functions ... 1160

**Appendix C Obsolete Functions ... 1165**

**Appendix D Known Causes of Trouble ... 1173**

D.1 Actual Bugs We Haven't Fixed Yet ... 1173
D.2 Reporting Bugs ... 1173
D.2.1 Have You Found a Bug? ... 1173
D.2.2 Where to Report Bugs ... 1174
D.2.3 How to Report Bugs ... 1174
D.2.4 Sending Patches for Octave ... 1175
D.3 How To Get Help with Octave ... 1176
D.4 How to Distinguish Between Octave and Matlab ... 1176

**Appendix E Installing Octave ... 1179**

E.1 Build Dependencies ... 1179
E.1.1 Obtaining the Dependencies Automatically ... 1179
E.1.2 Build Tools ... 1179
E.1.3 External Packages ... 1180
E.2 Running Configure and Make ... 1183
E.3 Compiling Octave with 64-bit Indexing ... 1187
E.4 Installation Problems ... 1190

**Appendix F Grammar and Parser ... 1193**

F.1 Keywords ... 1193
F.2 Parser ... 1201

**Appendix G GNU GENERAL PUBLIC LICENSE ... 1203**

**Concept Index ... 1215**

**Function Index ... 1221**

**Operator Index ... 1237**

**Graphics Properties Index ... 1239**

# Preface

Octave was originally intended to be companion software for an undergraduate-level textbook on chemical reactor design being written by James B. Rawlings of the University of Wisconsin-Madison and John G. Ekerdt of the University of Texas.

Clearly, Octave is now much more than just another 'courseware' package with limited utility beyond the classroom. Although our initial goals were somewhat vague, we knew that we wanted to create something that would enable students to solve realistic problems, and that they could use for many things other than chemical reactor design problems. We find that most students pick up the basics of Octave quickly, and are using it confidently in just a few hours.

Although it was originally intended to be used to teach reactor design, it has been used in several other undergraduate and graduate courses in the Chemical Engineering Department at the University of Texas, and the math department at the University of Texas has been using it for teaching differential equations and linear algebra as well. More recently, Octave has been used as the primary computational tool for teaching Stanford's online Machine Learning class (`http://ml-class.org`) taught by Andrew Ng. Tens of thousands of students participated in the course.

If you find Octave useful, please let us know. We are always interested to find out how Octave is being used.

Virtually everyone thinks that the name Octave has something to do with music, but it is actually the name of one of John W. Eaton's former professors who wrote a famous textbook on chemical reaction engineering, and who was also well known for his ability to do quick 'back of the envelope' calculations. We hope that this software will make it possible for many people to do more ambitious computations just as easily.

Everyone is encouraged to share this software with others under the terms of the GNU General Public License (see Appendix G [Copying], page 1203). You are also encouraged to help make Octave more useful by writing and contributing additional functions for it, and by reporting any problems you may have.

## Acknowledgements

Many people have contributed to Octave's development. The following people have helped code parts of Octave or aided in various other ways (listed alphabetically).

| | | |
|---|---|---|
| Ben Abbott | Drew Abbot | NVS Abhilash |
| Andy Adler | Adam H. Aitkenhead | Fernando Alvarruiz |
| Joakim Andén | Giles Anderson | Joel Andersson |
| Lachlan Andrew | Pedro Angelo | Damjan Angelovski |
| Muthiah Annamalai | Markus Appel | Leonardo Araujo |
| Branden Archer | Willem Atsma | Marco Atzeri |
| Ander Aurrekoetxea | Shai Ayal | Sahil Badyal |
| Jeff Bai | Roger Banks | Ben Barrowes |
| Alexander Barth | David Bateman | Heinz Bauschke |
| Miguel Bazdresch | Julien Bect | Stefan Beller |
| Roman Belov | Markus Bergholz | Karl Berry |
| Andreas Bertsatos | Atri Bhattacharya | Ethan Biery |

David Billinghurst
Don Bindner
Jakub Bogusz
Gaël Bonithon
Leonardo S. Borges
Moritz Borgmann
Paul Boven
Richard Bovey
John Bradshaw
Marcus Brinkmann
Max Brister
Remy Bruno
Stefan Brüns
Clemens Buchacher
Ansgar Burchard
Antonius Burgers
Marco Caliari
Daniel Calvelo
John C. Campbell
Juan Pablo Carbajal
Jean-Francois Cardoso
Joao Cardoso
Larrie Carr
David Castelow
Vincent Cautaerts
Marco Cecchetti
Corbin Champion
Clinton Chee
Albert Chin-A-Young
Sunghyun Cho
Carsten Clark
Catalin Codreanu
J. D. Cole
Jacopo Corno
Andre da Costa Barros
Martin Costabel
Michael Creel
Richard Crozier
Jeff Cunningham
Martin Dalecki
Jacob Dawid
Jorge Barros de Abreu
Carlo de Falco
Thomas D. Dean
Philippe Defert
Bill Denney
Fabian Deutsch
Christos Dimitrakakis
Pantxo Diribarne
Vivek Dogra
John Donoghue
David M. Doolin
Carnë Draug
Sergey Dudoladov
Pascal A. Dupuis
John W. Eaton
Dirk Eddelbuettel
Pieter Eendebak
Paul Eggert
Stephen Eglen
Peter Ekberg
Abdallah K. Elshamy
Garrett Euler
Edmund Grimley Evans
Rolf Fabian
Francesco Faccio
Gunnar Farnebäck
Massimiliano Fasi
Stephen Fegan
Ramon Garcia Fernandez
Kasper H. Filtenborg
Torsten Finke
David Finkel
Guillaume Flandin
Colin Foster
Jose Daniel Munoz Frias
Brad Froehle
Castor Fu
Eduardo Gallestey
Walter Gautschi
Klaus Gebhardt
Driss Ghaddab
Eugenio Gianniti
Hartmut Gimpel
Michele Ginesi
Nicolo Giorgetti
Arun Giridhar
Michael D. Godfrey
Dave Goel
Michael Goffioul
Glenn Golden
Tomislav Goles
Keith Goodman
Brian Gough
Alexander Graf
Michael C. Grant
Steffen Groot
Etienne Grossmann
David Grundberg
Kyle Guinn
Vaibhav Gupta
Peter Gustafson
Kai Habel
Patrick Häcker
William P. Y. Hadisoeseno
Jaroslav Hajek
Benjamin Hall
Alexander Hansen
Kim Hansen
Gene Harvey
Søren Hauberg
Dave Hawthorne
Oliver Heimlich
Daniel Heiserer
Piotr Held
Martin Helm
Stefan Hepp
Martin Hepperle
Jordi Gutiérrez Hermoso
Israel Herraiz
Yozo Hida
Christian Himpe
Ryan Hinton
Roman Hodek
A. Scottedward Hodel
Júlio Hoffimann
Richard Allan Holcombe
Tom Holroyd
David Hoover
Kurt Hornik
Craig Hudson
Christopher Hulbert
Cyril Humbert
John Hunt
Stefan Husmann
Teemu Ikonen
Alan W. Irwin
Allan Jacobs
Marcel Jacobse
Geoff Jacobsen
Vytautas Jančauskas

Andrew Janke
Robert Jenssen
Heikki Junes
Jarkko Kaleva
Lute Kamstra
Christof Kaufmann
Paul Kienzle
Erik Kjellson
Lasse Kliemann
Martin Köhler
Peter Konowski
Endre Kozma
Aravindh Krishnamoorthy
Piotr Krzyzanowski
Tetsuro Kurita
Miroslaw Kwasniak
Claude Lacoursiere
Bill Lash
Friedrich Leisch
Thorsten Liebig
Timo Lindfors
Yu Liu
Sebastien Loisel
Massimo Lorenzin
Hoxide Ma
Jens-Uwe Mager
Alexander Mamonov
Axel Mathéi
Christoph Mayer
Ronald van der Meer
Ed Meyer
Petr Mikulik
Serviscope Minor
Stephen Montgomery-Smith
Kai P. Mueller
Hannes Müller

Iain Murray
Carmen Navarrete
Al Niessner
Takuji Nishimura
Patrick Noffke
Krzesimir Nowak
Peter O'Gorman
Serkan Önder
Luis F. Ortiz

Nicholas R. Jankowski
Cai Jianming
Matthias Jüschke
Avinoam Kalma
Fotios Kasolis
Joel Keay
Lars Kindermann
Arno J. Klaassen
Geoffrey Knauth
Heine Kolltveit
Ken Kouno
Daniel Kraft
Oyvind Kristiansen
Volker Kuhlmann
Ben Kurtz
Rafael Laboissiere
Nithin Lakshmisha
Dirk Laurie
Michael Leitner
Torsten Lilge
Benjamin Lindner
David Livings
Andrej Lojdl
Emil Lucretiu
Colin Macdonald
Stefan Mahr
Ricardo Marranita
Makoto Matsumoto
Laurent Mazet
Markus Meisinger
Thorsten Meyer
Linton Miller
Stefan Monnier
Anthony Morast
Amod Mulay
Victor Munoz

Nicholas Musolino
Todd Neal
Felipe G. Nievinski
Akira Noda
Victor Norton
Michael O'Brien
Thorsten Ohl
Arno Onken
Carl Osterwisch

Mats Jansson
Steven G. Johnson
Atsushi Kajita
Mohamed Kamoun
Thomas Kasper
Mumit Khan
Aaron A. King
Alexander Klein
Hendrik Koerner
Thomas Kolman
Kacper Kowalik
Nir Krakauer
Artem Krosheninnikov
Ilya Kurdyukov
Philipp Kutin
Kai Labusch
Walter Landry
Maurice LeBrun
Johannes Leuschner
Jyh-miin Lin
Ross Lippert
Barbara Locsi
Erik de Castro Lopo
Yi-Hong Lyu
James Macnicol
Rob Mahurin
Orestes Mas
Tatsuro Matsuoka
G. D. McBain
Júlio Hoffimann Mendes
Stefan Miereis
Mike Miller
Rafael Monteiro
Antoine Moreau
Armin Müller
PrasannaKumar Muralidharan
Markus Mützel
Philip Nienhuis
Rick Niles
Kai Noda
Eric Norum
Cillian O'Driscoll
Kai T. Ohlhus
Valentin Ortega-Clavero
Janne Olavi Paanajärvi

Óvári
Jason Alan Palmer
Rolando Pereira
Jim Peterson
Nicholas Piper
Hans Ekkehard Plesser
Nathan Podlich
Jef Poskanzer
Tejaswi D. Prakash
Pooja Rao
Balint Reczey
Lukas Reichlin
Reinhard Resch
Jason Riedy
Petter Risholm
Dmitry Roshchin
Fabio Rossi
David Rörich
Ryan Rusaw
Juhani Saastamoinen
Ben Sapp
Michel D. Schmid
Nicol N. Schraudolph
Ludwig Schwardt
Dmitri A. Sergatskov
Ahsan Ali Shahid
Robert T. Short
John Smith
Peter L. Sondergaard
Quentin H. Spencer
Imad-Eddine Srairi
Russell Standish
Doug Stewart
Judd Storrs
Ivan Sutoris
Ariel Tankus
Matthew Tennyhg
Georg Thimm
Andrew Thornton
Christophe Tournery
Karsten Trulsen
Utkarsh Upadhyay
Peter Van Wieren
Gregory Vanuxem
Sébastien Villemot
Steven Waldrip
Andreas Weber

Scott Pakin
Gabriele Pannocchia
Per Persson
Johannes Pfeifer
Elias Pipping
Sergey Plotnikov
Orion Poplawski
Francesco Potortì
Jarno Rajahalme
James B. Rawlings
Joshua Redstone
Michael Reifenberger
Jens Restemeier
E. Joshua Rigler
Matthew W. Roberts
Peter Rosin
Mark van Rossum
Kevin Ruland
Olli Saarela
Radek Salac
Aleksej Saushev
Julian Schnidder
Sebastian Schubert
Thomas L. Scofield
Vanya Sergeev
Baylis Shanks
Paulo Silva
Julius Smith
Rüdiger Sonderfeld
Christoph Spiel
Andreas Stahel
Ryan Starret
Jen Stewart
Thomas Stuart
John Swensen
Falk Tannhäuser
Remi Thebault
Paul Thomas
Olaf Till
Thomas Treichl
David Turner
José Vallet
James R. Van Zandt
Mihas Varantsou
Marco Vitetta
Thomas Walter
Olaf Weber

José Luis García Pallero
Sylvain Pelissier
Primozz Peterlin
Danilo Piazzalunga
Robert Platt
Tom Poage
Ondrej Popp
Konstantinos Poulios
Eduardo Ramos
Eric S. Raymond
Andy Register
Ernst Reissner
Anthony Richardson
Sander van Rijn
Melvin Robinson
Andrew Ross
Joe Rothweiler
Kristian Rumberg
Toni Saarela
Mike Sander
Alois Schlögl
Sebastian Schöps
Lasse Schuirmann
Daniel J. Sebald
Marko Seric
Andriy Shinkarchuck
Joseph P. Skudlarek
Shan G. Smith
Joerg Specht
David Spies
Richard Stallman
Brett Stewart
Jonathan Stickel
Bernardo Sulzbach
Daisuke Takago
Duncan Temple Lang
Kris Thielemans
Corey Thomasson
Petter Tomner
Abhinav Tripathi
Frederick Umminger
Stefan van der Walt
Risto Vanhanen
Ivana Varekova
Daniel Wagenaar
Jun Wang
Thomas Weber

Rik Wehbring
Bob Weigel
Andreas Weingessel
Martin Weiser
Michael Weitzel
David Wells
Joachim Wiesemann
Alexander Wilms
Joe Winegarden
Georg Wiora
Eddy Xiao
Sahil Yadav
Fook Fah Yap
Sean Young
Michele Zaffalon
Serhiy Zahoriya
Johannes Zarl
Michael Zeising
Federico Zenith
Claudius Zingerli
Alex Zvoleff
Richard Zweig

Special thanks to the following people and organizations for supporting the development of Octave:

- The United States Department of Energy, through grant number DE-FG02-04ER25635.
- Ashok Krishnamurthy, David Hudak, Juan Carlos Chaves, and Stanley C. Ahalt of the Ohio Supercomputer Center.
- The National Science Foundation, through grant numbers CTS-0105360, CTS-9708497, CTS-9311420, CTS-8957123, and CNS-0540147.
- The industrial members of the Texas-Wisconsin Modeling and Control Consortium (TWMCC).
- The Paul A. Elfers Endowed Chair in Chemical Engineering at the University of Wisconsin-Madison.
- Digital Equipment Corporation, for an equipment grant as part of their External Research Program.
- Sun Microsystems, Inc., for an Academic Equipment grant.
- International Business Machines, Inc., for providing equipment as part of a grant to the University of Texas College of Engineering.
- Texaco Chemical Company, for providing funding to continue the development of this software.
- The University of Texas College of Engineering, for providing a Challenge for Excellence Research Supplement, and for providing an Academic Development Funds grant.
- The State of Texas, for providing funding through the Texas Advanced Technology Program under Grant No. 003658-078.
- Noel Bell, Senior Engineer, Texaco Chemical Company, Austin Texas.
- John A. Turner, Group Leader, Continuum Dynamics (CCS-2), Los Alamos National Laboratory, for registering the `https://octave.org` domain name.
- James B. Rawlings, Professor, University of Wisconsin-Madison, Department of Chemical and Biological Engineering.
- Richard Stallman, for writing GNU.

This project would not have been possible without the GNU software used in and to produce Octave.

## Citing Octave in Publications

In view of the many contributions made by numerous developers over many years it is common courtesy to cite Octave in publications when it has been used during the course of

research or the preparation of figures. The `citation` function can automatically generate a recommended citation text for Octave or any of its packages. See the help text below on how to use `citation`.

`citation`
`citation` *package*

Display instructions for citing GNU Octave or its packages in publications.

When called without an argument, display information on how to cite the core GNU Octave system.

When given a package name *package*, display information on citing the specific named package. Note that some packages may not yet have instructions on how to cite them.

The GNU Octave developers and its active community of package authors have invested a lot of time and effort in creating GNU Octave as it is today. Please give credit where credit is due and cite GNU Octave and its packages when you use them.

## How You Can Contribute to Octave

There are a number of ways that you can contribute to help make Octave a better system. Perhaps the most important way to contribute is to write high-quality code for solving new problems, and to make your code freely available for others to use. See https://www.octave.org/get-involved.html for detailed information.

If you find Octave useful, consider providing additional funding to continue its development. Even a modest amount of additional funding could make a significant difference in the amount of time that is available for development and support.

Donations supporting Octave development may be made on the web at https://www.octave.org/donate.

If you cannot provide funding or contribute code, you can still help make Octave better and more reliable by reporting any bugs you find and by offering suggestions for ways to improve Octave. See Appendix D [Trouble], page 1173, for tips on how to write useful bug reports.

## Distribution

Octave is *free* software. This means that everyone is free to use it and free to redistribute it on certain conditions. Octave is not, however, in the public domain. It is copyrighted and there are restrictions on its distribution, but the restrictions are designed to ensure that others will have the same freedom to use and redistribute Octave that you have. The precise conditions can be found in the GNU General Public License that comes with Octave and that also appears in Appendix G [Copying], page 1203.

To download a copy of Octave, please visit https://www.octave.org/download.html.

# 1 A Brief Introduction to Octave

GNU Octave is a high-level language primarily intended for numerical computations. It is typically used for such problems as solving linear and nonlinear equations, numerical linear algebra, statistical analysis, and for performing other numerical experiments. It may also be used as a batch-oriented language for automated data processing.

The current version of Octave executes in a graphical user interface (GUI). The GUI hosts an Integrated Development Environment (IDE) which includes a code editor with syntax highlighting, built-in debugger, documentation browser, as well as the interpreter for the language itself. A command-line interface for Octave is also available.

GNU Octave is freely redistributable software. You may redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation. The GPL is included in this manual, see Appendix G [Copying], page 1203.

This manual provides comprehensive documentation on how to install, run, use, and extend GNU Octave. Additional chapters describe how to report bugs and help contribute code.

This document corresponds to Octave version 11.3.0.

## 1.1 Running Octave

If you installed Octave from an installer program, it will likely have created some icons on your desktop for you to start Octave, either with the graphical user interface (GUI) or a command line interface (CLI). You can also typically find Octave in the "Start Menu" or equivalent of your computer. You can also type `octave` in a command shell; as long as you have Octave in your path, it will start.

If you start Octave with the GUI, the central window is the Octave's own command-line interface (also called a REPL by other programming languages for Read-Evaluate-Print-Loop). In this window Octave displays an initial message and then a prompt like `>>` or `octave:1>` indicating it is ready to accept input. If you have chosen the traditional command-line interface then the command prompt appears in the same window that was running a command shell. In either case, you can immediately begin typing Octave commands.

If you get into trouble, you can usually interrupt Octave by typing *`Control-C`*. Doing this will normally return you to Octave's prompt.

To exit Octave, type *`quit`* or *`exit`* at the Octave prompt.

On systems that support job control, you can suspend Octave by sending it a `SIGTSTP` signal, usually by typing *`Ctrl-z`*.

## 1.2 Simple Examples

The following chapters describe all of Octave's features in detail, but before doing that, it might be helpful to give a sampling of some of its capabilities.

If you are new to Octave, we recommend that you try these examples to begin learning Octave by using it. Lines marked like so, '`octave:13>`' or '`>> `', are lines you type, ending each with a carriage return. (Don't type the text '`octave:13>`' itself! That is only the

Octave prompt, which also looks like `>>` in the GUI. Octave will respond to your commands with an answer, or by displaying a graph.

### 1.2.1 Elementary Calculations

Octave can easily be used for basic numerical calculations. Octave knows about arithmetic operations (+,-,*,/), exponentiation (^), natural logarithms/exponents (log, exp), and the trigonometric functions (sin, cos, ...). Moreover, Octave calculations work on real or imaginary numbers (i,j). In addition, some mathematical constants such as the base of the natural logarithm (e) and the ratio of a circle's circumference to its diameter (pi) are pre-defined.

For example, to verify Euler's Identity,

$$e^{\imath\pi} = -1$$

type the following which will evaluate to `-1` within the tolerance of the calculation.

```
octave:1> exp (i*pi)
```

### 1.2.2 Creating a Matrix

Vectors and matrices are the basic building blocks for numerical analysis. To create a new matrix and store it in a variable so that you can refer to it later, type the command

```
octave:1> A = [ 1, 1, 2; 3, 5, 8; 13, 21, 34 ]
```

Octave will respond by printing the matrix in neatly aligned columns. Octave uses a comma or space to separate entries in a row, and a semicolon or carriage return to separate one row from the next. Ending a command with a semicolon tells Octave not to print the result of the command. For example,

```
octave:2> B = rand (3, 2);
```

will create a 3 row, 2 column matrix with each element set to a random value between zero and one.

To display the value of a variable, simply type the name of the variable at the prompt. For example, to display the value stored in the matrix `B`, type the command

```
octave:3> B
```

### 1.2.3 Matrix Arithmetic

Octave uses standard mathematical notation with the advantage over low-level languages that operators may act on scalars, vector, matrices, or N-dimensional arrays. For example, to multiply the matrix `A` by a scalar value, type the command

```
octave:4> 2 * A
```

To multiply the two matrices `A` and `B`, type the command

```
octave:5> A * B
```

and to form the matrix product $\mathtt{A}^T\mathtt{A}$, type the command

```
octave:6> A' * A
```

### 1.2.4 Solving Systems of Linear Equations

Systems of linear equations are ubiquitous in numerical analysis. To solve the set of linear equations `Ax = b`, use the left division operator, '\':

```
x = A \ b
```

This is conceptually equivalent to $\mathtt{A}^{-1}\mathtt{b}$, but avoids computing the inverse of a matrix directly.

If the coefficient matrix is singular, Octave will print a warning message and compute a minimum norm solution.

A simple example comes from chemistry and the need to obtain balanced chemical equations. Consider the burning of hydrogen and oxygen to produce water.

$$\mathrm{H}_2 + \mathrm{O}_2 \rightarrow \mathrm{H}_2\mathrm{O}$$

The equation above is not accurate. The Law of Conservation of Mass requires that the number of molecules of each type balance on the left- and right-hand sides of the equation. Writing the variable overall reaction with individual equations for hydrogen and oxygen one finds:

$$x_1\mathrm{H}_2 + x_2\mathrm{O}_2 \rightarrow \mathrm{H}_2\mathrm{O}$$

$$\mathrm{H}: \quad 2x_1 + 0x_2 \rightarrow 2$$

$$\mathrm{O}: \quad 0x_1 + 2x_2 \rightarrow 1$$

The solution in Octave is found in just three steps.

```
octave:1> A = [ 2, 0; 0, 2 ];
octave:2> b = [ 2; 1 ];
octave:3> x = A \ b
```

### 1.2.5 Integrating Differential Equations

Octave has built-in functions for solving nonlinear differential equations of the form

$$\frac{dx}{dt} = f(x,t), \qquad x(t = t_0) = x_0$$

For Octave to integrate equations of this form, you must first provide a definition of the function $f(x,t)$. This is straightforward, and may be accomplished by entering the function body directly on the command line. For example, the following commands define the right-hand side function for an interesting pair of nonlinear differential equations. Note that while you are entering a function, Octave responds with a different prompt, to indicate that it is waiting for you to complete your input.

```
octave:1> function xdot = f (x, t)
>
>  r = 0.25;
>  k = 1.4;
>  a = 1.5;
>  b = 0.16;
>  c = 0.9;
>  d = 0.8;
>
>  xdot(1) = r*x(1)*(1 - x(1)/k) - a*x(1)*x(2)/(1 + b*x(1));
>  xdot(2) = c*a*x(1)*x(2)/(1 + b*x(1)) - d*x(2);
>
> endfunction
```

Given the initial condition

```
octave:2> x0 = [1; 2];
```

and the set of output times as a column vector (note that the first output time corresponds to the initial condition given above)

```
octave:3> t = linspace (0, 50, 200)';
```

it is easy to integrate the set of differential equations:

```
octave:4> x = lsode ("f", x0, t);
```

The function `lsode` uses the Livermore Solver for Ordinary Differential Equations, described in A. C. Hindmarsh, "ODEPACK, a Systematized Collection of ODE Solvers", *Scientific Computing*, R. S. Stepleman et al. (Eds.), North-Holland, Amsterdam, 1983, pp. 55–64.

### 1.2.6 Producing Graphical Output

To display the solution of the previous example graphically, use the command

```
octave:1> plot (t, x)
```

Octave will automatically create a separate window to display the plot.

To save a plot once it has been displayed on the screen, use the print command. For example,

```
print foo.pdf
```

will create a file called `foo.pdf` that contains a rendering of the current plot in Portable Document Format. The command

```
help print
```

explains more options for the `print` command and provides a list of additional output file formats.

### 1.2.7 Help and Documentation

Octave has an extensive help facility. The same documentation that is available in printed form is also available from the Octave prompt, because both forms of the documentation are created from the same input file.

In order to get good help you first need to know the name of the command that you want to use. The name of this function may not always be obvious, but a good place to start is to

type `help --list`. This will show you all the operators, keywords, built-in functions, and loadable functions available in the current session of Octave. An alternative is to search the documentation using the `lookfor` function (described in Section 2.3 [Getting Help], page 21).

Once you know the name of the function you wish to use, you can get more help on the function by simply including the name as an argument to help. For example,

```
help plot
```

will display the help text for the `plot` function.

The part of Octave's help facility that allows you to read the complete text of the printed manual from within Octave normally uses a separate program called Info. When you invoke Info you will be put into a menu driven program that contains the entire Octave manual. Help for using Info is provided in this manual, see Section 2.3 [Getting Help], page 21.

### 1.2.8 Editing What You Have Typed

At the Octave prompt, you can recall, edit, and reissue previous commands using Emacs- or vi-style editing commands. The default keybindings use Emacs-style commands. For example, to recall the previous command, press *`Control-p`* (written *`C-p`* for short). Doing this will normally bring back the previous line of input. *`C-n`* will bring up the next line of input, *`C-b`* will move the cursor backward on the line, *`C-f`* will move the cursor forward on the line, etc.

A complete description of the command line editing capability is given in this manual, see Section 2.4 [Command Line Editing], page 26.

## 1.3 Conventions

This section explains the notation conventions that are used in this manual. You may want to skip this section and refer back to it later.

### 1.3.1 Fonts

Examples of Octave code appear in this font or form: `svd (a)`. Names that represent variables or function arguments appear in this font or form: *first-number*. Commands that you type at the shell prompt appear in this font or form: '`octave --no-init-file`'. Commands that you type at the Octave prompt sometimes appear in this font or form: *`foo --bar --baz`*. Specific keys on your keyboard appear in this font or form: `RET`.

### 1.3.2 Evaluation Notation

In the examples in this manual, results from expressions that you evaluate are indicated with '⇒'. For example:

```
sqrt (2)
     ⇒ 1.4142
```

You can read this as "`sqrt (2)` evaluates to 1.4142".

In some cases, matrix values that are returned by expressions are displayed like this

```
[1, 2; 3, 4] == [1, 3; 2, 4]
     ⇒ [ 1, 0; 0, 1 ]
```

and in other cases, they are displayed like this

```
eye (3)
     ⇒  1  0  0
        0  1  0
        0  0  1
```

in order to clearly show the structure of the result.

Sometimes to help describe one expression, another expression is shown that produces identical results. The exact equivalence of expressions is indicated with '≡'. For example:

```
rot90 ([1, 2; 3, 4], -1)
≡
rot90 ([1, 2; 3, 4], 3)
≡
rot90 ([1, 2; 3, 4], 7)
```

### 1.3.3 Printing Notation

Many of the examples in this manual print text when they are evaluated. In this manual the printed text resulting from an example is indicated by '⊣'. The value that is returned by evaluating the expression is displayed with '⇒' (`1` in the next example) and follows on a separate line.

```
printf ("foo %s\n", "bar")
     ⊣ foo bar
     ⇒ 1
```

### 1.3.4 Error Messages

Some examples signal errors. This normally displays an error message on your terminal. Error messages are shown on a line beginning with `error:`.

```
fieldnames ([1, 2; 3, 4])
error: fieldnames: Invalid input argument
```

### 1.3.5 Format of Descriptions

Functions and commands are described in this manual in a uniform format. The first line of a description contains the name of the item followed by its arguments, if any. If there are multiple ways to invoke the function then each allowable form is listed.

The description follows on succeeding lines, sometimes with examples.

#### 1.3.5.1 A Sample Function Description

In a function description, the name of the function being described appears first. It is followed on the same line by a list of parameters. The names used for the parameters are also used in the body of the description.

After all of the calling forms have been enumerated, the next line is a concise one-sentence summary of the function.

After the summary there may be documentation on the inputs and outputs, examples of function usage, notes about the algorithm used, and references to related functions.

Here is a description of an imaginary function `foo`:

`foo (`*x*`)`
`foo (`*x*`, `*y*`)`
`foo (`*x*`, `*y*`, ...)`
Subtract *x* from *y*, then add any remaining arguments to the result.

The input *x* must be a numeric scalar, vector, or array.

The optional input *y* defaults to 19 if it is not supplied.

Example:

```
foo (1, [3, 5], 3, 9)
     ⇒ [ 14, 16 ]
foo (5)
     ⇒ 14
```

More generally,

```
foo (w, x, y, ...)
≡
x - w + y + ...
```

**See also:** bar

Any parameter whose name contains the name of a type (e.g., *integer* or *matrix*) is expected to be of that type. Parameters named *object* may be of any type. Parameters with other sorts of names (e.g., *new_file*) are discussed specifically in the description of the function. In some sections, features common to parameters of several functions are described at the beginning.

### 1.3.5.2 A Sample Command Description

Commands are functions that may be called without surrounding their arguments in parentheses. Command descriptions have a format similar to function descriptions. For example, here is the description for Octave's `diary` command:

`diary`
`diary` *on*
`diary` *off*
`diary` *filename*
`[`*status*`,` *diaryfile*`] = diary`

Record a list of all commands *and* the output they produce, mixed together just as they appear on the terminal.

Valid options are:

on — Start recording a session in a file called `diary` in the current working directory.

off — Stop recording the session in the diary file.

*filename* — Record the session in the file named *filename*.

With no input or output arguments, `diary` toggles the current diary state.

If output arguments are requested, `diary` ignores inputs and returns the current status. The boolean *status* indicates whether recording is on or off, and *diaryfile* is the name of the file where the session is stored.

**See also:** history, evalc.

# 2 Getting Started

This chapter explains some of Octave's basic features, including how to start an Octave session, get help at the command prompt, edit the command line, and write Octave programs that can be executed as commands from your shell.

## 2.1 Invoking Octave from the Command Line

Normally, Octave is used interactively by running the program '`octave`' without any arguments. Once started, Octave reads commands from the terminal until you tell it to exit.

You can also specify the name of a file on the command line, and Octave will read and execute the commands from the named file and then exit when it is finished.

You can further control how Octave starts by using the command-line options described in the next section, and Octave itself can remind you of the options available. Type '`octave --help`' to display all available options and briefly describe their use ('`octave -h`' is a shorter equivalent).

### 2.1.1 Command Line Options

Here is a complete list of the command line options that Octave accepts.

`--built-in-docstrings-file` *`filename`*
: Specify the name of the file containing documentation strings for the built-in functions of Octave. This value is normally correct and should only need to be specified in extraordinary situations.

`--doc-cache-file` *`filename`*
: Specify the name of the documentation cache file to use. The value of *filename* specified on the command line will override any value of `OCTAVE_DOC_CACHE_FILE` found in the environment, but not any commands in the system or user startup files that use the `doc_cache_file` function.

`--echo-commands`
`-x`
: Echo commands as they are executed.

`--eval` *`code`*
`-e` *`code`*
: Evaluate *code* and exit when finished unless `--persist` is also specified.

`--exec-path` *`path`*
: Specify the path to search for programs to run. The value of *path* specified on the command line will override any value of `OCTAVE_EXEC_PATH` found in the environment, but not any commands in the system or user startup files that call the `EXEC_PATH` function.

`--gui`
`-g`
: Start the graphical user interface (GUI).

`--help`
`-h`
: Print short help message and exit.

`--image-path` *path*
: Add path to the head of the search path for images. The value of *path* specified on the command line will override any value of `OCTAVE_IMAGE_PATH` found in the environment, but not any commands in the system or user startup files that call the `IMAGE_PATH` function.

`--info-file` *filename*
: Specify the name of the info file to use. The value of *filename* specified on the command line will override any value of `OCTAVE_INFO_FILE` found in the environment, but not any commands in the system or user startup files that use the `info_file` function.

`--info-program` *program*
: Specify the name of the info program to use. The value of *program* specified on the command line will override any value of `OCTAVE_INFO_PROGRAM` found in the environment, but not any commands in the system or user startup files that use the `info_program` function.

`--init-trace`
: Print the name of each configuration as it is read and executed during initialization.

`--interactive`
`-i`
: Force interactive behavior. This can be useful for running Octave via a remote shell command or inside an Emacs shell buffer.

`--line-editing`
: Force readline use for command-line editing.

`--no-gui`
`-G`
: Disable the graphical user interface (GUI) and use the command line interface (CLI) instead. This is the default behavior, but this option may be useful to override a previous `--gui`.

`--no-history`
`-H`
: Disable recording of command-line history.

`--no-init-all`
`--norc`
`-f`
: Don't read any of the system or user initialization files at startup. This is equivalent to using both of the options `--no-site-file` and `--no-init-user`.

`--no-init-path`
: Don't initialize the search path for function files to include default locations.

`--no-init-site`
: Don't read the site-wide `octaverc` initialization files.

`--no-init-user`
: Don't read the user initialization files `~/.octaverc` and `.octaverc`.

`--no-line-editing`
: Disable command-line editing.

`--no-window-system`
`-W` Disable use of a windowing system including graphics. This forces a strictly terminal-only environment.

`--path` *path*
`-p` *path* Add path to the head of the search path for function files. The value of *path* specified on the command line will override any value of `OCTAVE_PATH` found in the environment, but not any commands in the system or user startup files that set the internal load path through one of the path functions.

`--persist`
Go to interactive mode after `--eval` or reading from a file named on the command line.

`--quiet`
`--silent`
`-q` Don't print the usual greeting and version message at startup.

`--texi-macros-file` *filename*
Specify the name of the file containing Texinfo macros for use by makeinfo.

`--traditional`
`--braindead`
For compatibility with MATLAB, set initial values for user preferences to the following values

```
PS1                           = ">> "
PS2                           = ""
PS4                           = ""
beep_on_error                 = true
confirm_recursive_rmdir       = false
crash_dumps_octave_core       = false
optimize_diagonal_matrix      = false
optimize_permutation_matrix   = false
optimize_range                = false
fixed_point_format            = true
history_timestamp_format_string = "%%-- %D %I:%M %p --%%"
print_struct_array_contents   = true
save_default_options          = "-mat-binary"
struct_levels_to_print        = 0
```

and disable the following warnings

```
Octave:abbreviated-property-match
Octave:colon-nonscalar-argument
Octave:data-file-in-path
Octave:empty-index
Octave:function-name-clash
Octave:possible-matlab-short-circuit-operator
```

Note that this does not enable the `Octave:language-extension` warning, which you might want if you want to be told about writing code that works in Octave but not MATLAB (see [warning], page 263, [warning_ids], page 265).

`--version`
`-v` Print program version information and exit.

*`file`* Execute commands from *file*. Exit when done unless `--persist` is also specified.

Octave also includes several functions which return information about the command line, including the number of arguments and all of the options.

*`args`* `= argv ()`
Return the command line arguments passed to Octave.

For example, if you invoked Octave using the command

```
octave --no-line-editing --quiet
```

`argv` would return a cell array of strings with the elements `--no-line-editing` and `--quiet`.

If you write an executable Octave script, `argv` will return the list of arguments passed to the script. See Section 2.6 [Executable Octave Programs], page 37, for an example of how to create an executable Octave script.

**See also:** [program_name], page 18, [cmdline_options], page 18.

*`opt_struct`* `= cmdline_options ()`
Return a structure containing detailed information about the command line arguments passed to Octave.

Programming Note: This function provides copious amounts of information about Octave's parsing of command line options and may be more useful for debugging Octave rather than for general use.

**See also:** [argv], page 18, [program_name], page 18.

*`name`* `= program_name ()`
Return the filename component of the value returned by `program_invocation_name`.

**See also:** [program_invocation_name], page 18, [argv], page 18.

*`name`* `= program_invocation_name ()`
Return the string that was typed at the shell prompt to run Octave.

The string may include path components as well as the program filename.

If executing a script from the command line (e.g., `octave foo.m`) or using an executable Octave script, the program name is set to the name of the script. See Section 2.6 [Executable Octave Programs], page 37, for an example of how to create an executable Octave script.

**See also:** [program_name], page 18, [argv], page 18.

Here is an example of using these functions to reproduce the command line which invoked Octave.

```
printf ("%s", program_name ());
arg_list = argv ();
for i = 1:nargin
  printf (" %s", arg_list{i});
endfor
printf ("\n");
```

See Section 6.3.3 [Indexing Cell Arrays], page 137, for an explanation of how to retrieve objects from cell arrays, and Section 11.2 [Defining Functions], page 205, for information about the variable `nargin`.

### 2.1.2 Startup Files

When Octave starts, it looks for commands to execute from the files in the following list. These files may contain any valid Octave commands, including function definitions.

*octave-home*`/share/octave/site/m/startup/octaverc`
: where *octave-home* is the directory in which Octave is installed (the default is `/usr/local`). This file is provided so that changes to the default Octave environment can be made globally for all users at your site for all versions of Octave you have installed. Care should be taken when making changes to this file since all users of Octave at your site will be affected. The default file may be overridden by the environment variable `OCTAVE_SITE_INITFILE`.

*octave-home*`/share/octave/`*version*`/m/startup/octaverc`
: where *octave-home* is the directory in which Octave is installed (the default is `/usr/local`), and *version* is the version number of Octave. This file is provided so that changes to the default Octave environment can be made globally for all users of a particular version of Octave. Care should be taken when making changes to this file since all users of Octave at your site will be affected. The default file may be overridden by the environment variable `OCTAVE_VERSION_INITFILE`.

*config-dir*`/octave/octaverc`
: where *config-dir* is the platform-dependent location for user local configuration files (e.g., `$XDG_CONFIG_HOME` on many Unix-like operating systems or `%APPDATA%` on Windows).

`~/.octaverc`
: This file is used to make personal changes to the default Octave environment.

`.octaverc`
: This file can be used to make changes to the default Octave environment for a particular project. Octave searches for this file in the current directory after it reads `~/.octaverc`. Any use of the `cd` command in the `~/.octaverc` file will affect the directory where Octave searches for `.octaverc`.

  If you start Octave in your home directory, commands from the file `~/.octaverc` will only be executed once.

`startup.m`
: This file is used to make personal changes to the default Octave environment. It is executed for MATLAB compatibility, but `~/.octaverc` is the preferred location for configuration changes.

A message will be displayed as each of the startup files is read if you invoke Octave with the `--verbose` option but without the `--quiet` option.

The startup files are always processed in the system's locale charset (independent of the m-file encoding that is set, for example, in the GUI properties). In other words, the

system's locale charset is in effect until a user manually sets the m-file encoding (e.g., in one of the startup files) and triggers re-parsing of any relevant m-files. Octave can be forced to use a new encoding with the function `mfile_encoding`:

```
mfile_encoding ("utf-8");  # set new encoding
clear ("functions");  # re-parse all .m files in the new encoding
```

This changes the encoding that is used to interpret all subsequently run startup and m-files (not including the currently executing file).

## 2.2 Quitting Octave

Shutdown is initiated with the `exit` or `quit` commands (they are equivalent). Similar to startup, Octave has a shutdown process that can be customized by user script files. During shutdown Octave will search for the script file `finish.m` in the function load path. Commands to save all workspace variables or cleanup temporary files may be placed there. Additional functions to execute on shutdown may be registered with `atexit`.

`quit`
`quit` *cancel*
`quit` *force*
`quit` ("*cancel*")
`quit` ("*force*")
`quit` (*status*)
`quit` (*status*, "*force*")
`exit` (...)

Quit the current Octave session.

If the optional integer value *status* is supplied, pass that value to the operating system as Octave's exit status. The default value is zero.

When exiting, Octave will attempt to run the m-file `finish.m` if it exists. User commands to save the workspace or clean up temporary files may be placed in that file. Alternatively, another m-file may be scheduled to run using `atexit`. If an error occurs while executing the `finish.m` file, Octave does not exit and control is returned to the command prompt.

If the optional argument `"cancel"` is provided, Octave does not exit and control is returned to the command prompt. This feature allows the `finish.m` file to cancel the quit process.

If the user preference to request confirmation before exiting, Octave will display a dialog and give the user an option to cancel the exit process.

If the optional argument `"force"` is provided, no confirmation is requested, and the execution of the `finish.m` file is skipped.

Programming Note: `exit` is an alias for `quit` and can be used interchangeably.

**See also:** [atexit], page 20.

`atexit` (*fcn*)
`atexit` (*fcn*, *true*)
`atexit` (*fcn*, *false*)

*status* = atexit (*fcn*, *false*)

Register a function to be called when Octave exits.

For example,

```
function last_words ()
  disp ("Bye bye");
endfunction
atexit ("last_words");
```

will print the message `"Bye bye"` when Octave exits.

The additional argument *flag* will register or unregister *fcn* from the list of functions to be called when Octave exits. If *flag* is true, the function is registered, and if *flag* is false, it is unregistered. For example, after registering the function `last_words` above,

```
atexit ("last_words", false);
```

will remove the function from the list and Octave will not call `last_words` when it exits.

The optional output *status* is only available when unregistering a function. The value is true if the unregistering was successful and false otherwise.

Programming Note: `atexit` only removes the first occurrence of a function from the list; if a function was placed in the list multiple times with `atexit`, it must also be removed from the list multiple times.

**See also:** [quit], page 20.

## 2.3 Commands for Getting Help

The entire text of this manual is available from the Octave prompt via the command *doc*. In addition, the documentation for individual user-written functions and variables is also available via the *help* command. This section describes the commands used for reading the manual and the documentation strings for user-supplied functions and variables. See Section 11.10 [Function Files], page 224, for more information about how to document the functions you write.

help *name*
help *–list*
help .
help
*help_text* = help (...)

Display the help text for *name*.

For example, the command *help help* prints a short message describing the `help` command.

Given the single argument `--list`, list all operators, keywords, built-in functions, and loadable functions available in the current session of Octave.

Given the single argument `.`, list all operators available in the current session of Octave.

If invoked without any arguments, `help` displays instructions on how to access help from the command line.

The help command can provide information about most operators, but *name* must be enclosed by single or double quotes to prevent the Octave interpreter from acting on *name*. For example, `help "+"` displays help on the addition operator.

**See also:** [doc], page 22, [lookfor], page 22, [which], page 156, [info], page 23.

`doc` *`function_name`*
`doc`

Display documentation for the function *function_name* directly from an online version of the printed manual, using the GNU Info browser.

If invoked without an argument, the manual is shown from the beginning.

For example, the command *`doc rand`* starts the GNU Info browser at the `rand` node in the online version of the manual.

Once the GNU Info browser is running, help for using it is available using the command *`C-h`*.

**See also:** [help], page 21.

`lookfor` *`str`*
`lookfor` *`-all`* *`str`*
`[`*`fcn`*`, `*`help1str`*`] = lookfor (`*`str`*`)`
`[`*`fcn`*`, `*`help1str`*`] = lookfor ("`*`-all`*`", `*`str`*`)`

Search for the string *str* in the documentation of all functions in the current function search path.

By default, `lookfor` looks for *str* in just the first sentence of the help string for each function found. The entire help text of each function can be searched by using the `"-all"` argument. All searches are case insensitive.

When called with no output arguments, `lookfor` prints the list of matching functions to the terminal. Otherwise, the output argument *fcns* contains the function names and *help1str* contains the first sentence from the help string of each function.

Programming Note: The ability of `lookfor` to correctly identify the first sentence of the help text is dependent on the format of the function's help. All Octave core functions are correctly formatted, but the same can not be guaranteed for external packages and user-supplied functions. Therefore, the use of the `"-all"` argument may be necessary to find related functions that are not a part of Octave.

The speed of lookup is greatly enhanced by having a cached documentation file. For more information, see [`doc_cache_create`], page 25.

**See also:** [help], page 21, [doc], page 22, [which], page 156, [path], page 229, [doc_cache_create], page 25.

To see what is new in the current release of Octave, use the `news` function.

`news`
`news` *`package`*

Display the current NEWS file for Octave or an installed package.

When called without an argument, display the NEWS file for Octave.

When given a package name *package*, display the current NEWS file for that package.

**See also:** [ver], page 1075, [pkg], page 1084.

`info` ()
: Display contact information for the GNU Octave community.

`warranty` ()
: Describe the conditions for copying and distributing Octave.

The following functions can be used to change which programs are used for displaying the documentation, and where the documentation can be found.

*val* = `info_file` ()
*old_val* = `info_file` (*new_val*)
*old_val* = `info_file` (*new_val*, "*local*")
: Query or set the internal variable that specifies the name of the Octave info file.

  The default value is *octave-home*`/share/info/octave.info`, in which *octave-home* is the root directory of the Octave installation. The default value may be overridden by the environment variable `OCTAVE_INFO_FILE`, or the command line argument `--info-file FNAME`.

  When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

  **See also:** [info_program], page 23, [doc], page 22, [help], page 21, [makeinfo_program], page 23.

*val* = `info_program` ()
*old_val* = `info_program` (*new_val*)
*old_val* = `info_program` (*new_val*, "*local*")
: Query or set the internal variable that specifies the name of the info program to run.

  The default value is `info`. The default value may be overridden by the environment variable `OCTAVE_INFO_PROGRAM`, or the command line argument `--info-program NAME`.

  When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

  **See also:** [info_file], page 23, [doc], page 22, [help], page 21, [makeinfo_program], page 23.

*val* = `makeinfo_program` ()
*old_val* = `makeinfo_program` (*new_val*)
*old_val* = `makeinfo_program` (*new_val*, "*local*")
: Query or set the internal variable that specifies the name of the program that Octave runs to format help text containing Texinfo markup commands.

  The default value is `makeinfo`.

  When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

  **See also:** [texi_macros_file], page 24, [info_file], page 23, [info_program], page 23, [doc], page 22, [help], page 21.

*val* = texi_macros_file ()
*old_val* = texi_macros_file (*new_val*)
*old_val* = texi_macros_file (*new_val*, "*local*")

Query or set the internal variable that specifies the name of the file containing Texinfo macros that are prepended to documentation strings before they are passed to makeinfo.

The default value is *octave-home*/share/octave/*version*/etc/macros.texi, in which *octave-home* is the root directory of the Octave installation, and *version* is the Octave version number. The default value may be overridden by the environment variable OCTAVE_TEXI_MACROS_FILE, or the command line argument --texi-macros-file FNAME.

When called from inside a function with the "local" option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [makeinfo_program], page 23.

*val* = doc_cache_file ()
*old_val* = doc_cache_file (*new_val*)
*old_val* = doc_cache_file (*new_val*, "*local*")

Query or set the internal variable that specifies the name of the Octave documentation cache file.

A cache file significantly improves the performance of the lookfor command. The default value is *octave-home*/share/octave/*version*/etc/doc-cache, in which *octave-home* is the root directory of the Octave installation, and *version* is the Octave version number. The default value may be overridden by the environment variable OCTAVE_DOC_CACHE_FILE, or the command line argument --doc-cache-file FNAME.

When called from inside a function with the "local" option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [doc_cache_create], page 25, [lookfor], page 22, [info_program], page 23, [doc], page 22, [help], page 21, [makeinfo_program], page 23.

**See also:** [lookfor], page 22.

*val* = built_in_docstrings_file ()
*old_val* = built_in_docstrings_file (*new_val*)
*old_val* = built_in_docstrings_file (*new_val*, "*local*")

Query or set the internal variable that specifies the name of the file containing docstrings for built-in Octave functions.

The default value is *octave-home*/share/octave/*version*/etc/built-in-docstrings, in which *octave-home* is the root directory of the Octave installation, and *version* is the Octave version number. The default value may be overridden by the environment variable OCTAVE_BUILT_IN_DOCSTRINGS_FILE, or the command line argument --built-in-docstrings-file FNAME.

Note: This variable is only used when Octave is initializing itself. Modifying it during a running session of Octave will have no effect.

*val* = suppress_verbose_help_message ()
*old_val* = suppress_verbose_help_message (*new_val*)
*old_val* = suppress_verbose_help_message (*new_val*, "*local*")

Query or set the internal variable that controls whether Octave will add additional help information to the end of the output from the `help` command and usage messages for built-in commands.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

The following functions are principally used internally by Octave for generating the documentation. They are documented here for completeness and because they may occasionally be useful for users.

doc_cache_create (*out_file*, *directory*)
doc_cache_create (*out_file*)
doc_cache_create ()

Generate documentation cache for all functions in *directory*.

A documentation cache is generated for all functions in *directory* which may be a single string or a cell array of strings. The cache is used to speed up the function `lookfor`.

The cache is saved in the file *out_file* which defaults to the value `doc-cache` if not given.

If no directory is given (or it is the empty matrix), a cache for built-in functions, operators, and keywords is generated.

**See also:** [doc_cache_file], page 24, [lookfor], page 22, [path], page 229.

[*text*, *format*] = get_help_text (*name*)

Return the raw help text of function *name*.

The raw help text is returned in *text* and the format in *format*. The format is a string which is one of `"texinfo"`, `"html"`, or `"plain text"`.

**See also:** [get_help_text_from_file], page 25.

[*text*, *format*] = get_help_text_from_file (*fname*)

Return the raw help text from the file *fname*.

The raw help text is returned in *text* and the format in *format*. The format is a string which is one of `"texinfo"`, `"html"`, or `"plain text"`.

**See also:** [get_help_text], page 25.

*text* = get_first_help_sentence (*name*)
*text* = get_first_help_sentence (*name*, *max_len*)
[*text*, *status*] = get_first_help_sentence (...)

Return the first sentence of a function's help text.

The first sentence is defined as the text after the function declaration until either the first period (`"."`) or the first appearance of two consecutive newlines (`"\n\n"`). The text is truncated to a maximum length of *max_len*, which defaults to 80. If the text

must be truncated the last three characters of the text are replaced with `"..."` to indicate that more text was available.

The optional output argument *status* returns the status reported by `makeinfo`. If only one output argument is requested, and *status* is nonzero, a warning is displayed.

As an example, the first sentence of this help text is

```
get_first_help_sentence ("get_first_help_sentence")
⊣ ans = Return the first sentence of a function's help text.
```

## 2.4 Command Line Editing

Octave uses the GNU Readline library to provide an extensive set of command-line editing and history features. Only the most common features are described in this manual. In addition, all of the editing functions can be bound to different key strokes at the user's discretion. This manual assumes no changes from the default Emacs bindings. See the GNU Readline Library manual for more information on customizing Readline and for a complete feature list.

To insert printing characters (letters, digits, symbols, etc.), simply type the character. Octave will insert the character at the cursor and advance the cursor forward.

Many of the command-line editing functions operate using control characters. For example, the character *`Control-a`* moves the cursor to the beginning of the line. To type *`C-a`*, hold down `CTRL` and then press `a`. In the following sections, control characters such as *`Control-a`* are written as *`C-a`*.

Another set of command-line editing functions use Meta characters. To type *`M-u`*, hold down the `META` key and press `u`. Depending on the keyboard, the `META` key may be labeled `ALT` or even `WINDOWS`. If your terminal does not have a `META` key, you can still type Meta characters using two-character sequences starting with *`ESC`*. Thus, to enter *`M-u`*, you would type `ESC u`. The *`ESC`* character sequences are also allowed on terminals with real Meta keys. In the following sections, Meta characters such as *`Meta-u`* are written as *`M-u`*.

### 2.4.1 Cursor Motion

The following commands allow you to position the cursor.

*`C-b`* Move back one character.

*`C-f`* Move forward one character.

`BACKSPACE`
Delete the character to the left of the cursor.

`DEL` Delete the character underneath the cursor.

*`C-d`* Delete the character underneath the cursor.

*`M-f`* Move forward a word.

*`M-b`* Move backward a word.

*`C-a`* Move to the start of the line.

*`C-e`* Move to the end of the line.

`C-l` Clear the screen, reprinting the current line at the top.

`C-_`
`C-/` Undo the last action. You can undo all the way back to an empty line.

`M-r` Undo all changes made to this line. This is like typing the 'undo' command enough times to get back to the beginning.

The above table describes the most basic possible keystrokes that you need in order to do editing of the input line. On most terminals, you can also use the left and right arrow keys in place of `C-f` and `C-b` to move forward and backward.

Notice how `C-f` moves forward a character, while `M-f` moves forward a word. It is a loose convention that control keystrokes operate on characters while meta keystrokes operate on words.

The function `clc` will allow you to clear the screen from within Octave programs.

`clc` ()
`home` ()
Clear the terminal screen and move the cursor to the upper left corner.

Programming Note: `home` is an alias for `clc` and can be used interchangeably.

### 2.4.2 Killing and Yanking

*Killing* text means to delete the text from the line, but to save it away for later use, usually by *yanking* it back into the line. If the description for a command says that it 'kills' text, then you can be sure that you can get the text back in a different (or the same) place later.

Here is the list of commands for killing text.

`C-k` Kill the text from the current cursor position to the end of the line.

`M-d` Kill from the cursor to the end of the current word, or if between words, to the end of the next word.

`M-DEL` Kill from the cursor to the start of the previous word, or if between words, to the start of the previous word.

`C-w` Kill from the cursor to the previous whitespace. This is different than `M-DEL` because the word boundaries differ.

And, here is how to *yank* the text back into the line. Yanking means to copy the most-recently-killed text from the kill buffer.

`C-y` Yank the most recently killed text back into the buffer at the cursor.

`M-y` Rotate the kill-ring, and yank the new top. You can only do this if the prior command is `C-y` or `M-y`.

When you use a kill command, the text is saved in a *kill-ring*. Any number of consecutive kills save all of the killed text together, so that when you yank it back, you get it in one clean sweep. The kill ring is not line specific; the text that you killed on a previously typed line is available to be yanked back later, when you are typing another line.

### 2.4.3 Commands for Changing Text

The following commands can be used for entering characters that would otherwise have a special meaning (e.g., `TAB`, *C-q*, etc.), or for quickly correcting typing mistakes.

*C-q*
*C-v* Add the next character that you type to the line verbatim. This is how to insert things like *C-q* for example.

*M-TAB* Insert a tab character.

*C-t* Drag the character before the cursor forward over the character at the cursor, also moving the cursor forward. If the cursor is at the end of the line, then transpose the two characters before it.

*M-t* Drag the word behind the cursor past the word in front of the cursor moving the cursor over that word as well.

*M-u* Uppercase the characters following the cursor to the end of the current (or following) word, moving the cursor to the end of the word.

*M-l* Lowercase the characters following the cursor to the end of the current (or following) word, moving the cursor to the end of the word.

*M-c* Uppercase the character following the cursor (or the beginning of the next word if the cursor is between words), moving the cursor to the end of the word.

### 2.4.4 Commands for Completion

The following commands allow Octave to complete command and variable names for you.

`TAB` Attempt to do completion on the text before the cursor. Octave can complete the names of commands and variables.

*M-?* List the possible completions of the text before the cursor.

*val* = completion_append_char ()
*old_val* = completion_append_char (*new_val*)
*old_val* = completion_append_char (*new_val*, "*local*")

Query or set the internal character variable that is appended to successful command-line completion attempts.

The default value is `" "` (a single space).

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

*completion_list* = completion_matches ("*hint*")

Generate possible word completions for Octave given the character sequence *hint*.

This function is provided for the benefit of programs like Emacs which might be controlling Octave and handling user input. For example:

```
completion_matches ("sine")
⇒
sinetone
sinewave
```

Programming Note: The current command number in Octave is not incremented when this function is called. This is a feature, not a bug.

### 2.4.5 Commands for Manipulating the History

Octave normally keeps track of the commands you type so that you can recall previous commands to edit or execute them again. When you exit Octave, the most recent commands you have typed, up to the number specified by the variable `history_size`, are saved in a file. When Octave starts, it loads an initial list of commands from the file named by the variable `history_file`.

Here are the commands for simple browsing and searching the history list.

LFD
RET
: Accept the current line regardless of where the cursor is. If the line is non-empty, add it to the history list. If the line was a history line, then restore the history line to its original state.

*C-p*
: Move 'up' through the history list.

*C-n*
: Move 'down' through the history list.

*M-<*
: Move to the first line in the history.

*M->*
: Move to the end of the input history, i.e., the line you are entering!

*C-r*
: Search backward starting at the current line and moving 'up' through the history as necessary. This is an incremental search.

*C-s*
: Search forward starting at the current line and moving 'down' through the history as necessary.

On most terminals, you can also use the up and down arrow keys in place of *C-p* and *C-n* to move through the history list.

In addition to the keyboard commands for moving through the history list, Octave provides three functions for viewing, editing, and re-running chunks of commands from the history list.

`history`
`history` *opt1* ...
*H* = `history` ()
*H* = `history` (*opt1*, ...)
: If invoked with no arguments, `history` displays a list of commands that you have executed.

  Valid options are:

  *n*
  -*n*
  : Display only the most recent *n* lines of history.

  -c
  : Clear the history list.

  -q
  : Don't number the displayed lines of history. This is useful for cutting and pasting commands using the X Window System.

`-r` *file* Read the file *file*, appending its contents to the current history list. If the name is omitted, use the default history file (see [history_file], page 32).

`-w` *file* Write the current history to the file *file*. If the name is omitted, use the default history file (see [history_file], page 32).

For example, to display the five most recent commands that you have typed without displaying line numbers, use the command *history -q 5*.

If invoked with a single output argument, the history will be saved to that argument as a cell string and will not be output to screen.

**See also:** [edit_history], page 30, [run_history], page 30, [history_file], page 32, [history_save], page 31.

`edit_history`
`edit_history` *cmd_number*
`edit_history` *first last*

Edit the history list using the editor named by the variable `EDITOR`.

The commands to be edited are first copied to a temporary file. When you exit the editor, Octave executes the commands that remain in the file. It is often more convenient to use `edit_history` to define functions rather than attempting to enter them directly on the command line. The block of commands is executed as soon as you exit the editor. To avoid executing any commands, simply delete all the lines from the buffer before leaving the editor.

When invoked with no arguments, edit the previously executed command; With one argument, edit the specified command *cmd_number*; With two arguments, edit the list of commands between *first* and *last*. Command number specifiers may also be negative where -1 refers to the most recently executed command. The following are equivalent and edit the most recently executed command.

```
edit_history
edit_history -1
```

When using ranges, specifying a larger number for the first command than the last command reverses the list of commands before they are placed in the buffer to be edited.

**See also:** [run_history], page 30, [history], page 29.

`run_history`
`run_history` *cmd_number*
`run_history` *first last*

Run commands from the history list.

When invoked with no arguments, run the previously executed command;

With one argument, run the specified command *cmd_number*;

With two arguments, run the list of commands between *first* and *last*. Command number specifiers may also be negative where -1 refers to the most recently executed command. For example, the command

```
run_history
     OR
run_history -1
```

executes the most recent command again. The command

```
run_history 13 169
```

executes commands 13 through 169.

Specifying a larger number for the first command than the last command reverses the list of commands before executing them. For example:

```
disp (1)
disp (2)
run_history -1 -2
⇒
 2
 1
```

**See also:** [edit_history], page 30, [history], page 29.

Octave also allows you customize the details of when, where, and how history is saved.

*val* = history_save ()
*old_val* = history_save (*new_val*)
*old_val* = history_save (*new_val*, "*local*")

Query or set the internal variable that controls whether commands entered on the command line are saved in the history file.

When called from inside a function with the "local" option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [history_control], page 31, [history_file], page 32, [history_size], page 32, [history_timestamp_format_string], page 32.

*val* = history_control ()
*old_val* = history_control (*new_val*)

Query or set the internal variable that specifies how commands are saved to the history list.

The default value is an empty character string, but may be overridden by the environment variable OCTAVE_HISTCONTROL.

The value of history_control is a colon-separated list of values controlling how commands are saved on the history list. If the list of values includes ignorespace, lines which begin with a space character are not saved in the history list. A value of ignoredups causes lines matching the previous history entry to not be saved. A value of ignoreboth is shorthand for ignorespace and ignoredups. A value of erasedups causes all previous lines matching the current line to be removed from the history list before that line is saved. Any value not in the above list is ignored. If history_control is the empty string, all commands are saved on the history list, subject to the value of history_save.

**See also:** [history_file], page 32, [history_size], page 32, [history_timestamp_format_string], page 32, [history_save], page 31.

*val* = history_file ()
*old_val* = history_file (*new_val*)

Query or set the internal variable that specifies the name of the file used to store command history.

All future commands issued during the current Octave session will be written to this new file (if the current setting of `history_save` allows for this).

The default value is `$DATA/octave/history`, where `$DATA` is the platform-specific location for (roaming) user data files (e.g., `$XDG_DATA_HOME` or, if that is not set, `~/.local/share` on Unix-like operating systems or `%APPDATA%` on Windows). The default value may be overridden by the environment variable `OCTAVE_HISTFILE`.

Programming Notes:

If you want to permanently change the location of Octave's history file you need to issue the `history_file` command in every new Octave session. This can be achieved by using Octave's `.octaverc` startup file.

If you also want to read the saved history commands of past Octave sessions from this different history file, then you need to use the additional command `history -r` after setting the new value of the history file. Example code in Octave's startup file to do this might look like this:

```
history_file ("~/new/.octave_hist");
if (exist (history_file ()))
  history ("-r", history_file());
endif
```

**See also:** [history], page 29, [history_control], page 31, [history_save], page 31, [history_size], page 32, [history_timestamp_format_string], page 32.

*val* = history_size ()
*old_val* = history_size (*new_val*)

Query or set the internal variable that specifies how many entries to store in the history file.

The default value is `1000`, but may be overridden by the environment variable `OCTAVE_HISTSIZE`.

**See also:** [history_file], page 32, [history_timestamp_format_string], page 32, [history_save], page 31.

*val* = history_timestamp_format_string ()
*old_val* = history_timestamp_format_string (*new_val*)
*old_val* = history_timestamp_format_string (*new_val*, "*local*")

Query or set the internal variable that specifies the format string for the comment line that is written to the history file when Octave exits.

The format string is passed to `strftime`. The default value is

```
"# Octave VERSION, %a %b %d %H:%M:%S %Y %Z"
```

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [strftime], page 1027, [history_file], page 32, [history_size], page 32, [history_save], page 31.

*val* = EDITOR ()
*old_val* = EDITOR (*new_val*)
*old_val* = EDITOR (*new_val*, "*local*")

Query or set the internal variable that specifies the default text editor when using the CLI.

The default value is taken from the environment variable `EDITOR` when Octave starts. If the environment variable is not initialized, `EDITOR` will be set to `"emacs"`.

*Note:* This setting applies when running the CLI. When using the Octave GUI the default editor is specified in the Editor tab of Preferences.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [edit], page 225, [edit_history], page 30.

### 2.4.6 Customizing `readline`

Octave uses the GNU Readline library for command-line editing and history features. Readline is very flexible and can be modified through a configuration file of commands (See the GNU Readline library for the exact command syntax). The default configuration file is normally `~/.inputrc`.

Octave provides two commands for initializing Readline and thereby changing the command line behavior.

readline_read_init_file ()
readline_read_init_file (*file*)

Read the readline library initialization file *file*.

If *file* is omitted, read the default initialization file (normally `~/.inputrc`).

See Section "Readline Init File" in *GNU Readline Library*, for details.

**See also:** [readline_re_read_init_file], page 33.

readline_re_read_init_file ()

Re-read the last readline library initialization file that was read.

See Section "Readline Init File" in *GNU Readline Library*, for details.

**See also:** [readline_read_init_file], page 33.

### 2.4.7 Customizing the Prompt

The following variables are available for customizing the appearance of the command-line prompts. Octave allows the prompt to be customized by inserting a number of backslash-escaped special characters that are decoded as follows:

‘`\t`’ The time.

‘`\d`’ The date.

‘`\n`’ Begins a new line by printing the equivalent of a carriage return followed by a line feed.

‘`\s`’ The name of the program (usually just ‘`octave`’).

‘`\w`’ The current working directory.

‘`\W`’ The basename of the current working directory.

‘`\u`’ The username of the current user.

‘`\h`’ The hostname, up to the first ‘.’.

‘`\H`’ The hostname.

‘`\#`’ The command number of this command, counting from when Octave starts.

‘`\!`’ The history number of this command. This differs from ‘`\#`’ by the number of commands in the history list when Octave starts.

‘`\$`’ If the effective UID is 0, a ‘`#`’, otherwise a ‘`$`’.

‘`\nnn`’ The character whose character code in octal is *nnn*.

‘`\\`’ A backslash.

*val* = PS1 ()
*old_val* = PS1 (*new_val*)
*old_val* = PS1 (*new_val*, "*local*")

Query or set the primary prompt string.

When executing interactively, Octave displays the primary prompt when it is ready to read a command.

The default value of the primary prompt string is `'octave:\#> '`. To change it, use a command like

```
PS1 ('\u@\H> ')
```

which will result in the prompt ‘`boris@kremvax> `’ for the user ‘`boris`’ logged in on the host ‘`kremvax.kgb.su`’. Note that two backslashes are required to enter a backslash into a double-quoted character string. See Chapter 5 [Strings], page 75.

You can also use ANSI escape sequences if your terminal supports them. This can be useful for coloring the prompt. For example,

```
PS1 ('\[\033[01;31m\]\s:\#> \[\033[0m\]')
```

will give the default Octave prompt a red coloring.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [PS2], page 34, [PS4], page 35.

*val* = PS2 ()
*old_val* = PS2 (*new_val*)
*old_val* = PS2 (*new_val*, "*local*")

Query or set the secondary prompt string.

The secondary prompt is printed when Octave is expecting additional input to complete a command. For example, if you are typing a `for` loop that spans several lines, Octave will print the secondary prompt at the beginning of each line after the first. The default value of the secondary prompt string is `"> "`.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [PS1], page 34, [PS4], page 35.

*val* = PS4 ()
*old_val* = PS4 (*new_val*)
*old_val* = PS4 (*new_val*, "*local*")

Query or set the character string used to prefix output produced when echoing commands is enabled.

The default value is `"+ "`. See Section 2.4.8 [Diary and Echo Commands], page 35, for a description of echoing commands.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [echo], page 36, [PS1], page 34, [PS2], page 34.

### 2.4.8 Diary and Echo Commands

Octave's diary feature allows you to keep a log of all or part of an interactive session by recording the input you type and the output that Octave produces in a separate file.

diary
diary *on*
diary *off*
diary *filename*
[*status*, *diaryfile*] = diary

Record a list of all commands *and* the output they produce, mixed together just as they appear on the terminal.

Valid options are:

on — Start or resume recording a session. If no *filename* has previously been specified recording will take place in a file called `diary` in the current working directory.

off — Stop recording the session in the diary file.

*filename* — Start or resume recording the session in the file named *filename*.

With no input or output arguments, `diary` toggles the current diary state.

If output arguments are requested, `diary` ignores inputs and returns the current status. The boolean *status* indicates whether recording is on or off, and *diaryfile* is the name of the file where the session is stored.

**See also:** [history], page 29, [evalc], page 187.

Sometimes it is useful to see the commands in a function or script as they are being evaluated. This can be especially helpful for debugging some kinds of problems.

`echo`
`echo` *on*
`echo` *off*
`echo` *on all*
`echo` *off all*
`echo` *function on*
`echo` *function off*

Control whether commands are displayed as they are executed.

Valid options are:

`on` Enable echoing of commands as they are executed in script files.

`off` Disable echoing of commands as they are executed in script and function files.

`on all` Enable echoing of commands as they are executed in script files and functions.

`off all` Disable echoing of commands as they are executed in script files and functions.

*function* `on`
Enable echoing of commands as they are executed in the named function.

*function* `off`
Disable echoing of commands as they are executed in the named function.

With no arguments, `echo` toggles the current echo state.

Programming Note: Echoing all commands can be a simple way to debug an easy coding problem. However, the amount of output can grow quite quickly. For more difficult problems the built-in debugger (`help debug`) is more useful.

**See also:** [PS4], page 35.

## 2.5 How Octave Reports Errors

Octave reports two kinds of errors for invalid programs.

A *parse error* occurs if Octave cannot understand something you have typed. For example, if you misspell a keyword,

```
octave:13> function z = f (x, y) z = x ||| 2; endfunction
```

Octave will respond immediately with a message like this:

```
parse error:

  syntax error

>>> function z = f (x, y) z = x ||| y; endfunction
                                 ^
```

For most parse errors, Octave uses a caret (‘^’) to mark the point on the line where it was unable to make sense of your input. In this case, Octave generated an error message because the keyword for the logical or operator (`||`) was misspelled. It marked the error at the third ‘`|`’ because the code leading up to this was correct but the final ‘`|`’ was not understood.

Another class of error message occurs at evaluation time. These errors are called *run-time errors*, or sometimes *evaluation errors*, because they occur when your program is being *run*, or *evaluated*. For example, if after correcting the mistake in the previous function definition, you type

```
octave:13> f ()
```

Octave will respond with

```
error: ‘x' undefined near line 1 column 24
error: called from:
error:   f at line 1, column 22
```

This error message has several parts, and gives quite a bit of information to help you locate the source of the error. The messages are generated from the point of the innermost error, and provide a traceback of enclosing expressions and function calls.

In the example above, the first line indicates that a variable named ‘`x`’ was found to be undefined near line 1 and column 24 of some function or expression. For errors occurring within functions, lines are counted from the beginning of the file containing the function definition. For errors occurring outside of an enclosing function, the line number indicates the input line number, which is usually displayed in the primary prompt string.

The second and third lines of the error message indicate that the error occurred within the function `f`. If the function `f` had been called from within another function, for example, `g`, the list of errors would have ended with one more line:

```
error:   g at line 1, column 17
```

These lists of function calls make it fairly easy to trace the path your program took before the error occurred, and to correct the error before trying again.

## 2.6 Executable Octave Programs

Once you have learned Octave, you may want to write self-contained Octave scripts, using the ‘`#!`’ script mechanism. You can do this on GNU systems and on many Unix systems[1].

Self-contained Octave scripts are useful when you want to write a program which users can invoke without knowing that the program is written in the Octave language. Octave scripts are also used for batch processing of data files. Once an algorithm has been developed and tested in the interactive portion of Octave, it can be committed to an executable script and used again and again on new data files.

As a trivial example of an executable Octave script, you might create a text file named `hello`, containing the following lines:

```
#! octave-interpreter-name -qf
# a sample Octave program
printf ("Hello, world!\n");
```

[1] The ‘`#!`’ mechanism works on Unix systems derived from Berkeley Unix, System V Release 4, and some System V Release 3 systems.

(where *octave-interpreter-name* should be replaced with the full path and name of your Octave binary). Note that this will only work if '`#!`' appears at the very beginning of the file. After making the file executable (with the `chmod` command on Unix systems), you can simply type:

```
hello
```

at the shell, and the system will arrange to run Octave as if you had typed:

```
octave hello
```

The line beginning with '`#!`' lists the full path and filename of an interpreter to be run, and an optional initial command line argument to pass to that interpreter. The operating system then runs the interpreter with the given argument and the full argument list of the executed program. The first argument in the list is the full filename of the Octave executable. The rest of the argument list will either be options to Octave, or data files, or both. The '`-qf`' options are usually specified in stand-alone Octave programs to prevent them from printing the normal startup message, and to keep them from behaving differently depending on the contents of a particular user's `~/.octaverc` file. See Section 2.1 [Invoking Octave from the Command Line], page 15.

Note that some operating systems may place a limit on the number of characters that are recognized after '`#!`'. Also, the arguments appearing in a '`#!`' line are parsed differently by various shells/systems. The majority of them group all the arguments together in one string and pass it to the interpreter as a single argument. In this case, the following script:

```
#! octave-interpreter-name -q -f # comment
```

is equivalent to typing at the command line:

```
octave "-q -f # comment"
```

which will produce an error message. Unfortunately, it is not possible for Octave to determine whether it has been called from the command line or from a '`#!`' script, so some care is needed when using the '`#!`' mechanism.

### 2.6.1 Passing Arguments to Executable Scripts

Note that when Octave is started from an executable script, the built-in function `argv` returns a cell array containing the command line arguments passed to the executable Octave script, not the arguments passed to the Octave interpreter on the '`#!`' line of the script. For example, the following program will reproduce the command line that was used to execute the script, not '`-qf`'.

```
#! /bin/octave -qf
printf ("%s", program_name ());
arg_list = argv ();
for i = 1:nargin
  printf (" %s", arg_list{i});
endfor
printf ("\n");
```

### 2.6.2 Dual-Purpose Executable Scripts and Octave Functions

To write m-files that can act as executable programs when called from the shell or as normal functions when called from within Octave, use default input arguments initialized with the `argv` function.

If a function is called from the shell Octave will not pass any input parameters to the function and therefore the default argument is used. But when a function is called from the interpreter any arguments *are* passed to the function and these override the default.

Additionally, the file must end with the extension `.m` so that the interpreter will recognize it as an Octave function. Finally, the output from `argv` is a cell array of strings. It may be necessary to convert this to a numeric value with `str2double` or `str2num` before processing.

As a complete example, consider the following code located in the file `mysin.m`.

```
#! /bin/octave -qf
function retval = mysin (x = str2double (argv(){end}))
  retval = sin (x)
endfunction
```

This can be called from the shell with

```
mysin.m 1.5
```

or from Octave with

```
mysin (1.5)
```

## 2.7 Comments in Octave Programs

A *comment* is some text that is included in a program for the sake of human readers, and which is NOT an executable part of the program. Comments can explain what the program does, and how it works. Nearly all programming languages have provisions for comments, because programs are typically hard to understand without them.

### 2.7.1 Single Line Comments

In the Octave language, a comment starts with either the sharp sign character, '`#`', or the percent symbol '`%`' and continues to the end of the line. Any text following the sharp sign or percent symbol is ignored by the Octave interpreter and not executed. The following example shows whole line and partial line comments.

```
function countdown
  # Count down for main rocket engines
  disp (3);
  disp (2);
  disp (1);
  disp ("Blast Off!");  # Rocket leaves pad
endfunction
```

### 2.7.2 Block Comments

Entire blocks of code can be commented by enclosing the code between matching '`#{`' and '`#}`' or '`%{`' and '`%}`' markers. For example,

```
function quick_countdown
  # Count down for main rocket engines
  disp (3);
 #{
  disp (2);
  disp (1);
 #}
  disp ("Blast Off!");  # Rocket leaves pad
endfunction
```

will produce a very quick countdown from `'3'` to `"Blast Off"` as the lines `"disp (2);"` and `"disp (1);"` won't be executed.

The block comment markers must appear alone as the only characters on a line (excepting whitespace) in order to be parsed correctly.

### 2.7.3 Comments and the Help System

The `help` command (see Section 2.3 [Getting Help], page 21) is able to find the first block of comments in a function and return those as a documentation string. This means that the same commands used to get help on built-in functions are available for properly formatted user-defined functions. For example, after defining the function `f` below,

```
function xdot = f (x, t)

# usage: f (x, t)
#
# This function defines the right-hand
# side functions for a set of nonlinear
# differential equations.

  r = 0.25;
  ...
endfunction
```

the command *`help f`* produces the output

```
usage: f (x, t)

This function defines the right-hand
side functions for a set of nonlinear
differential equations.
```

Although it is possible to put comment lines into keyboard-composed, throw-away Octave programs, it usually isn't very useful because the purpose of a comment is to help you or another person understand the program at a later time.

The `help` parser currently only recognizes single line comments (see Section 2.7.1 [Single Line Comments], page 39) and not block comments for the initial help text.

# 3 Data Types

All versions of Octave include a number of built-in data types, including real and complex scalars and matrices, character strings, a data structure type, and an array that can contain all data types.

It is also possible to define new specialized data types by writing a small amount of C++ code. On some systems, new data types can be loaded dynamically while Octave is running, so it is not necessary to recompile all of Octave just to add a new type. See Appendix A [External Code Interface], page 1097, for more information about Octave's dynamic linking capabilities. Section 3.2 [User-defined Data Types], page 46, describes what you must do to define a new data type for Octave.

*typestr* = typeinfo (*expr*)
*cstr* = typeinfo ()

Return the type of the expression *expr*, as a string.

If *expr* is omitted, return a cell array of strings containing all the currently installed data types.

**See also:** [class], page 41, [isa], page 41.

## 3.1 Built-in Data Types

The standard built-in data types are real and complex scalars and matrices, ranges, character strings, a data structure type, and cell arrays. Additional built-in data types may be added in future versions. If you need a specialized data type that is not currently provided as a built-in type, you are encouraged to write your own user-defined data type and contribute it for distribution in a future release of Octave.

The data type of a variable can be determined and changed through the use of the following functions.

*classname* = class (*obj*)
*cls* = class (*s*, *classname*)
*cls* = class (*s*, *classname*, *parent1*, ...)

Return the class of the object *obj*, or create a class with fields from structure *s* and name (string) *classname*.

Additional arguments name a list of parent classes from which the new class is derived.

**See also:** [typeinfo], page 41, [isa], page 41.

*tf* = isa (*obj*, *classname*)

Return true if *obj* is an object from the class *classname*.

*classname* may also be one of the following class categories:

`"float"` Floating point value comprising classes `"double"` and `"single"`.

`"integer"`
Integer value comprising classes (u)int8, (u)int16, (u)int32, (u)int64.

`"numeric"`
Numeric value comprising either a floating point or integer value.

If *classname* is a cell array of string, a logical array of the same size is returned, containing true for each class to which *obj* belongs to.

**See also:** [class], page 41, [typeinfo], page 41.

`y = cast (x, "type")`
`y = cast (x, "like", var)`

Convert *x* to data type *type*.

The input *x* may be a scalar, vector, or matrix of a class that is convertible to the target class (see below).

If a variable *var* is specified after `"like"`, *x* is converted to the same data type and sparsity attribute. If *var* is complex, *x* will be complex, too.

*var* may be and *type* may name any of the following built-in numeric classes:

```
"double"
"single"
"logical"
"char"
"int8"
"int16"
"int32"
"int64"
"uint8"
"uint16"
"uint32"
"uint64"
```

The value *x* may be modified to fit within the range of the new type.

Examples:

```
cast (-5, "uint8")
   ⇒ 0
cast (300, "int8")
   ⇒ 127
```

Programming Note: This function relies on the object *x* having a conversion method named *type*. User-defined classes may implement only a subset of the full list of types shown above. In that case, it may be necessary to call cast twice in order to reach the desired type. For example, the conversion to double is nearly always implemented, but the conversion to uint8 might not be. In that case, the following code will work:

```
cast (cast (user_defined_val, "double"), "uint8")
```

**See also:** [typecast], page 42, [int8], page 59, [uint8], page 60, [int16], page 60, [uint16], page 60, [int32], page 60, [uint32], page 60, [int64], page 60, [uint64], page 60, [double], page 51, [single], page 58, [logical], page 66, [char], page 83, [class], page 41, [typeinfo], page 41.

`y = typecast (x, "class")`

Return a new array *y* resulting from interpreting the data of *x* in memory as data of the numeric class *class*.

Both the class of *x* and *class* must be one of the built-in numeric classes:

```
"logical"
"char"
"int8"
"int16"
"int32"
"int64"
"uint8"
"uint16"
"uint32"
"uint64"
"double"
"single"
"double complex"
"single complex"
```

the last two are only used with *class*; they indicate that a complex-valued result is requested. Complex arrays are stored in memory as consecutive pairs of real numbers. The sizes of integer types are given by their bit counts. Both logical and char are typically one byte wide; however, this is not guaranteed by C++. If your system is IEEE 754 conformant, single and double will be 4 bytes and 8 bytes wide, respectively. `"logical"` is not allowed for *class*.

If the input is a row vector, the return value is a row vector, otherwise it is a column vector.

If the bit length of *x* is not divisible by that of *class*, an error occurs.

An example of the use of typecast on a little-endian machine is

```
x = uint16 ([1, 65535]);
typecast (x, "uint8")
⇒ [   1,   0, 255, 255]
```

**See also:** [cast], page 42, [bitpack], page 43, [bitunpack], page 44, [swapbytes], page 43.

`y = swapbytes (x)`

Swap the byte order on values, converting from little endian to big endian and vice versa.

For example:

```
swapbytes (uint16 (1:4))
⇒   256   512   768  1024
```

**See also:** [typecast], page 42, [cast], page 42.

`y = bitpack (x, class)`

Return a new array *y* resulting from interpreting the logical array *x* as raw bit patterns for data of the numeric class *class*.

*class* must be one of the built-in numeric classes:

```
"double"
"single"
"double complex"
"single complex"
"char"
"int8"
"int16"
"int32"
"int64"
"uint8"
"uint16"
"uint32"
"uint64"
```

The number of elements of *x* should be divisible by the bit length of *class*. If it is not, excess bits are discarded. Bits come in increasing order of significance, i.e., `x(1)` is bit 0, `x(2)` is bit 1, etc.

The result is a row vector if *x* is a row vector, otherwise it is a column vector.

**See also:** [bitunpack], page 44, [typecast], page 42.

`y = bitunpack (x)`

Return a logical array *y* corresponding to the raw bit patterns of *x*.

*x* must belong to one of the built-in numeric classes:

```
"double"
"single"
"char"
"int8"
"int16"
"int32"
"int64"
"uint8"
"uint16"
"uint32"
"uint64"
```

The result is a row vector if *x* is a row vector; otherwise, it is a column vector.

**See also:** [bitpack], page 43, [typecast], page 42.

### 3.1.1 Numeric Objects

Octave's built-in numeric objects include real, complex, and integer scalars and matrices. All built-in floating point numeric data is currently stored as double precision numbers. On systems that use IEEE 754 floating point format, values in the range of approximately $2.2251 \times 10^{-308}$ to $1.7977 \times 10^{308}$ can be stored, and the relative precision is approximately $2.2204 \times 10^{-16}$. The exact values are given by the variables `realmin`, `realmax`, and `eps`, respectively.

Matrix objects can be of any size, and can be dynamically reshaped and resized. It is easy to extract individual rows, columns, or submatrices using a variety of powerful indexing features. See Section 8.1 [Index Expressions], page 159.

See Chapter 4 [Numeric Data Types], page 51, for more information.

### 3.1.2 Missing Data

It is possible to represent missing data explicitly in Octave using NA (short for "Not Available"). This is helpful in distinguishing between a property of the data (i.e., some of it was not recorded) and calculations on the data which generated an error (i.e., created NaN values). In short, if you do not get the result you expect is it your data or your algorithm?

The missing data marker is a special case of the representation of NaN. Because of that, it can only be used with data represented by floating point numbers—no integer, logical, or char values.

In general, use NA and the test `isna`, to describe the dataset or to reduce the dataset to only valid entries. Numerical calculations with NA will generally `"poison"` the results and conclude with an output NA. However, this can not be guaranteed on all platforms and NA may be replaced by NaN.

Example 1 : Describing the dataset

```
data = [1, NA, 3];
percent_missing = 100 * sum (isna (data(:))) / numel (data);
printf ('%2.0f%% of the dataset is missing\n', percent_missing);
 ⊣ 33% of the dataset is missing
```

Example 2 : Restrict calculations to valid data

```
raw_data = [1, NA, 3];
printf ('mean of raw data is %.1f\n', mean (raw_data));
 ⊣ mean of raw data is NA
valid_data = raw_data (! isna (raw_data));
printf ('mean of valid data is %.1f\n', mean (valid_data));
 ⊣ mean of valid data is 2.0
```

`x = NA`
`x = NA (`*n*`)`
`x = NA (`*m*`, `*n*`, ...)`
`x = NA ([`*m*`, `*n*`, ...])`
`x = NA (..., `*class*`)`
`x = NA (..., "`*like*`", `*var*`)`

Return a scalar, matrix, or N-dimensional array whose elements are all equal to the special constant NA (Not Available) used to designate missing values.

Note that `NA` always compares not equal to `NA` (`NA != NA`). To find `NA` values, use the `isna` function.

If called with no arguments, return the scalar value `NA`.

If invoked with a single scalar integer argument *n*, return a square NxN matrix.

If invoked with two or more scalar integer arguments, or a vector of integer values, return an array with the given dimensions.

The optional argument *class* specifies the class of the return array. The only valid options are `"double"` (default) or `"single"`.

Programming Note: The missing data marker `NA` is a special case of the representation of `NaN`. Numerical calculations with `NA` will generally "poison" the results and conclude with an output of `NA`. However, this can not be guaranteed on all platforms and `NA` may be replaced by `NaN`. See Section 3.1.2 [Missing Data], page 45.

**See also:** [isna], page 46.

*tf* = `isna (x)`

Return a logical array which is true where the elements of $x$ are NA (missing) values and false where they are not.

For example:

```
isna ([13, Inf, NA, NaN])
     ⇒ [ 0, 0, 1, 0 ]
```

**See also:** [isnan], page 567, [isinf], page 567, [isfinite], page 567.

### 3.1.3 String Objects

A character string in Octave consists of a sequence of characters enclosed in either double-quote or single-quote marks. Internally, Octave currently stores strings as matrices of characters. All the indexing operations that work for matrix objects also work for strings.

See Chapter 5 [Strings], page 75, for more information.

### 3.1.4 Data Structure Objects

Octave's data structure type can help you to organize related objects of different types. The current implementation uses an associative array with indices limited to strings, but the syntax is more like C-style structures.

See Section 6.1 [Structures], page 117, for more information.

### 3.1.5 Cell Array Objects

A Cell Array in Octave is general array that can hold any number of different data types.

See Section 6.3 [Cell Arrays], page 131, for more information.

## 3.2 User-defined Data Types

Someday I hope to expand this to include a complete description of Octave's mechanism for managing user-defined data types. Until this feature is documented here, you will have to make do by reading the code in the `ov.h`, `ops.h`, and related files from Octave's `src` directory.

## 3.3 Object Sizes

The following functions allow you to determine the size of a variable or expression. These functions are defined for all objects. They return $-1$ when the operation doesn't make sense. For example, Octave's data structure type doesn't have rows or columns, so the `rows` and `columns` functions return $-1$ for structure arguments.

*n* = `ndims (A)`

Return the number of dimensions of $A$.

For any array, the result will always be greater than or equal to 2. Trailing singleton dimensions are not counted, i.e., trailing dimensions $d$ greater than 2 for which `size (A, d) = 1`.

```
ndims (ones (4, 1, 2, 1))
    ⇒ 3
```

**See also:** [size], page 48.

*nc* = `columns (A)`
*nc* = `width (A)`

Return the number of columns of $A$.

This is equivalent to `size (A, 2)`.

Programming Note: `width` is an alias for `columns` and can be used interchangeably.

**See also:** [rows], page 47, [size], page 48, [length], page 48, [numel], page 47, [isscalar], page 70, [isvector], page 70, [ismatrix], page 69.

*nr* = `rows (A)`
*nr* = `height (A)`

Return the number of rows of $A$.

This is equivalent to `size (A, 1)`.

Programming Note: `height` is an alias for `rows` and can be used interchangeably.

**See also:** [columns], page 47, [size], page 48, [length], page 48, [numel], page 47, [isscalar], page 70, [isvector], page 70, [ismatrix], page 69.

*n* = `numel (A)`
*n* = `numel (A, idx1, idx2, ...)`

Return the number of elements in the object $A$.

Optionally, if indices *idx1*, *idx2*, ... are supplied, return the number of elements that would result from the indexing

```
A(idx1, idx2, ...)
```

Note that the indices do not have to be scalar numbers. For example,

```
a = 1;
b = ones (2, 3);
numel (a, b)
```

will return 6, as this is the number of ways to index with $b$. Or the index could be the string `":"` which represents the colon operator. For example,

```
A = ones (5, 3);
numel (A, 2, ":")
```

will return 3 as the second row has three column entries.

This method is also called when an object appears as lvalue with cs-list indexing, i.e., `object{...}` or `object(...).field`.

**See also:** [size], page 48, [length], page 48, [ndims], page 46.

*n* = `length (`*A*`)`

Return the length of the object *A*.

The length is 0 for empty objects, 1 for scalars, and the number of elements for vectors. For matrix or N-dimensional objects, the length is the number of elements along the largest dimension (equivalent to `max (size (A))`).

**See also:** [numel], page 47, [size], page 48.

*sz* = `size (`*A*`)`
*dim_sz* = `size (`*A*`,` *dim*`)`
*dim_sz* = `size (`*A*`,` *d1*`,` *d2*`, ...)`
[*rows*`,` *cols*`, ...,` *dim_N_sz*] = `size (...)`

Return a row vector with the size (number of elements) of each dimension for the object *A*.

When given a second argument, *dim*, return the size of the corresponding dimension. If *dim* is a vector, return each of the corresponding dimensions. Multiple dimensions may also be specified as separate arguments.

With a single output argument, `size` returns a row vector. When called with multiple output arguments, `size` returns the size of dimension N in the Nth argument. The number of rows, dimension 1, is returned in the first argument, the number of columns, dimension 2, is returned in the second argument, etc. If there are more dimensions in *A* than there are output arguments, `size` returns the total number of elements in the remaining dimensions in the final output argument. If the requested dimension *dim* is greater than the number of dimensions in *A*, `size` returns 1 (not 0).

Example 1: single row vector output

```
size ([1, 2; 3, 4; 5, 6])
   ⇒ [ 3, 2 ]
```

Example 2: number of elements in 2nd dimension (columns)

```
size ([1, 2; 3, 4; 5, 6], 2)
    ⇒ 2
```

Example 3: number of output arguments == number of dimensions

```
[nr, nc] = size ([1, 2; 3, 4; 5, 6])
    ⇒ nr = 3
    ⇒ nc = 2
```

Example 4: number of output arguments < number of dimensions

```
[nr, remainder] = size (ones (2, 3, 4, 5))
    ⇒ nr = 2
    ⇒ remainder = 60
```

Example 5: number of elements in dimension > number of actual dimensions

```
sz4 = size (ones (2, 3), 4)
    ⇒ sz4 = 1
```

**See also:** [numel], page 47, [ndims], page 46, [length], page 48, [rows], page 47, [columns], page 47, [size_equal], page 49, [common_size], page 567.

*tf* = **isempty** (*A*)
: Return true if *A* is an empty object (any one of its dimensions is zero).

  **See also:** [isnull], page 49, [isa], page 41.

*tf* = **isnull** (*x*)
: Return true if *x* is a special null array, string, or single quoted string.

  Indexed assignment with such a null value on the right-hand side should delete array elements. This function is used in place of `isempty` when overloading the indexed assignment method (`subsasgn`) for user-defined classes. `isnull` is used to distinguish between these two cases:

  `A(I) = []`

  and

  `X = []; A(I) = X`

  In the first assignment, the right-hand side is `[]` which is a special null value. As long as the index *I* is not empty, this code should delete elements from *A* rather than perform assignment.

  In the second assignment, the right-hand side is empty (because *X* is `[]`), but it is **not** null. This code should assign the empty value to elements in *A*.

  An example from Octave's built-in char class demonstrates the interpreter behavior when `isnull` is used correctly.

```
str = "Hello World";
nm = "Wally";
str(7:end) = nm                 # indexed assignment
  ⇒ str = Hello Wally
str(7:end) = ""                 # indexed deletion
  ⇒ str = Hello
```

  **See also:** [isempty], page 49, [isindex], page 164.

*sz* = **sizeof** (*val*)
: Return the size of *val* in bytes.

  **See also:** [whos], page 151.

*TF* = **size_equal** (*A*, *B*)
*TF* = **size_equal** (*A*, *B*, ...)
: Return true if the dimensions of all arguments agree.

  Trailing singleton dimensions are ignored. When called with a single argument, or no argument, `size_equal` returns true.

  **See also:** [size], page 48, [numel], page 47, [ndims], page 46, [common_size], page 567.

*B* = **squeeze** (*A*)
: Remove singleton dimensions from *A* and return the result.

  Note that for compatibility with MATLAB, all objects have a minimum of two dimensions and row vectors are left unchanged.

  **See also:** [reshape], page 573.

# 4 Numeric Data Types

A *numeric constant* may be a scalar, a vector, or a matrix, and it may contain complex values.

The simplest form of a numeric constant, a scalar, is a single number. Note that by default numbers are represented within Octave by IEEE 754 double precision (binary64) floating-point format (complex constants are stored as pairs of binary64 values). It is, however, possible to represent true integers as described in Section 4.4 [Integer Data Types], page 59.

If the numeric constant is a real integer, it can be defined in decimal, hexadecimal, or binary notation. An important difference, however, is that decimal constants will be stored as binary64 values while using hexadecimal or binary notation will result in a true integer with a storage class just large enough to hold the specified number. Hexadecimal notation starts with '`0x`' or '`0X`', binary notation starts with '`0b`' or '`0B`', otherwise decimal notation is assumed. As a consequence, '`0b`' is not a hexadecimal number, in fact, it is not a valid number at all.

For better readability, digits may be partitioned by the underscore separator '`_`', which is ignored by the Octave interpreter. Here are some examples of real-valued integer constants, which all represent the same value:

```
42            # decimal notation, binary64
0x2A          # hexadecimal notation
0b101010      # binary notation
0b10_1010     # underscore notation
round (42.1)  # also binary64
```

In decimal notation, the numeric constant may be denoted as decimal fraction or even in scientific (exponential) notation. Note that this is not possible for hexadecimal or binary notation. Again, in the following example all numeric constants represent the same value:

```
.105
1.05e-1
.00105e+2
```

Unlike most programming languages, complex numeric constants are denoted as the sum of real and imaginary parts. The imaginary part is denoted by a real-valued numeric constant followed immediately by a complex value indicator ('`i`', '`j`', '`I`', or '`J`' which represents $\sqrt{-1}$). No spaces are allowed between the numeric constant and the complex value indicator. All complex values are stored as pairs of binary64 values and the use of hexadecimal or binary notation does *not* result in a true integer. Some examples of complex numeric constants that all represent the same value:

```
3 + 42i
3 + 42j
3 + 42I
3 + 42J
3.0 + 42.0i
3.0 + 0x2Ai
3.0 + 0b10_1010i
0.3e1 + 420e-1i
```

`y = double (x)`
: Convert $x$ to double precision type.

  **See also:** [single], page 58.

`z = complex (x)`
`z = complex (re, im)`
: Return a complex value from real arguments.

  With 1 real argument $x$, return the complex result `x + 0i`.

  With 2 real arguments, return the complex result `re + imi`. `complex` can often be more convenient than expressions such as `a + b*i`. For example:

```
complex ([1, 2], [3, 4])
  ⇒ [ 1 + 3i   2 + 4i ]
```

  **See also:** [real], page 606, [imag], page 606, [iscomplex], page 69, [abs], page 605, [arg], page 605.

Octave automatically "narrows" complex types to real types when the imaginary part is zero. The interpreter performs this narrowing whenever it gets the opportunity, in order to save memory, and to emulate MATLAB behavior. For example, if a variable initially has the complex value `4 + 2i` and subsequently the value `2i` is subtracted from it, then it becomes the real type `4` and not the complex type `4 + 0i`. Testing whether the resultant type is complex using the expression `iscomplex (4+2i - 2i)` returns false.

To ensure that a value is of complex type even with a zero imaginary part, use `complex (4, 0)` instead of `4 + 0i`. But even then, the interpreter will still narrow the type to real at the next change done to it (such as adding 1 to it), unless its imaginary part truly becomes nonzero.

## 4.1 Matrices

It is easy to define a matrix of values in Octave. The size of the matrix is determined automatically, so it is not necessary to explicitly state the dimensions. The expression

```
a = [1, 2; 3, 4]
```

results in the matrix

$$a = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$$

Elements of a matrix may be arbitrary expressions, provided that the dimensions all make sense when combining the various pieces. For example, given the above matrix, the expression

```
[ a, a ]
```

produces the matrix

```
ans =

  1  2  1  2
  3  4  3  4
```

but the expression

```
[ a, 1 ]
```

produces the error

```
error: number of rows must match (1 != 2) near line 13, column 6
```

(assuming that this expression was entered as the first thing on line 13, of course).

Inside the square brackets that delimit a matrix expression, Octave looks at the surrounding context to determine whether spaces and newline characters should be converted into element and row separators, or simply ignored, so an expression like

```
a = [ 1 2
      3 4 ]
```

will work. However, some possible sources of confusion remain. For example, in the expression

```
[ 1 - 1 ]
```

the '`-`' is treated as a binary operator and the result is the scalar 0, but in the expression

```
[ 1 -1 ]
```

the '`-`' is treated as a unary operator and the result is the vector `[ 1, -1 ]`. Similarly, the expression

```
[ sin (pi) ]
```

will be parsed as

```
[ sin, (pi) ]
```

and will result in an error since the `sin` function will be called with no arguments. To get around this, you must omit the space between `sin` and the opening parenthesis, or enclose the expression in a set of parentheses:

```
[ (sin (pi)) ]
```

Whitespace surrounding the single quote character ('`'`', used as a transpose operator and for delimiting character strings) can also cause confusion. Given `a = 1`, the expression

```
[ 1 a' ]
```

results in the single quote character being treated as a transpose operator and the result is the vector `[ 1, 1 ]`, but the expression

```
[ 1 a ' ]
```

produces the error message

```
parse error:

  syntax error

>>> [ 1 a ' ]
              ^
```

because not doing so would cause trouble when parsing the valid expression

```
[ a 'foo' ]
```

For clarity, it is probably best to always use commas and semicolons to separate matrix elements and rows.

The maximum number of elements in a matrix is fixed when Octave is compiled. The allowable number can be queried with the function `sizemax`. Note that other factors, such as the amount of memory available on your machine, may limit the maximum size of matrices to something smaller.

*max_numel* = sizemax ()
: Return the largest value allowed for the size of an array.

  If Octave is compiled with 64-bit indexing, the result is of class int64, otherwise it is of class int32. The maximum array size is slightly smaller than the maximum value allowable for the relevant class as reported by `intmax`.

  **See also:** [intmax], page 60.

When you type a matrix or the name of a variable whose value is a matrix, Octave responds by printing the matrix in with neatly aligned rows and columns. If the rows of the matrix are too large to fit on the screen, Octave splits the matrix and displays a header before each section to indicate which columns are being displayed. You can use the following variables to control the format of the output.

*val* = output_precision ()
*old_val* = output_precision (*new_val*)
*old_val* = output_precision (*new_val*, "*local*")
: Query or set the internal variable that specifies the minimum number of significant figures to display for numeric output.

  Note that regardless of the value set for `output_precision`, the number of digits of precision displayed is limited to 16 for double precision values and 7 for single precision values. Also, calls to the `format` function that change numeric display can also change the set value for `output_precision`.

  When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

  **See also:** [format], page 288, [fixed_point_format], page 55.

It is possible to achieve a wide range of output styles by using different values of `output_precision`. Reasonable combinations can be set using the `format` function. See Section 14.1 [Basic Input and Output], page 287.

*val* = split_long_rows ()
*old_val* = split_long_rows (*new_val*)
*old_val* = split_long_rows (*new_val*, "*local*")
: Query or set the internal variable that controls whether rows of a matrix may be split when displayed to a terminal window.

  If the rows are split, Octave will display the matrix in a series of smaller pieces, each of which can fit within the limits of your terminal width and each set of rows is labeled so that you can easily see which columns are currently being displayed. For example:

```
octave:13> rand (2,10)
ans =

 Columns 1 through 6:

  0.75883  0.93290  0.40064  0.43818  0.94958  0.16467
  0.75697  0.51942  0.40031  0.61784  0.92309  0.40201

 Columns 7 through 10:

  0.90174  0.11854  0.72313  0.73326
  0.44672  0.94303  0.56564  0.82150
```

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [format], page 288.

Octave automatically switches to scientific notation when values become very large or very small. This guarantees that you will see several significant figures for every value in a matrix. If you would prefer to see all values in a matrix printed in a fixed point format, you can use the function `fixed_point_format`. But doing so is not recommended, because it can produce output that can easily be misinterpreted.

*val* = fixed_point_format ()
*old_val* = fixed_point_format (*new_val*)
*old_val* = fixed_point_format (*new_val*, "*local*")

Query or set the internal variable that controls whether Octave will use a scaled format to print matrix values.

The scaled format prints a scaling factor on the first line of output chosen such that the largest matrix element can be written with a single leading digit. For example:

```
fixed_point_format (true)
logspace (1, 7, 5)'
ans =

  1.0e+07  *

  0.00000
  0.00003
  0.00100
  0.03162
  1.00000
```

Notice that the first value appears to be 0 when it is actually 1. Because of the possibility for confusion you should be careful about enabling `fixed_point_format`.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [format], page 288, [output_precision], page 54.

### 4.1.1 Empty Matrices

A matrix may have one or both dimensions zero, and operations on empty matrices are handled as described by Carl de Boor in "An Empty Exercise", *SIGNUM*, Vol. 25, pp. 2–6, 1990 and C. N. Nett and W. M. Haddad, in "A System-Theoretic Appropriate Realization of the Empty Matrix Concept", *IEEE Transactions on Automatic Control*, Vol. 38, No. 5, May 1993. Briefly, given a scalar $s$, an $m \times n$ matrix $M_{m\times n}$, and an $m \times n$ empty matrix $[]_{m\times n}$ (with either one or both dimensions equal to zero), the following are true:

$$\begin{aligned} s \cdot []_{m\times n} &= []_{m\times n} \cdot s = []_{m\times n} \\ []_{m\times n} + []_{m\times n} &= []_{m\times n} \\ []_{0\times m} \cdot M_{m\times n} &= []_{0\times n} \\ M_{m\times n} \cdot []_{n\times 0} &= []_{m\times 0} \\ []_{m\times 0} \cdot []_{0\times n} &= 0_{m\times n} \end{aligned}$$

By default, dimensions of the empty matrix are printed along with the empty matrix symbol, '[]'. The built-in variable print_empty_dimensions controls this behavior.

*val* = print_empty_dimensions ()
*old_val* = print_empty_dimensions (*new_val*)
*old_val* = print_empty_dimensions (*new_val*, "*local*")

Query or set the internal variable that controls whether the dimensions of empty matrices are printed along with the empty matrix symbol, '[]'.

For example, the expression

```
zeros (3, 0)
```

will print

```
ans = [](3x0)
```

When called from inside a function with the "local" option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [format], page 288.

Empty matrices may also be used in assignment statements as a convenient way to delete rows or columns of matrices. See Section 8.6 [Assignment Expressions], page 179.

When Octave parses a matrix expression, it examines the elements of the list to determine whether they are all constants. If they are, it replaces the list with a single matrix constant.

## 4.2 Ranges

A *range* is a convenient way to write a row vector with evenly spaced elements. A range expression is defined by the value of the first element in the range, an optional value for the increment between elements, and a maximum value which the elements of the range will not exceed. The base, increment, and limit are separated by colons (the ':' character) and

may contain any arithmetic expressions and function calls. If the increment is omitted, it is assumed to be 1. For example, the range

```
1 : 5
```

defines the set of values `[ 1, 2, 3, 4, 5 ]`, and the range

```
1 : 3 : 5
```

defines the set of values `[ 1, 4 ]`.

Although a range constant specifies a row vector, Octave does *not* normally convert range constants to vectors unless it is necessary to do so. This allows you to write a constant like `1 : 10000` without using 80,000 bytes of storage on a typical workstation.

A common example of when it does become necessary to convert ranges into vectors occurs when they appear within a vector (i.e., inside square brackets). For instance, whereas

```
x = 0 : 0.1 : 1;
```

defines $x$ to be a variable of type `double_range` and occupies 24 bytes of memory, the expression

```
y = [ 0 : 0.1 : 1];
```

defines $y$ to be of type `matrix` and occupies 88 bytes of memory.

This space saving optimization may be disabled using the function `optimize_range`.

*val* = optimize_range ()
*old_val* = optimize_range (*new_val*)
*old_val* = optimize_range (*new_val*, "*local*")

Query or set whether a special space-efficient format is used for storing ranges.

The default value is true. If this option is set to false, Octave will store ranges as full matrices.

When called from inside a function with the `"local"` option, the setting is changed locally for the function and any subroutines it calls. The original setting is restored when exiting the function.

**See also:** [optimize_diagonal_matrix], page 717, [optimize_permutation_matrix], page 717.

Note that the upper (or lower, if the increment is negative) bound on the range is not always included in the set of values. This can be useful in some contexts. For example:

```
## x is some predefined range or vector or matrix or array
x(1:2:end) += 1;   # increment all  odd-numbered elements
x(2:2:end) -= 1;   # decrement all even-numbered elements
```

The above code works correctly whether $x$ has an odd number of elements or not: there is no need to treat the two cases differently.

Octave uses floating point arithmetic to compute the values in the range. As a result, defining ranges with floating-point values can result in pitfalls like these:

```
a = -2
b = (0.3 - 0.2 - 0.1)
x = a : b
```

Due to floating point rounding, $b$ may or may not equal zero exactly, and if it does not, it may be above zero or below zero, hence the final range $x$ may or may not include zero as its final value. Similarly:

```
x = 1.80 : 0.05 : 1.90
y = 1.85 : 0.05 : 1.90
```

is not as predictable as it looks. As of Octave 8.3, the results obtained are that $x$ has three elements (1.80, 1.85, and 1.90), and $y$ has only one element (1.85 but not 1.90). Thus, when using floating points in ranges, changing the start of the range can easily affect the end of the range even though the ending value was not touched in the above example.

To avoid such pitfalls with floating-points in ranges, you can use one of the following patterns. This change to the previous code:

```
x = (0:2) * 0.05 + 1.80
y = (0:1) * 0.05 + 1.85
```

use integers to construct the range and then converts to floating point making the overall expression safe and repeatable across platforms, compilers, and compiler settings. If you know the number of elements, you can use the `linspace` function (see Section 16.3 [Special Utility Matrices], page 580), which will include the endpoints of a range. You can also make judicious use of `round`, `floor`, `ceil`, `fix`, etc. to set the limits and the increment without getting interference from floating-point rounding. For example, the earlier example can be made safer with one of the following:

```
a = -2
b = round ((0.3 - 0.2 - 0.1) * 1e12) / 1e12   # rounds to 12 digits
c = floor (0.3 - 0.2 - 0.1)                   # floors as integer
d = floor ((0.3 - 0.2 - 0.1) * 1e12) / 1e12   # floors at 12 digits
x = a : b
y = a : c
z = a : d
```

If the result of the range expression is empty then Octave returns an empty `matrix`, not an empty `double_range`. Similarly, if there is just a single element in the range then Octave returns a `scalar`, not a `double_range` with one element.

## 4.3 Single Precision Data Types

Octave includes support for single precision data types, and most of the functions in Octave accept single precision values and return single precision answers. A single precision variable is created with the `single` function.

`y = single (x)`
: Convert $x$ to single precision type.

  **See also:** [double], page 51.

for example:

```
sngl = single (rand (2, 2))
     ⇒ sngl =
        0.37569   0.92982
        0.11962   0.50876
class (sngl)
    ⇒ single
```

Many functions can also return single precision values directly. For example

```
ones (2, 2, "single")
zeros (2, 2, "single")
eye (2, 2,  "single")
rand (2, 2, "single")
NaN (2, 2, "single")
NA (2, 2, "single")
Inf (2, 2, "single")
```

will all return single precision matrices.

## 4.4 Integer Data Types

Octave supports integer matrices as an alternative to using double precision. It is possible to use both signed and unsigned integers represented by 8, 16, 32, or 64 bits. It should be noted that most computations require floating point data, meaning that integers will often change type when involved in numeric computations. For this reason integers are most often used to store data, and not for calculations.

In general, most integer matrices are created by casting existing matrices to integers. The following example shows how to cast a matrix into 32-bit integers.

```
float = rand (2, 2)
     ⇒ float = 0.37569   0.92982
               0.11962   0.50876
integer = int32 (float)
     ⇒ integer = 0  1
                 0  1
```

As can be seen, floating point values are rounded to the nearest integer when converted.

*tf* = **isinteger** (*x*)

Return true if *x* is an integer object (int8, uint8, int16, etc.).

Note that `isinteger (14)` is false because numeric constants in Octave are double precision floating point values.

**See also:** [isfloat], page 69, [ischar], page 77, [islogical], page 69, [isstring], page 77, [isnumeric], page 69, [isa], page 41.

*y* = **int8** (*x*)

Convert *x* to 8-bit integer type.

**See also:** [uint8], page 60, [int16], page 60, [uint16], page 60, [int32], page 60, [uint32], page 60, [int64], page 60, [uint64], page 60.

`y = uint8 (x)`
Convert *x* to unsigned 8-bit integer type.

**See also:** [int8], page 59, [int16], page 60, [uint16], page 60, [int32], page 60, [uint32], page 60, [int64], page 60, [uint64], page 60.

`y = int16 (x)`
Convert *x* to 16-bit integer type.

**See also:** [int8], page 59, [uint8], page 60, [uint16], page 60, [int32], page 60, [uint32], page 60, [int64], page 60, [uint64], page 60.

`y = uint16 (x)`
Convert *x* to unsigned 16-bit integer type.

**See also:** [int8], page 59, [uint8], page 60, [int16], page 60, [int32], page 60, [uint32], page 60, [int64], page 60, [uint64], page 60.

`y = int32 (x)`
Convert *x* to 32-bit integer type.

**See also:** [int8], page 59, [uint8], page 60, [int16], page 60, [uint16], page 60, [uint32], page 60, [int64], page 60, [uint64], page 60.

`y = uint32 (x)`
Convert *x* to unsigned 32-bit integer type.

**See also:** [int8], page 59, [uint8], page 60, [int16], page 60, [uint16], page 60, [int32], page 60, [int64], page 60, [uint64], page 60.

`y = int64 (x)`
Convert *x* to 64-bit integer type.

**See also:** [int8], page 59, [uint8], page 60, [int16], page 60, [uint16], page 60, [int32], page 60, [uint32], page 60, [uint64], page 60.

`y = uint64 (x)`
Convert *x* to unsigned 64-bit integer type.

**See also:** [int8], page 59, [uint8], page 60, [int16], page 60, [uint16], page 60, [int32], page 60, [uint32], page 60, [int64], page 60.

*Imax* `= intmax ()`
*Imax* `= intmax ("`*type*`")`
*Imax* `= intmax (`*var*`)`
Return the largest integer that can be represented by a specific integer type.

The input is either a string `"`*type*`"` specifying an integer type, or it is an existing integer variable *var*.

Possible values for *type* are

`"int8"` signed 8-bit integer.

`"int16"` signed 16-bit integer.

`"int32"` signed 32-bit integer.

`"int64"` signed 64-bit integer.

`"uint8"` unsigned 8-bit integer.

`"uint16"` unsigned 16-bit integer.

`"uint32"` unsigned 32-bit integer.

`"uint64"` unsigned 64-bit integer.

The default for *type* is `"int32"`.

Example Code - query an existing variable

```
x = int8 (1);
intmax (x)
  ⇒ 127
```

**See also:** [intmin], page 61, [flintmax], page 61.

*Imin* = intmin ()
*Imin* = intmin ("*type*")
*Imin* = intmin (*var*)

Return the smallest integer that can be represented by a specific integer type.

The input is either a string `"type"` specifying an integer type, or it is an existing integer variable *var*.

Possible values for *type* are

`"int8"` signed 8-bit integer.

`"int16"` signed 16-bit integer.

`"int32"` signed 32-bit integer.

`"int64"` signed 64-bit integer.

`"uint8"` unsigned 8-bit integer.

`"uint16"` unsigned 16-bit integer.

`"uint32"` unsigned 32-bit integer.

`"uint64"` unsigned 64-bit integer.

The default for *type* is `"int32"`.

Example Code - query an existing variable

```
x = int8 (1);
intmin (x)
  ⇒ -128
```

**See also:** [intmax], page 60, [flintmax], page 61.

*Imax* = flintmax ()
*Imax* = flintmax ("*double*")
*Imax* = flintmax ("*single*")
*Imax* = flintmax (*var*)

Return the largest integer that can be represented consecutively in a floating point value.

The input is either a string specifying a floating point type, or it is an existing floating point variable *var*.

The default type is `"double"`, but `"single"` is a valid option. On IEEE 754 compatible systems, `flintmax` is $2^{53}$ for `"double"` and $2^{24}$ for `"single"`.

Example Code - query an existing variable

```
x = single (1);
flintmax (x)
  ⇒ 16777216
```

**See also:** [intmax], page 60, [realmax], page 646, [realmin], page 646.

### 4.4.1 Hexadecimal and Binary Integer Constants

The use of hexadecimal or binary notation to define a number will automatically create an unsigned integer using a representation that is just large enough to hold the specified value. For example:

```
0b101        # uint8
0x100        # uint16
0xDEADBEEF   # uint32
0x1DEADBEEF  # uint64
```

The storage class can be specified by adding a suffix. Use '`s`' for signed integers and '`u`' for unsigned integers along with a size ('`8`', '`16`', '`32`', '`64`'). The use of the underscore separator '`_`' can improve readability. For example:

```
0b101s16        # int16
0b101_s16       # int16, value and representation separated
0xDEADBEEFs32   # int32
0xDEADBEEF_u64  # uint64
```

Note that when defining matrices of integer constants the overall matrix will have the storage class of its first element. The matrix `[0x1; 0x100; 0x10000]` will be of type `uint8` and the larger values will be truncated because of the saturation semantics of integer values. To avoid this issue either: 1) declare the first integer to be of the desired size such as `0x1u32; 0x100; 0x10000]`, or 2) pad constants in array expressions with leading zeros so that they use the same number of digits for each value such as `[0x00_00_01; 0x00_01_00; 0x01_00_00]`.

### 4.4.2 Integer Arithmetic

While many numerical computations can't be carried out in integers, Octave does support basic operations like addition and multiplication on integers. The operators `+`, `-`, `.*`, and `./` work on integers of the same type. So, it is possible to add two 32-bit integers, but not to add a 32-bit integer and a 16-bit integer.

When doing integer arithmetic one should consider the possibility of underflow and overflow. This happens when the result of the computation can't be represented using the chosen integer type. As an example it is not possible to represent the result of $10 - 20$ when using unsigned integers. Octave makes sure that the result of integer computations is the integer that is closest to the true result. So, the result of $10 - 20$ when using unsigned integers is zero.

When doing integer division Octave will round the result to the nearest integer. This is different from most programming languages, where the result is often floored to the nearest integer. So, the result of `int32 (5) ./ int32 (8)` is `1`.

*C* = `idivide` (*A*, *B*, *op*)

Integer division with different rounding rules.

The standard behavior of integer division such as `A ./ B` is to round the result to the nearest integer. This is not always the desired behavior and `idivide` permits integer element-by-element division to be performed with different treatment for the fractional part of the division as determined by the *op* flag. *op* is a string with one of the values:

`"fix"` Calculate `A ./ B` with the fractional part rounded towards zero.

`"round"` Calculate `A ./ B` with the fractional part rounded towards the nearest integer.

`"floor"` Calculate `A ./ B` with the fractional part rounded towards negative infinity.

`"ceil"` Calculate `A ./ B` with the fractional part rounded towards positive infinity.

If *op* is not given it defaults to `"fix"`. An example demonstrating these rounding rules is

```
idivide (int8 ([-3, 3]), int8 (4), "fix")
  ⇒   0   0
idivide (int8 ([-3, 3]), int8 (4), "round")
  ⇒  -1   1
idivide (int8 ([-3, 3]), int8 (4), "floor")
  ⇒  -1   0
idivide (int8 ([-3, 3]), int8 (4), "ceil")
  ⇒   0   1
```

**See also:** [ceil], page 616, [floor], page 616, [fix], page 616, [round], page 617, [ldivide], page 172, [rdivide], page 173.

## 4.5 Bit Manipulations

Octave provides a number of functions for the manipulation of numeric values on a bit by bit basis. The basic functions to set and obtain the values of individual bits are `bitset` and `bitget`.

*B* = `bitset` (*A*, *n*)
*B* = `bitset` (*A*, *n*, *val*)

Set or reset bit(s) at position *n* of the unsigned integers in *A*.

The least significant bit is $n = 1$. $val = 0$ resets bits and $val = 1$ sets bits. If no *val* is specified it defaults to 1 (set bit). All inputs must be the same size or scalars.

Example 1: Set multiple bits

```
x = bitset (1, 3:5)
  ⇒ x =

   5    9   17

dec2bin (x)
  ⇒
     00101
     01001
     10001
```

Example 2: Reset and set bits

```
x = bitset ([15 14], 1, [0 1])
  ⇒ x =

   14    15
```

**See also:** [bitand], page 64, [bitor], page 65, [bitxor], page 65, [bitget], page 64, [bitcmp], page 65, [bitshift], page 65, [intmax], page 60, [flintmax], page 61.

`b = bitget (A, n)`

Return the bit value at position(s) $n$ of the unsigned integers in $A$.

The least significant bit is $n = 1$.

```
bitget (100, 8:-1:1)
⇒ 0  1  1  0  0  1  0  0
```

**See also:** [bitand], page 64, [bitor], page 65, [bitxor], page 65, [bitset], page 63, [bitcmp], page 65, [bitshift], page 65, [intmax], page 60, [flintmax], page 61.

The arguments to all of Octave's bitwise operations can be scalar or arrays, except for `bitcmp`, whose $k$ argument must a scalar. In the case where more than one argument is an array, then all arguments must have the same shape, and the bitwise operator is applied to each of the elements of the argument individually. If at least one argument is a scalar and one an array, then the scalar argument is duplicated. Therefore

```
bitget (100, 8:-1:1)
```

is the same as

```
bitget (100 * ones (1, 8), 8:-1:1)
```

It should be noted that all values passed to the bit manipulation functions of Octave are treated as integers. Therefore, even though the example for `bitset` above passes the floating point value `10`, it is treated as the bits `[1, 0, 1, 0]` rather than the bits of the native floating point format representation of `10`.

As the maximum value that can be represented by a number is important for bit manipulation, particularly when forming masks, Octave supplies two utility functions: `flintmax` for floating point integers, and `intmax` for integer objects (`uint8`, `int64`, etc.).

Octave also includes the basic bitwise 'and', 'or', and 'exclusive or' operators.

`z = bitand (x, y)`

Return the bitwise AND of non-negative integers.

$x$, $y$ must be in the range [0,intmax]

**See also:** [bitor], page 65, [bitxor], page 65, [bitset], page 63, [bitget], page 64, [bitcmp], page 65, [bitshift], page 65, [intmax], page 60, [flintmax], page 61.

`z = bitor (x, y)`

Return the bitwise OR of non-negative integers $x$ and $y$.

**See also:** [bitor], page 65, [bitxor], page 65, [bitset], page 63, [bitget], page 64, [bitcmp], page 65, [bitshift], page 65, [intmax], page 60, [flintmax], page 61.

`z = bitxor (x, y)`

Return the bitwise XOR of non-negative integers $x$ and $y$.

**See also:** [bitand], page 64, [bitor], page 65, [bitset], page 63, [bitget], page 64, [bitcmp], page 65, [bitshift], page 65, [intmax], page 60, [flintmax], page 61.

The bitwise 'not' operator is a unary operator that performs a logical negation of each of the bits of the value. For this to make sense, the mask against which the value is negated must be defined. Octave's bitwise 'not' operator is `bitcmp`.

`C = bitcmp (A, k)`

Return the $k$-bit complement of integers in $A$.

If $k$ is omitted `k = log2 (flintmax) + 1` is assumed.

```
bitcmp (7,4)
  ⇒ 8
dec2bin (11)
  ⇒ 1011
dec2bin (bitcmp (11, 6))
  ⇒ 110100
```

**See also:** [bitand], page 64, [bitor], page 65, [bitxor], page 65, [bitset], page 63, [bitget], page 64, [bitcmp], page 65, [bitshift], page 65, [flintmax], page 61.

Octave also includes the ability to left-shift and right-shift values bitwise.

`B = bitshift (A, k)`
`B = bitshift (A, k, n)`

Return a $k$ bit shift of $n$-digit unsigned integers in $A$.

A positive $k$ leads to a left shift; A negative value to a right shift.

If $n$ is omitted it defaults to 64. $n$ must be in the range [1,64].

```
bitshift (eye (3), 1)
⇒
2 0 0
0 2 0
0 0 2

bitshift (10, [-2, -1, 0, 1, 2])
⇒ 2   5  10  20  40
```

**See also:** [bitand], page 64, [bitor], page 65, [bitxor], page 65, [bitset], page 63, [bitget], page 64, [bitcmp], page 65, [intmax], page 60, [flintmax], page 61.

Bits that are shifted out of either end of the value are lost. Octave also uses arithmetic shifts, where the sign bit of the value is kept during a right shift. For example:

```
bitshift (-10, -1)
⇒ -5
bitshift (int8 (-1), -1)
⇒ -1
```

Note that `bitshift (int8 (-1), -1)` is `-1` since the bit representation of `-1` in the `int8` data type is `[1, 1, 1, 1, 1, 1, 1, 1]`.

## 4.6 Logical Values

Octave has built-in support for logical values, i.e., variables that are either `true` or `false`. When comparing two variables, the result will be a logical value whose value depends on whether or not the comparison is true.

The basic logical operations are `&`, `|`, and `!`, which correspond to "Logical And", "Logical Or", and "Logical Negation". These operations all follow the usual rules of logic.

It is also possible to use logical values as part of standard numerical calculations. In this case `true` is converted to `1`, and `false` to 0, both represented using double precision floating point numbers. So, the result of `true*22 - false/6` is 22.

Logical values can also be used to index matrices and cell arrays. When indexing with a logical array the result will be a vector containing the values corresponding to `true` parts of the logical array. See [Logical Indexing], page 161.

Logical values can also be constructed by casting numeric objects to logical values, or by using the `true` or `false` functions.

*TF* = `logical` (*x*)

Convert the numeric object *x* to logical type.

Any nonzero values will be converted to true (1) while zero values will be converted to false (0). The non-numeric value NaN cannot be converted and will produce an error.

Compatibility Note: Octave accepts complex values as input, whereas MATLAB issues an error.

**See also:** [double], page 51, [single], page 58, [char], page 83.

*T* = `true`
*T* = `true` (*n*)
*T* = `true` (*m*, *n*, ...)
*T* = `true` ([*m*, *n*, ...])
*T* = `true` (..., "*like*", *var*)

Return a scalar, matrix, or N-dimensional array whose elements are all logical `1`.

If called with no arguments, return the scalar value logical `1`.

If invoked with a single scalar integer argument *n*, return a square NxN matrix.

If invoked with two or more scalar integer arguments, or a vector of integer values, return an array with the given dimensions.

If a logical variable *var* is specified after `"like"`, the output $T$ will have the same sparsity as *var*.

Programming Notes: The code `true (...)` is faster (30X) and more memory efficient than `logical (ones (...))`.

Any negative dimensions are treated as zero, and any zero dimensions will result in an empty matrix. This odd behavior is for MATLAB compatibility.

**See also:** [false], page 67, [logical], page 66, [ones], page 581, [zeros], page 581.

*F* = `false`
*F* = `false` (*n*)
*F* = `false` (*m*, *n*, ...)
*F* = `false` ([*m*, *n*, ...])
*F* = `false` (..., "*like*", *var*)

Return a scalar, matrix, or N-dimensional array whose elements are all logical `0`.

If called with no arguments, return the scalar value logical `0`.

If invoked with a single scalar integer argument *n*, return a square NxN matrix.

If invoked with two or more scalar integer arguments, or a vector of integer values, return an array with the given dimensions.

If a logical variable *var* is specified after `"like"`, the output $F$ will have the same sparsity as *var*.

Programming Note: The code `false (...)` is faster (30X) and more memory efficient than `logical (zeros (...))`.

Any negative dimensions are treated as zero, and any zero dimensions will result in an empty matrix. This odd behavior is for MATLAB compatibility.

**See also:** [true], page 66, [logical], page 66, [ones], page 581, [zeros], page 581.

## 4.7 Automatic Conversion of Data Types

Many operators and functions can work with mixed data types. For example,

```
uint8 (1) + 1
    ⇒ 2

single (1) + 1
    ⇒ 2

min (single (1), 0)
   ⇒ 0
```

where the results are respectively of types uint8, single, and single respectively. This is done for MATLAB compatibility. Valid mixed operations are defined as follows:

| **Mixed Operation** | **Result** |
|---|---|
| double OP single | single |
| double OP integer | integer |
| double OP char | double |

| | |
|---|---|
| double OP logical | double |
| single OP integer | integer |
| single OP char | single |
| single OP logical | single |

When functions expect a double but are passed other types, automatic conversion is function-dependent:

```
a = det (int8 ([1 2; 3 4]))
    ⇒ a = -2
class (a)
    ⇒ double

a = eig (int8 ([1 2; 3 4]))
    ⇒ error: eig: wrong type argument 'int8 matrix'
```

When two operands are both integers but of different widths, then some cases convert them to the wider bitwidth, and other cases throw an error:

```
a = min (int8 (100), int16 (200))
    ⇒ 100
class (a)
    ⇒ int16

int8 (100) + int16 (200)
   ⇒ error: binary operator '+' not implemented
   for 'int8 scalar' by 'int16 scalar' operations
```

For two integer operands, they typically need to both be signed or both be unsigned. Mixing signed and unsigned usually causes an error, even if they are of the same bitwidth.

```
min (int16 (100), uint16 (200))
   ⇒ error: min: cannot compute min (int16 scalar, uint16 scalar)
```

In the case of mixed type indexed assignments, the type is not changed. For example,

```
x = ones (2, 2);
x(1, 1) = single (2)
   ⇒ x = 2   1
         1   1
```

where `x` remains of the double precision type.

## 4.8 Predicates for Numeric Objects

Since the type of a variable may change during the execution of a program, it can be necessary to do type checking at run-time. Doing this also allows you to change the behavior of a function depending on the type of the input. As an example, this naive implementation of `abs` returns the absolute value of the input if it is a real number, and the length of the input if it is a complex number.

```
function a = abs (x)
  if (isreal (x))
    a = sign (x) .* x;
  elseif (iscomplex (x))
    a = sqrt (real(x).^2 + imag(x).^2);
  endif
endfunction
```

The following functions are available for determining the type of a variable.

*tf* = `isnumeric` (*x*)

Return true if *x* is a numeric object, i.e., an integer, real, or complex array.

Logical and character arrays are not considered to be numeric.

**See also:** [isinteger], page 59, [isfloat], page 69, [isreal], page 69, [iscomplex], page 69, [ischar], page 77, [islogical], page 69, [isstring], page 77, [iscell], page 132, [isstruct], page 125, [isa], page 41.

*tf* = `islogical` (*x*)
*tf* = `isbool` (*x*)

Return true if *x* is a logical object.

Programming Note: `isbool` is an alias for `islogical` and can be used interchangeably.

**See also:** [ischar], page 77, [isfloat], page 69, [isinteger], page 59, [isstring], page 77, [isnumeric], page 69, [isa], page 41.

*tf* = `isfloat` (*x*)

Return true if *x* is a floating-point numeric object.

Objects of class double or single are floating-point objects.

**See also:** [isinteger], page 59, [ischar], page 77, [islogical], page 69, [isnumeric], page 69, [isstring], page 77, [isa], page 41.

*tf* = `isreal` (*x*)

Return true if *x* is a non-complex matrix or scalar.

For compatibility with MATLAB, this includes logical and character matrices.

**See also:** [iscomplex], page 69, [isnumeric], page 69, [isa], page 41.

*tf* = `iscomplex` (*x*)

Return true if *x* is a complex-valued numeric object.

**See also:** [isreal], page 69, [isnumeric], page 69, [ischar], page 77, [isfloat], page 69, [islogical], page 69, [isstring], page 77, [isa], page 41.

*tf* = `ismatrix` (*x*)

Return true if *x* is a 2-D array.

A matrix is an array of any type where `ndims (x) == 2` and for which `size (x)` returns `[M, N]` with non-negative M and N.

**See also:** [isscalar], page 70, [isvector], page 70, [iscell], page 132, [isstruct], page 125, [issparse], page 734, [isa], page 41.

`tf = isvector (x)`

Return true if $x$ is a vector.

A vector is a 2-D array of any type where one of the dimensions is equal to 1 (either 1xN or Nx1). As a consequence of this definition, a 1x1 object (a scalar) is also a vector.

**See also:** [isscalar], page 70, [ismatrix], page 69, [iscolumn], page 70, [isrow], page 70, [size], page 48.

`tf = isrow (x)`

Return true if $x$ is a row vector.

A row vector is a 2-D array of any type for which `size (x)` returns `[1, N]` with non-negative N.

**See also:** [iscolumn], page 70, [isscalar], page 70, [isvector], page 70, [ismatrix], page 69, [size], page 48.

`tf = iscolumn (x)`

Return true if $x$ is a column vector.

A column vector is a 2-D array of any type for which `size (x)` returns `[N, 1]` with non-negative N.

**See also:** [isrow], page 70, [isscalar], page 70, [isvector], page 70, [ismatrix], page 69, [size], page 48.

`tf = isscalar (x)`

Return true if $x$ is a scalar.

A scalar is a single-element object of any type for which `size (x)` returns `[1, 1]`.

**See also:** [isvector], page 70, [ismatrix], page 69, [size], page 48.

`tf = issquare (x)`

Return true if $x$ is a 2-D square array.

A square array is a 2-D array of any type for which `size (x)` returns `[N, N]` where N is a non-negative integer.

**See also:** [isscalar], page 70, [isvector], page 70, [ismatrix], page 69, [size], page 48.

`tf = issymmetric (A)`
`tf = issymmetric (A, tol)`
`tf = issymmetric (A, "skew")`
`tf = issymmetric (A, "skew", tol)`

Return true if $A$ is a symmetric or skew-symmetric numeric matrix within the tolerance specified by *tol*.

The default tolerance is zero (uses faster code).

The type of symmetry to check may be specified with the additional input `"nonskew"` (default) for regular symmetry or `"skew"` for skew-symmetry.

Background: A matrix is symmetric if the transpose of the matrix is equal to the original matrix: `A == A.'`. If a tolerance is given then symmetry is determined by `norm (A - A.', Inf) / norm (A, Inf) < tol`.

A matrix is skew-symmetric if the transpose of the matrix is equal to the negative of the original matrix: `A == -A.'`. If a tolerance is given then skew-symmetry is determined by `norm (A + A.', Inf) / norm (A, Inf) < tol`.

**See also:** [ishermitian], page 71, [isdefinite], page 71.

`tf = ishermitian (A)`
`tf = ishermitian (A, tol)`
`tf = ishermitian (A, "skew")`
`tf = ishermitian (A, "skew", tol)`

Return true if *A* is a Hermitian or skew-Hermitian numeric matrix within the tolerance specified by *tol*.

The default tolerance is zero (uses faster code).

The type of symmetry to check may be specified with the additional input `"nonskew"` (default) for regular Hermitian or `"skew"` for skew-Hermitian.

Background: A matrix is Hermitian if the complex conjugate transpose of the matrix is equal to the original matrix: `A == A'`. If a tolerance is given then the calculation is `norm (A - A', Inf) / norm (A, Inf) < tol`.

A matrix is skew-Hermitian if the complex conjugate transpose of the matrix is equal to the negative of the original matrix: `A == -A'`. If a tolerance is given then the calculation is `norm (A + A', Inf) / norm (A, Inf) < tol`.

**See also:** [issymmetric], page 70, [isdefinite], page 71.

`tf = isdefinite (A)`
`tf = isdefinite (A, tol)`

Return true if *A* is symmetric positive definite numeric matrix within the tolerance specified by *tol*.

If *tol* is omitted, use a tolerance of `100 * eps * norm (A, "fro")`.

Background: A positive definite matrix has eigenvalues which are all greater than zero. A positive semi-definite matrix has eigenvalues which are all greater than or equal to zero. The matrix *A* is very likely to be positive semi-definite if the following two conditions hold for a suitably small tolerance *tol*.

```
isdefinite (A) ⇒ 0
isdefinite (A + 5*tol, tol) ⇒ 1
```

**See also:** [issymmetric], page 70, [ishermitian], page 71.

`tf = isbanded (A, lower, upper)`

Return true if *A* is a numeric matrix with entries confined between *lower* diagonals below the main diagonal and *upper* diagonals above the main diagonal.

*lower* and *upper* must be non-negative integers.

**See also:** [isdiag], page 71, [istril], page 72, [istriu], page 72, [bandwidth], page 648.

`tf = isdiag (A)`

Return true if *A* is a diagonal numeric matrix which is defined as a 2-D array where all elements above and below the main diagonal are zero.

**See also:** [isbanded], page 71, [istril], page 72, [istriu], page 72, [diag], page 579, [bandwidth], page 648.

*tf* = `istril (A)`

Return true if $A$ is a lower triangular numeric matrix.

A lower triangular matrix has nonzero entries only on the main diagonal and below.

**See also:** [istriu], page 72, [isbanded], page 71, [isdiag], page 71, [tril], page 577, [bandwidth], page 648.

*tf* = `istriu (A)`

Return true if $A$ is an upper triangular numeric matrix.

An upper triangular matrix has nonzero entries only on the main diagonal and above.

**See also:** [isdiag], page 71, [isbanded], page 71, [istril], page 72, [triu], page 578, [bandwidth], page 648.

*tf* = `isprime (x)`

Return a logical array which is true where the elements of $x$ are prime numbers and false where they are not.

A prime number is conventionally defined as a positive integer greater than 1 (e.g., 2, 3, ...) which is divisible only by itself and 1. Octave extends this definition to include both negative integers and complex values. A negative integer is prime if its positive counterpart is prime. This is equivalent to `isprime (abs (x))`.

If `class (x)` is complex, then primality is tested in the domain of Gaussian integers (https://en.wikipedia.org/wiki/Gaussian_integer). Some non-complex integers are prime in the ordinary sense, but not in the domain of Gaussian integers. For example, $5 = (1+2i) * (1-2i)$ shows that 5 is not prime because it has a factor other than itself and 1. Exercise caution when testing complex and real values together in the same matrix.

Examples:

```
isprime (1:6)
  ⇒  0  1  1  0  1  0

isprime ([i, 2, 3, 5])
  ⇒  0  0  1  0
```

Programming Note: `isprime` is suitable for all $x$ in the range $\text{abs}(x) < 2^{64}$. Cast inputs larger than `flintmax` to `uint64`.

For larger inputs, use 'sym' if you have the Symbolic package installed and loaded:

```
isprime (sym ('587453897092589025253904 50') + (0:4))
   ⇒  0  1  0  0  0
```

Compatibility Note: MATLAB does not extend the definition of prime numbers and will produce an error if given negative or complex inputs.

**See also:** [primes], page 627, [factor], page 625, [gcd], page 625, [lcm], page 626.

*tf* = `isuniform (v)`
[*tf*, *delta*] = `isuniform (v)`

Return true if the real vector $v$ is uniformly spaced and false otherwise.

A vector is uniform if the mean difference (*delta*) between all elements is the same to within a tolerance of `4 * eps (max (abs (v)))`.

The optional output *delta* is the uniform difference between elements. If the vector is not uniform then *delta* is `NaN`. *delta* is of the same class as *v* for floating point inputs and of class double for integer, logical, and character inputs.

Programming Notes: The output is always false for the special cases of an empty input or a scalar input. If any element is `NaN` then the output is false. If *delta* is smaller than the calculated relative tolerance then an absolute tolerance of `eps` is used.

**See also:** [linspace], page 584, [colon], page 982.

If instead of knowing properties of variables, you wish to know which variables are defined and to gather other information about the workspace itself, see Section 7.3 [Status of Variables], page 150.

# 5 Strings

A *string constant* consists of a sequence of characters enclosed in either double-quote or single-quote marks. For example, both of the following expressions

```
"parrot"
'parrot'
```

represent the string whose contents are '`parrot`'. Strings in Octave can be of any length.

Since the single-quote mark is also used for the transpose operator (see Section 8.3 [Arithmetic Ops], page 170) but double-quote marks have no other purpose in Octave, it is best to use double-quote marks to denote strings.

Strings can be concatenated using the notation for defining matrices. For example, the expression

```
[ "foo" , "bar" , "baz" ]
```

produces the string whose contents are '`foobarbaz`'. See Chapter 4 [Numeric Data Types], page 51, for more information about creating matrices.

While strings can in principle store arbitrary content, most functions expect them to be UTF-8 encoded Unicode strings.

Furthermore, it is possible to create a string without actually writing a text. The function `blanks` creates a string of a given length consisting only of blank characters (ASCII code 32).

*str* = **blanks** (*n*)

Return a string of *n* blanks.

For example:

```
blanks (10);
whos ans
     ⇒
      Attr Name        Size                     Bytes  Class
      ==== ====        ====                     =====  =====
           ans         1x10                        10  char
```

**See also:** [repmat], page 582.

## 5.1 Escape Sequences in String Constants

In double-quoted strings, the backslash character is used to introduce *escape sequences* that represent other characters. For example, '`\n`' embeds a newline character in a double-quoted string and '`\"`' embeds a double quote character. In single-quoted strings, backslash is not a special character. Here is an example showing the difference:

```
double ("\n")
    ⇒ 10
double ('\n')
    ⇒ [ 92 110 ]
```

Here is a table of all the escape sequences used in Octave (within double quoted strings). They are the same as those used in the C programming language.

`\\` Represents a literal backslash, '`\`'.

`\"` Represents a literal double-quote character, '"'.

`\'` Represents a literal single-quote character, "'".

`\0` Represents the null character, control-@, ASCII code 0.

`\a` Represents the "alert" character, control-g, ASCII code 7.

`\b` Represents a backspace, control-h, ASCII code 8.

`\f` Represents a formfeed, control-l, ASCII code 12.

`\n` Represents a newline, control-j, ASCII code 10.

`\r` Represents a carriage return, control-m, ASCII code 13.

`\t` Represents a horizontal tab, control-i, ASCII code 9.

`\v` Represents a vertical tab, control-k, ASCII code 11.

`\`*nnn* Represents the octal value *nnn*, where *nnn* are one to three digits between 0 and 7. For example, the code for the ASCII ESC (escape) character is '`\033`'.

`\x`*hh*`...` Represents the hexadecimal value *hh*, where *hh* are hexadecimal digits ('`0`' through '`9`' and either '`A`' through '`F`' or '`a`' through '`f`'). Like the same construct in ANSI C, the escape sequence continues until the first non-hexadecimal digit is seen. However, using more than two hexadecimal digits produces undefined results.

In a single-quoted string there is only one escape sequence: you may insert a single quote character using two single quote characters in succession. For example,

```
'I can''t escape'
    ⇒ I can't escape
```

In scripts the two different string types can be distinguished if necessary by using `is_dq_string` and `is_sq_string`.

*tf* = `is_dq_string` (*x*)
Return true if *x* is a double-quoted character string.

**See also:** [is_sq_string], page 76, [ischar], page 77.

*tf* = `is_sq_string` (*x*)
Return true if *x* is a single-quoted character string.

**See also:** [is_dq_string], page 76, [ischar], page 77.

## 5.2 Character Arrays

The string representation used by Octave is an array of characters, so internally the string `"ddddddddddd"` is actually a row vector of length 10 containing the value 100 in all places (100 is the ASCII code of `"d"`). This lends itself to the obvious generalization to character matrices. Using a matrix of characters, it is possible to represent a collection of same-length strings in one variable. The convention used in Octave is that each row in a character matrix is a separate string, but letting each column represent a string is equally possible.

The easiest way to create a character matrix is to put several strings together into a matrix.

```
collection = [ "String #1"; "String #2" ];
```

This creates a 2-by-9 character matrix.

The function `ischar` can be used to test if an object is a character matrix.

*tf* = **ischar** (*x*)
: Return true if *x* is a character array.

  **See also:** [isfloat], page 69, [isinteger], page 59, [islogical], page 69, [isnumeric], page 69, [isstring], page 77, [iscellstr], page 140, [isa], page 41.

*tf* = **isstring** (*s*)
: Return true if *s* is a string array.

  A string array is a data type that stores strings (row vectors of characters) at each element in the array. It is distinct from character arrays which are N-dimensional arrays where each element is a single 1x1 character. It is also distinct from cell arrays of strings which store strings at each element, but use cell indexing '`{}`' to access elements rather than string arrays which use ordinary array indexing '`()`'.

  Programming Note: Octave does not yet implement string arrays so this function will always return false.

  **See also:** [ischar], page 77, [iscellstr], page 140, [isfloat], page 69, [isinteger], page 59, [islogical], page 69, [isnumeric], page 69, [isa], page 41.

To test if an object is a string (i.e., a 1xN row vector of characters and not a character matrix) you can use the `ischar` function in combination with the `isrow` function as in the following example:

```
ischar (collection)
     ⇒ 1

ischar (collection) && isrow (collection)
     ⇒ 0

ischar ("my string") && isrow ("my string")
     ⇒ 1
```

One relevant question is, what happens when a character matrix is created from strings of different length. The answer is that Octave puts blank characters at the end of strings shorter than the longest string. It is possible to use a different character than the blank character using the `string_fill_char` function.

*val* = **string_fill_char** ()
*old_val* = **string_fill_char** (*new_val*)
*old_val* = **string_fill_char** (*new_val*, "*local*")
: Query or set the internal variable used to pad all rows of a character matrix to the same length.

The value must be a single character and the default is " " (a single space). For example:

```
string_fill_char ("X");
[ "these"; "are"; "strings" ]
    ⇒  "theseXX"
       "areXXXX"
       "strings"
```

When called from inside a function with the "local" option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

Another useful function to control the text justification in this case is the strjust function.

*str* = strjust (*s*)
*str* = strjust (*s*, *pos*)

Return the text, *s*, justified according to *pos*, which may be "left", "center", or "right".

If *pos* is omitted it defaults to "right".

Null characters are replaced by spaces. All other character data are treated as non-white space.

Example:

```
strjust (["a"; "ab"; "abc"; "abcd"])
     ⇒
        "   a"
        "  ab"
        " abc"
        "abcd"
```

**See also:** [deblank], page 79, [strrep], page 94, [strtrim], page 79, [untabify], page 80.

This shows a problem with character matrices. It simply isn't possible to represent strings of different lengths. The solution is to use a cell array of strings, which is described in Section 6.3.4 [Cell Arrays of Strings], page 140.

## 5.3 String Operations

Octave supports a wide range of functions for manipulating strings. Since a string is just a matrix, simple manipulations can be accomplished using standard operators. The following example shows how to replace all blank characters with underscores.

```
quote = ...
  "First things first, but not necessarily in that order";
quote( quote == " " ) = "_"
⇒ quote =
    First_things_first,_but_not_necessarily_in_that_order
```

For more complex manipulations, such as searching, replacing, and general regular expressions, the following functions come with Octave.

### 5.3.1 Common String Operations

The following functions are useful to perform common String operations.

`y = lower (s)`
`y = tolower (s)`

Return a copy of the string or cell string *s*, with each uppercase character replaced by the corresponding lowercase one; non-alphabetic characters are left unchanged.

For example:

```
lower ("MiXeD cAsE 123")
    ⇒ "mixed case 123"
```

Programming Note: `tolower` is an alias for `lower` and either name can be used in Octave.

**See also:** [upper], page 79.

`y = upper (s)`
`y = toupper (s)`

Return a copy of the string or cell string *s*, with each lowercase character replaced by the corresponding uppercase one; non-alphabetic characters are left unchanged.

For example:

```
upper ("MiXeD cAsE 123")
    ⇒ "MIXED CASE 123"
```

Programming Note: `toupper` is an alias for `upper` and either name can be used in Octave.

**See also:** [lower], page 79.

`s = deblank (s)`

Remove trailing whitespace and nulls from *s*.

If *s* is a matrix, *deblank* trims each row to the length of the longest string. If *s* is a cell array of strings, operate recursively on each string element.

Examples:

```
deblank ("    abc  ")
     ⇒  "    abc"

deblank ([" abc   "; "   def   "])
     ⇒  [" abc  " ; "   def"]
```

**See also:** [strtrim], page 79.

`s = strtrim (s)`

Remove leading and trailing whitespace from *s*.

If *s* is a matrix, *strtrim* trims each row to the length of longest string. If *s* is a cell array of strings, operate recursively on each string element.

For example:

```
strtrim ("    abc  ")
     ⇒  "abc"

strtrim ([" abc   "; "   def   "])
     ⇒  ["abc  "  ; "  def"]
```

**See also:** [deblank], page 79.

*s* = strtrunc (*s*, *n*)

Truncate the character string *s* to length *n*.

If *s* is a character matrix, then the number of columns is adjusted.

If *s* is a cell array of strings, then the operation is performed on each cell element and the new cell array is returned.

*str* = untabify (*t*)
*str* = untabify (*t*, *tw*)
*str* = untabify (*t*, *tw*, *deblank*)

Replace TAB characters in *t* with spaces.

The input, *t*, may be either a 2-D character array, or a cell array of character strings. The output is the same class as the input.

The tab width is specified by *tw*, and defaults to eight.

If the optional argument *deblank* is true, then the spaces will be removed from the end of the character data.

The following example reads a file and writes an untabified version of the same file with trailing spaces stripped.

```
fid = fopen ("tabbed_script.m");
text = char (fread (fid, "uchar")');
fclose (fid);
fid = fopen ("untabified_script.m", "w");
text = untabify (strsplit (text, "\n"), 8, true);
fprintf (fid, "%s\n", text{:});
fclose (fid);
```

**See also:** [strjust], page 78, [strsplit], page 86, [deblank], page 79.

*newstr* = do_string_escapes (*string*)

Convert escape sequences in *string* to the characters they represent.

Escape sequences begin with a leading backslash (`'\'`) followed by 1–3 characters (.e.g., `"\n"` => newline).

**See also:** [undo_string_escapes], page 80.

*newstr* = undo_string_escapes (*string*)

Convert special characters in *string* back to their escaped forms.

For example, the expression

```
bell = "\a";
```

assigns the value of the alert character (control-g, ASCII code 7) to the string variable `bell`. If this string is printed, the system will ring the terminal bell (if it is possible).

This is normally the desired outcome. However, sometimes it is useful to be able to print the original representation of the string, with the special characters replaced by their escape sequences. For example,

```
octave:13> undo_string_escapes (bell)
ans = \a
```

replaces the unprintable alert character with its printable representation.

**See also:** [do_string_escapes], page 80.

### 5.3.2 Concatenating Strings

Strings can be concatenated using matrix notation (see Chapter 5 [Strings], page 75, Section 5.2 [Character Arrays], page 76) which is often the most natural method. For example:

```
fullname = [fname ".txt"];
email = ["<" user "@" domain ">"];
```

In each case it is easy to see what the final string will look like. This method is also the most efficient. When using matrix concatenation the parser immediately begins joining the strings without having to process the overhead of a function call and the input validation of the associated function.

The `newline` function can be used to join strings such that they appear as multiple lines of text when displayed.

*c* = `newline`

Return the character corresponding to a newline.

This is equivalent to `"\n"`.

Example Code

```
joined_string = [newline "line1" newline "line2"]
⇒
line1
line2
```

**See also:** [strcat], page 84, [strjoin], page 88, [strsplit], page 86.

In addition, there are several other functions for concatenating string objects which can be useful in specific circumstances: `char`, `strvcat`, `strcat`, and `cstrcat`. Finally, the general purpose concatenation functions can be used: see [cat], page 571, [horzcat], page 572, and [vertcat], page 572.

- All string concatenation functions except `cstrcat` convert numerical input into character data by taking the corresponding UTF-8 character for each element (or multi-byte sequence), as in the following example:

  ```
  char ([98, 97, 110, 97, 110, 97])
    ⇒ banana
  ```

  For conversion between locale encodings and UTF-8, see [unicode2native], page 99, and [native2unicode], page 99.

- `char` and `strvcat` concatenate vertically, while `strcat` and `cstrcat` concatenate horizontally. For example:

```
char ("an apple", "two pears")
    ⇒ an apple
      two pears

strcat ("oc", "tave", " is", " good", " for you")
     ⇒ octave is good for you
```

- `char` generates an empty row in the output for each empty string in the input. `strvcat`, on the other hand, eliminates empty strings.

```
char ("orange", "green", "", "red")
    ⇒ orange
      green

      red

strvcat ("orange", "green", "", "red")
    ⇒ orange
      green
      red
```

- All string concatenation functions except `cstrcat` also accept cell array data (see Section 6.3 [Cell Arrays], page 131). `char` and `strvcat` convert cell arrays into character arrays, while `strcat` concatenates within the cells of the cell arrays:

```
char ({"red", "green", "", "blue"})
     ⇒ red
       green

       blue

strcat ({"abc"; "ghi"}, {"def"; "jkl"})
     ⇒
       {
         [1,1] = abcdef
         [2,1] = ghijkl
       }
```

- `strcat` removes trailing white space in the arguments (except within cell arrays), while `cstrcat` leaves white space untouched. Both kinds of behavior can be useful as can be seen in the examples:

```
strcat (["dir1";"directory2"], ["/";"/"], ["file1";"file2"])
     ⇒ dir1/file1
       directory2/file2

cstrcat (["thirteen apples"; "a banana"], [" 5$";" 1$"])
      ⇒ thirteen apples 5$
        a banana        1$
```

Note that in the above example for `cstrcat`, the white space originates from the internal representation of the strings in a string array (see Section 5.2 [Character Arrays], page 76).

*C* = `char` (*A*)
*C* = `char` (*A*, ...)
*C* = `char` (*str1*, *str2*, ...)
*C* = `char` (*cell_array*)
'' = `char` ()

Create a string array from one or more numeric matrices, character matrices, or cell arrays.

Arguments are concatenated vertically. The returned values are padded with blanks as needed to make each row of the string array have the same length. Empty input strings are significant and will concatenated in the output.

For numerical input, each element is converted to the corresponding ASCII character. A range error results if an input is outside the ASCII range (0-255).

For cell arrays, each element is concatenated separately. Cell arrays converted through `char` can mostly be converted back with `cellstr`. For example:

```
char ([97, 98, 99], "", {"98", "99", 100}, "str1", ["ha", "lf"])
   ⇒ ["abc "
      "    "
      "98  "
      "99  "
      "d   "
      "str1"
      "half"]
```

**See also:** [strvcat], page 83, [cellstr], page 140.

*C* = `strvcat` (*A*)
*C* = `strvcat` (*A*, ...)
*C* = `strvcat` (*str1*, *str2*, ...)
*C* = `strvcat` (*cell_array*)

Create a character array from one or more numeric matrices, character matrices, or cell arrays.

Arguments are concatenated vertically. The returned values are padded with blanks as needed to make each row of the string array have the same length. Unlike `char`, empty strings are removed and will not appear in the output.

For numerical input, each element is converted to the corresponding ASCII character. A range error results if an input is outside the ASCII range (0-255).

For cell arrays, each element is concatenated separately. Cell arrays converted through `strvcat` can mostly be converted back with `cellstr`. For example:

```
strvcat ([97, 98, 99], "", {"98", "99", 100}, "str1", ["ha", "lf"])
      ⇒ ["abc "
          "98  "
          "99  "
          "d   "
          "str1"
          "half"]
```

**See also:** [char], page 83, [strcat], page 84, [cstrcat], page 84.

*str* = strcat (*s1*, *s2*, ...)

Return a string containing all the arguments concatenated horizontally.

If the arguments are cell strings, `strcat` returns a cell string with the individual cells concatenated. For numerical input, each element is converted to the corresponding ASCII character. Trailing white space for any character string input is eliminated before the strings are concatenated. Note that cell string values do **not** have whitespace trimmed.

For example:

```
strcat ("|", " leading space is preserved", "|")
    ⇒ | leading space is preserved|

strcat ("|", "trailing space is eliminated ", "|")
    ⇒ |trailing space is eliminated|

strcat ("homogeneous space |", "  ", "| is also eliminated")
    ⇒ homogeneous space || is also eliminated

s = [ "ab"; "cde" ];
strcat (s, s, s)
    ⇒
        "ababab   "
        "cdecdecde"

s = { "ab"; "cd " };
strcat (s, s, s)
    ⇒
        {
          [1,1] = ababab
          [2,1] = cd cd cd
        }
```

**See also:** [cstrcat], page 84, [char], page 83, [strvcat], page 83.

*str* = cstrcat (*s1*, *s2*, ...)

Return a string containing all the arguments concatenated horizontally with trailing white space preserved.

For example:

```
cstrcat ("ab   ", "cd")
      ⇒ "ab   cd"
```

```
s = [ "ab"; "cde" ];
cstrcat (s, s, s)
      ⇒ "ab ab ab "
        "cdecdecde"
```

**See also:** [strcat], page 84, [char], page 83, [strvcat], page 83.

### 5.3.3 Splitting and Joining Strings

*str* = substr (*s*, *offset*)
substr (*s*, *offset*, *len*)

Return the substring of *s* which starts at character number *offset* and is *len* characters long.

Position numbering for offsets begins with 1. If *offset* is negative, extraction starts that far from the end of the string.

If *len* is omitted, the substring extends to the end of *s*. A negative value for *len* extracts to within *len* characters of the end of the string

Examples:

```
substr ("This is a test string", 6, 9)
     ⇒ "is a test"
substr ("This is a test string", -11)
     ⇒ "test string"
substr ("This is a test string", -11, -7)
     ⇒ "test"
```

This function is patterned after the equivalent function in Perl.

[*tok*, *rem*] = strtok (*str*)
[*tok*, *rem*] = strtok (*str*, *delim*)

Find all characters in the string *str* up to, but not including, the first character which is in the string *delim*.

*str* may also be a cell array of strings in which case the function executes on every individual string and returns a cell array of tokens and remainders.

Leading delimiters are ignored. If *delim* is not specified, whitespace is assumed.

If *rem* is requested, it contains the remainder of the string, starting at the first delimiter.

Examples:

```
strtok ("this is the life")
     ⇒ "this"

[tok, rem] = strtok ("14*27+31", "+-*/")
     ⇒
        tok = 14
        rem = *27+31
```

**See also:** [index], page 92, [strsplit], page 86, [strchr], page 92, [isspace], page 114.

`[`*`cstr`*`] = strsplit (`*`str`*`)`
`[`*`cstr`*`] = strsplit (`*`str`*`, `*`del`*`)`
`[`*`cstr`*`] = strsplit (...,` *`name`*`,` *`value`*`)`
`[`*`cstr`*`,` *`matches`*`] = strsplit (...)`

Split the string *str* using the delimiters specified by *del* and return a cell string array of substrings.

If a delimiter is not specified the string is split at whitespace `{" ", "\f", "\n", "\r", "\t", "\v"}`. Otherwise, the delimiter, *del* must be a string or cell array of strings. By default, consecutive delimiters in the input string *s* are collapsed into one resulting in a single split.

Supported *name*/*value* pair arguments are:

- *collapsedelimiters* which may take the value of `true` (default) or `false`.
- *delimitertype* which may take the value of `"simple"` (default) or `"regularexpression"`. A simple delimiter matches the text exactly as written. Otherwise, the syntax for regular expressions outlined in `regexp` is used.

The optional second output, *matches*, returns the delimiters which were matched in the original string.

Examples with simple delimiters:

```
strsplit ("a b c")
      ⇒
          {
            [1,1] = a
            [1,2] = b
            [1,3] = c
          }

strsplit ("a,b,c", ",")
      ⇒
          {
            [1,1] = a
            [1,2] = b
            [1,3] = c
          }

strsplit ("a foo b,bar c", {" ", ",", "foo", "bar"})
      ⇒
          {
            [1,1] = a
            [1,2] = b
            [1,3] = c
          }

strsplit ("a,,b, c", {",", " "}, "collapsedelimiters", false)
      ⇒
          {
```

```
  [1,1] = a
  [1,2] =
  [1,3] = b
  [1,4] =
  [1,5] = c
}
```

Examples with regularexpression delimiters:

```
strsplit ("a foo b,bar c", ',|\s|foo|bar', ...
          "delimitertype", "regularexpression")
⇒
{
  [1,1] = a
  [1,2] = b
  [1,3] = c
}

strsplit ("a,,b, c", '[, ]', "collapsedelimiters", false, ...
          "delimitertype", "regularexpression")
⇒
{
  [1,1] = a
  [1,2] =
  [1,3] = b
  [1,4] =
  [1,5] = c
}

strsplit ("a,\t,b, c", {',', '\s'}, "delimitertype", "regularexpression")
⇒
{
  [1,1] = a
  [1,2] = b
  [1,3] = c
}

strsplit ("a,\t,b, c", {',', ' ', '\t'}, "collapsedelimiters", false)
⇒
{
  [1,1] = a
  [1,2] =
  [1,3] =
  [1,4] = b
  [1,5] =
  [1,6] = c
}
```

**See also:** [ostrsplit], page 87, [strjoin], page 88, [strtok], page 85, [regexp], page 95.

`[`*cstr*`] = ostrsplit (`*s*, *sep*`)`
`[`*cstr*`] = ostrsplit (`*s*, *sep*, *strip_empty*`)`

Split the string *s* using one or more separators *sep* and return a cell array of strings.

Consecutive separators and separators at boundaries result in empty strings, unless *strip_empty* is true. The default value of *strip_empty* is false.

2-D character arrays are split at separators and at the original column boundaries.

Example:

```
ostrsplit ("a,b,c", ",")
      ⇒
          {
            [1,1] = a
            [1,2] = b
            [1,3] = c
          }

ostrsplit (["a,b" ; "cde"], ",")
      ⇒
          {
            [1,1] = a
            [1,2] = b
            [1,3] = cde
          }
```

**See also:** [strsplit], page 86, [strtok], page 85.

*str* = strjoin (*cstr*)
*str* = strjoin (*cstr*, *delimiter*)

Join the elements of the cell string array, *cstr*, into a single string.

If no *delimiter* is specified, the elements of *cstr* are separated by a space.

If *delimiter* is specified as a string, the cell string array is joined using the string. Escape sequences are supported.

If *delimiter* is a cell string array whose length is one less than *cstr*, then the elements of *cstr* are joined by interleaving the cell string elements of *delimiter*. Escape sequences are not supported.

```
strjoin ({'Octave','Scilab','Lush','Yorick'}, '*')
      ⇒ 'Octave*Scilab*Lush*Yorick'
```

**See also:** [strsplit], page 86.

### 5.3.4 Searching in Strings

Since a string is a character array, comparisons between strings work element by element as the following example shows:

```
GNU = "GNU's Not UNIX";
spaces = (GNU == " ")
     ⇒ spaces =
       0   0   0   0   0   1   0   0   0   1   0   0   0   0
```

To determine if two strings are identical it is necessary to use the `strcmp` function. It compares complete strings and is case sensitive. `strncmp` compares only the first `N` characters (with `N` given as a parameter). `strcmpi` and `strncmpi` are the corresponding functions for case-insensitive comparison.

*tf* = strcmp (*str1*, *str2*)

Return 1 if the character strings *str1* and *str2* are the same, and 0 otherwise.

If either *str1* or *str2* is a cell array of strings, then an array of the same size is returned, containing the values described above for every member of the cell array. The other argument may also be a cell array of strings (of the same size or with only one element), char matrix or character string.

**Caution:** For compatibility with MATLAB, Octave's strcmp function returns 1 if the character strings are equal, and 0 otherwise. This is just the opposite of the corresponding C library function.

**See also:** [strcmpi], page 89, [strncmp], page 89, [strncmpi], page 89.

*tf* = `strncmp` (*str1*, *str2*, *n*)

Return 1 if the first *n* characters of strings *str1* and *str2* are the same, and 0 otherwise.

```
strncmp ("abce", "abcd", 3)
      ⇒ 1
```

If either *str1* or *str2* is a cell array of strings, then an array of the same size is returned, containing the values described above for every member of the cell array. The other argument may also be a cell array of strings (of the same size or with only one element), char matrix or character string.

```
strncmp ("abce", {"abcd", "bca", "abc"}, 3)
     ⇒ [1, 0, 1]
```

**Caution:** For compatibility with MATLAB, Octave's strncmp function returns true if the character strings are equal, and false otherwise. This is just the opposite of the corresponding C library function. In addition Octave's strncmp function returns true if N = 0.

MATLAB incompatibility : Octave's strncmp function produces an error if N < 0, while MATLAB treats N < 0 the same as N = 0, always returning true.

**See also:** [strncmpi], page 89, [strcmp], page 88, [strcmpi], page 89.

*tf* = `strcmpi` (*str1*, *str2*)

Return 1 if the character strings *str1* and *str2* are the same, disregarding case of alphabetic characters, and 0 otherwise.

If either *str1* or *str2* is a cell array of strings, then an array of the same size is returned, containing the values described above for every member of the cell array. The other argument may also be a cell array of strings (of the same size or with only one element), char matrix or character string.

**Caution:** For compatibility with MATLAB, Octave's strcmpi function returns 1 if the character strings are equal, and 0 otherwise. This is just the opposite of the corresponding C library function.

**Caution:** National alphabets are not supported.

**See also:** [strcmp], page 88, [strncmp], page 89, [strncmpi], page 89.

*tf* = `strncmpi` (*str1*, *str2*, *n*)

Return 1 if the first *n* character of *s1* and *s2* are the same, disregarding case of alphabetic characters, and 0 otherwise.

If either *str1* or *str2* is a cell array of strings, then an array of the same size is returned, containing the values described above for every member of the cell array.

The other argument may also be a cell array of strings (of the same size or with only one element), char matrix or character string.

**Caution:** For compatibility with MATLAB, Octave's strncmpi function returns true if the character strings are equal, and false otherwise. This is just the opposite of the corresponding C library function. In addition Octave's strncmpi function returns true if N = 0.

MATLAB incompatibility : Octave's strncmpi function produces an error if N < 0, while MATLAB treats N < 0 the same as N = 0, always returning true.

**Caution:** National alphabets are not supported.

**See also:** [strncmp], page 89, [strcmp], page 88, [strcmpi], page 89.

Despite those comparison functions, there are more specialized function to find the index position of a search pattern within a string.

*retval* = startsWith (*str*, *pattern*)
*retval* = startsWith (*str*, *pattern*, "*IgnoreCase*", *ignore_case*)

Check whether string(s) start with pattern(s).

Return an array of logical values that indicates which string(s) in the input *str* (a single string or cell array of strings) begin with the input *pattern* (a single string or cell array of strings).

If the value of the parameter "IgnoreCase" is true, then the function will ignore the letter case of *str* and *pattern*. By default, the comparison is case sensitive.

Examples:

```
## one string and one pattern while considering case
startsWith ("hello", "he")
      ⇒  1

## one string and one pattern while ignoring case
startsWith ("hello", "HE", "IgnoreCase", true)
      ⇒  1

## multiple strings and multiple patterns while considering case
startsWith ({"lab work.pptx", "data.txt", "foundations.ppt"},
            {"lab", "data"})
      ⇒  1  1  0

## multiple strings and one pattern while considering case
startsWith ({"DATASHEET.ods", "data.txt", "foundations.ppt"},
            "data", "IgnoreCase", false)
      ⇒  0  1  0

## multiple strings and one pattern while ignoring case
startsWith ({"DATASHEET.ods", "data.txt", "foundations.ppt"},
            "data", "IgnoreCase", true)
      ⇒  1  1  0
```

**See also:** [endsWith], page 91, [regexp], page 95, [strncmp], page 89, [strncmpi], page 89.

*retval* = endsWith (*str*, *pattern*)
*retval* = endsWith (*str*, *pattern*, "*IgnoreCase*", *ignore_case*)

Check whether string(s) end with pattern(s).

Return an array of logical values that indicates which string(s) in the input *str* (a single string or cell array of strings) end with the input *pattern* (a single string or cell array of strings).

If the value of the parameter `"IgnoreCase"` is true, then the function will ignore the letter case of *str* and *pattern*. By default, the comparison is case sensitive.

Examples:

```
## one string and one pattern while considering case
endsWith ("hello", "lo")
      ⇒  1

## one string and one pattern while ignoring case
endsWith ("hello", "LO", "IgnoreCase", true)
      ⇒  1

## multiple strings and multiple patterns while considering case
endsWith ({"tests.txt", "mydoc.odt", "myFunc.m", "results.pptx"},
          {".docx", ".odt", ".txt"})
      ⇒  1  1  0  0

## multiple strings and one pattern while considering case
endsWith ({"TESTS.TXT", "mydoc.odt", "result.txt", "myFunc.m"},
          ".txt", "IgnoreCase", false)
      ⇒  0  0  1  0

## multiple strings and one pattern while ignoring case
endsWith ({"TESTS.TXT", "mydoc.odt", "result.txt", "myFunc.m"},
          ".txt", "IgnoreCase", true)
      ⇒  1  0  1  0
```

**See also:** [startsWith], page 90, [regexp], page 95, [strncmp], page 89, [strncmpi], page 89.

*v* = findstr (*s*, *t*)
*v* = findstr (*s*, *t*, *overlap*)

This function is obsolete. Use `strfind` instead.

Return the vector of all positions in the longer of the two strings *s* and *t* where an occurrence of the shorter of the two starts.

If the optional argument *overlap* is true (default), the returned vector can include overlapping positions. For example:

```
findstr ("abababa", "a")
     ⇒ [1, 3, 5];
findstr ("abababa", "aba", 0)
     ⇒ [1, 5]
```

**Caution:** `findstr` is obsolete. Use `strfind` in all new code.

**See also:** [strfind], page 93, [strmatch], page 93, [strcmp], page 88, [strncmp], page 89, [strcmpi], page 89, [strncmpi], page 89, [find], page 568.

*idx* = `strchr` (*str*, *chars*)
*idx* = `strchr` (*str*, *chars*, *n*)
*idx* = `strchr` (*str*, *chars*, *n*, *direction*)
[*i*, *j*] = `strchr` (...)

Search through the string *str* for occurrences of characters from the set *chars*.

The return value(s), as well as the *n* and *direction* arguments behave identically as in `find`.

This will be faster than using `regexp` in most cases.

**See also:** [find], page 568.

*n* = `index` (*s*, *t*)
*n* = `index` (*s*, *t*, *direction*)

Return the position of the first occurrence of the string *t* in the string *s*, or 0 if no occurrence is found.

*s* may also be a string array or cell array of strings.

For example:

```
index ("Teststring", "t")
    ⇒ 4
```

If *direction* is `"first"`, return the first element found. If *direction* is `"last"`, return the last element found.

**See also:** [find], page 568, [rindex], page 92.

*n* = `rindex` (*s*, *t*)

Return the position of the last occurrence of the character string *t* in the character string *s*, or 0 if no occurrence is found.

*s* may also be a string array or cell array of strings.

For example:

```
rindex ("Teststring", "t")
     ⇒ 6
```

The `rindex` function is equivalent to `index` with *direction* set to `"last"`.

**See also:** [find], page 568, [index], page 92.

*idx* = `unicode_idx` (*str*)

Return an array with the indices for each UTF-8 encoded character in *str*.

```
unicode_idx ("aäbc")
     ⇒ [1, 2, 2, 3, 4]
```

*idx* = strfind (*str*, *pattern*)
*idx* = strfind (*cellstr*, *pattern*)
*idx* = strfind (..., "*overlaps*", *val*)
*idx* = strfind (..., "*forcecelloutput*", *val*)

Search for *pattern* in the string *str* and return the starting index of every such occurrence in the vector *idx*.

If there is no such occurrence, or if *pattern* is longer than *str*, or if *pattern* itself is empty, then *idx* is the empty array `[]`.

The optional argument `"overlaps"` determines whether the pattern can match at every position in *str* (true), or only for unique occurrences of the complete pattern (false). The default is true.

If a cell array of strings *cellstr* is specified then *idx* is a cell array of vectors, as specified above.

The optional argument `"forcecelloutput"` forces *idx* to be returned as a cell array of vectors. The default is false.

Examples:

```
strfind ("abababa", "aba")
     ⇒ [1, 3, 5]

strfind ("abababa", "aba", "overlaps", false)
     ⇒ [1, 5]

strfind ({"abababa", "bebebe", "ab"}, "aba")
     ⇒
        {
          [1,1] =

             1   3   5

          [1,2] = [](1x0)
          [1,3] = [](1x0)
        }

strfind ("abababa", "aba", "forcecelloutput", true)
     ⇒
        {
          [1,1] =

             1   3   5
        }
```

**See also:** [regexp], page 95, [regexpi], page 98, [find], page 568.

*idx* = strmatch (*s*, *A*)
*idx* = strmatch (*s*, *A*, "*exact*")

This function is obsolete. **Use an alternative** such as `strncmp` or `strcmp` instead.

Return indices of entries of *A* which begin with the string *s*.

The second argument *A* must be a string, character matrix, or a cell array of strings.

If the third argument `"exact"` is not given, then *s* only needs to match *A* up to the length of *s*. Trailing spaces and nulls in *s* and *A* are ignored when matching.

For example:

```
strmatch ("apple", "apple juice")
     ⇒ 1

strmatch ("apple", ["apple  "; "apple juice"; "an apple"])
     ⇒ [1; 2]

strmatch ("apple", ["apple  "; "apple juice"; "an apple"], "exact")
     ⇒ [1]
```

**Caution:** `strmatch` is obsolete (and can produce incorrect results in MATLAB when used with cell arrays of strings. Use `strncmp` (normal case) or `strcmp` (`"exact"` case) in all new code. Other replacement possibilities, depending on application, include `regexp` or `validatestring`.

**See also:** [strncmp], page 89, [strcmp], page 88, [regexp], page 95, [strfind], page 93, [validatestring], page 217.

### 5.3.5 Searching and Replacing in Strings

*newstr* = strrep (*str*, *ptn*, *rep*)
*newstr* = strrep (*cellstr*, *ptn*, *rep*)
*newstr* = strrep (*...*, "*overlaps*", *val*)

Replace all occurrences of the pattern *ptn* in the string *str* with the string *rep* and return the result.

The optional argument `"overlaps"` determines whether the pattern can match at every position in *str* (true), or only for unique occurrences of the complete pattern (false). The default is true.

*s* may also be a cell array of strings, in which case the replacement is done for each element and a cell array is returned.

Example:

```
strrep ("This is a test string", "is", "&%$")
    ⇒  "Th&%$ &%$ a test string"
```

**See also:** [regexprep], page 98, [strfind], page 93.

*newstr* = erase (*str*, *ptn*)

Delete all occurrences of *ptn* within *str*.

*str* and *ptn* can be ordinary strings, cell array of strings, or character arrays.

Examples

```
## string, single pattern
erase ("Hello World!", " World")
    ⇒ "Hello!"

## cellstr, single pattern
erase ({"Hello", "World!"}, "World")
    ⇒ {"Hello", "!"}

## string, multiple patterns
erase ("The Octave interpreter is fabulous", ...
       {"interpreter ", "The "})
    ⇒ "Octave is fabulous"

## cellstr, multiple patterns
erase ({"The ", "Octave interpreter ", "is fabulous"}, ...
       {"interpreter ", "The "})
    ⇒ {"", "Octave ", "is fabulous"}
```

Programming Note: `erase` deletes the first instance of a pattern in a string when there are overlapping occurrences. For example:

```
erase ("abababa", "aba")
    ⇒ "b"
```

For processing overlaps, see [strrep], page 94.

**See also:** [strrep], page 94, [regexprep], page 98.

`[`*s*`, `*e*`, `*te*`, `*m*`, `*t*`, `*nm*`, `*sp*`] = regexp (`*str*`, `*pat*`)`
`[...] = regexp (`*str*`, `*pat*`, "`*opt1*`", ...)`

Regular expression string matching.

Search for *pat* in UTF-8 encoded *str* and return the positions and substrings of any matches, or empty values if there are none.

The matched pattern *pat* can include any of the standard regex operators, including:

`.` Match any character

`* + ? {}` Repetition operators, representing

- `*` Match zero or more times
- `+` Match one or more times
- `?` Match zero or one times
- `{`*n*`}` Match exactly *n* times
- `{`*n*`,}` Match *n* or more times
- `{`*m*`,`*n*`}` Match between *m* and *n* times

`[...] [^...]`

List operators. The pattern will match any character listed between "`[`" and "`]`". If the first character is "`^`" then the pattern is inverted and any character except those listed between brackets will match.

Escape sequences defined below can also be used inside list operators. For example, a template for a floating point number might be `[-+.\d]+`.

`() (?:)` Grouping operator. The first form, parentheses only, also creates a token.

`|` Alternation operator. Match one of a choice of regular expressions. The alternatives must be delimited by the grouping operator `()` above.

`^ $` Anchoring operators. Requires pattern to occur at the start (`^`) or end (`$`) of the string.

In addition, the following escaped characters have special meaning.

`\d` Match any digit

`\D` Match any non-digit

`\s` Match any whitespace character

`\S` Match any non-whitespace character

`\w` Match any word character

`\W` Match any non-word character

`\<` Match the beginning of a word

`\>` Match the end of a word

`\B` Match within a word

Implementation Note: For compatibility with MATLAB, escape sequences in *pat* (e.g., `"\n"` => newline) are expanded even when *pat* has been defined with single quotes. To disable expansion use a second backslash before the escape sequence (e.g., `"\\n"`) or use the `regexptranslate` function.

The outputs of `regexp` default to the order given below

*s* The start indices of each matching substring

*e* The end indices of each matching substring

*te* The extents of each matched token surrounded by `(...)` in *pat*

*m* A cell array of the text of each match

*t* A cell array of the text of each token matched

*nm* A structure containing the text of each matched named token, with the name being used as the fieldname. A named token is denoted by `(?<name>...)`.

*sp* A cell array of the text not returned by match, i.e., what remains if you split the string based on *pat*.

Particular output arguments, or the order of the output arguments, can be selected by additional *opt* arguments. These are strings and the correspondence between the output arguments and the optional argument are

| | |
|---|---|
| `'start'` | *s* |
| `'end'` | *e* |

| | |
|---|---|
| `'tokenExtents'` | *te* |
| `'match'` | *m* |
| `'tokens'` | *t* |
| `'names'` | *nm* |
| `'split'` | *sp* |

Additional arguments are summarized below.

'`once`' Return only the first occurrence of the pattern.

'`matchcase`'
Make the matching case sensitive. (default)
Alternatively, use (?-i) in the pattern.

'`ignorecase`'
Ignore case when matching the pattern to the string.
Alternatively, use (?i) in the pattern.

'`stringanchors`'
Match the anchor characters at the beginning and end of the string. (default)
Alternatively, use (?-m) in the pattern.

'`lineanchors`'
Match the anchor characters at the beginning and end of the line.
Alternatively, use (?m) in the pattern.

'`dotall`' The pattern `.` matches all characters including the newline character. (default)
Alternatively, use (?s) in the pattern.

'`dotexceptnewline`'
The pattern `.` matches all characters except the newline character.
Alternatively, use (?-s) in the pattern.

'`literalspacing`'
All characters in the pattern, including whitespace, are significant and are used in pattern matching. (default)
Alternatively, use (?-x) in the pattern.

'`freespacing`'
The pattern may include arbitrary whitespace and also comments beginning with the character '`#`'.
Alternatively, use (?x) in the pattern.

'`noemptymatch`'
Zero-length matches are not returned. (default)

'`emptymatch`'
Return zero-length matches.
`regexp ('a', 'b*', 'emptymatch')` returns `[1 2]` because there are zero or more `'b'` characters at positions 1 and end-of-string.

Stack Limitation Note: Pattern searches are done with a recursive function which can overflow the program stack when there are a high number of matches. For example,

```
regexp (repmat ('a', 1, 1e5), '(a)+')
```

may lead to a segfault. As an alternative, consider constructing pattern searches that reduce the number of matches (e.g., by creatively using set complement), and then further processing the return variables (now reduced in size) with successive `regexp` searches.

Octave's `regexp` implementation is based on the Perl Compatible Regular Expressions library (`https://www.pcre.org/`). For a more comprehensive list of `regexp` operator syntax see the "PCRE Syntax quick-reference summary".

**See also:** [regexpi], page 98, [strfind], page 93, [regexprep], page 98.

`[`*s*`, `*e*`, `*te*`, `*m*`, `*t*`, `*nm*`, `*sp*`] = regexpi (`*str*`, `*pat*`)`
`[...] = regexpi (`*str*`, `*pat*`, "`*opt1*`", ...)`

Case insensitive regular expression string matching.

Search for *pat* in UTF-8 encoded *str* and return the positions and substrings of any matches, or empty values if there are none. See [regexp], page 95, for details on the syntax of the search pattern.

**See also:** [regexp], page 95.

*outstr* `= regexprep (`*string*`, `*pat*`, `*repstr*`)`
*outstr* `= regexprep (`*string*`, `*pat*`, `*repstr*`, "`*opt1*`", ...)`

Replace occurrences of pattern *pat* in *string* with *repstr*.

The pattern is a regular expression as documented for `regexp`. See [regexp], page 95.

All strings must be UTF-8 encoded.

The replacement string may contain `$i`, which substitutes for the ith set of parentheses in the match string. For example,

```
regexprep ("Bill Dunn", '(\w+) (\w+)', '$2, $1')
```

returns `"Dunn, Bill"`

Options in addition to those of `regexp` are

'`once`' Replace only the first occurrence of *pat* in the result.

'`warnings`'
This option is present for compatibility but is ignored.

Implementation Note: For compatibility with MATLAB, escape sequences in *pat* (e.g., `"\n"` => newline) are expanded even when *pat* has been defined with single quotes. To disable expansion use a second backslash before the escape sequence (e.g., `"\\n"`) or use the `regexptranslate` function.

**See also:** [regexp], page 95, [regexpi], page 98, [strrep], page 94.

*str* `= regexptranslate (`*op*`, `*s*`)`

Translate a string for use in a regular expression.

This may include either wildcard replacement or special character escaping.

The behavior is controlled by *op* which can take the following values

`"wildcard"`
: The wildcard characters `.`, `*`, and `?` are replaced with wildcards that are appropriate for a regular expression. For example:

```
regexptranslate ("wildcard", "*.m")
     ⇒ '.*\.m'
```

`"escape"`
: The characters `$.?[]`, that have special meaning for regular expressions are escaped so that they are treated literally. For example:

```
regexptranslate ("escape", "12.5")
     ⇒ '12\.5'
```

**See also:** [regexp], page 95, [regexpi], page 98, [regexprep], page 98.

## 5.4 Converting Strings

Octave offers several kinds of conversion functions for Strings.

### 5.4.1 String encoding

*native_bytes* = unicode2native (*utf8_str*, *codepage*)
*native_bytes* = unicode2native (*utf8_str*)
: Convert UTF-8 string *utf8_str* to byte stream using *codepage*.

The character vector *utf8_str* is converted to a byte stream *native_bytes* using the code page given by *codepage*. The string *codepage* must be an identifier of a valid code page. Examples for valid code pages are `"ISO-8859-1"`, `"Shift-JIS"`, or `"UTF-16"`. For a list of supported code pages, see https://www.gnu.org/software/libiconv. If *codepage* is omitted or empty, the system default codepage is used.

If any of the characters cannot be mapped into the codepage *codepage*, they are replaced with the appropriate substitution sequence for that codepage.

**See also:** [native2unicode], page 99.

*utf8_str* = native2unicode (*native_bytes*, *codepage*)
*utf8_str* = native2unicode (*native_bytes*)
: Convert byte stream *native_bytes* to UTF-8 using *codepage*.

The numbers in the vector *native_bytes* are rounded and clipped to integers between 0 and 255. This byte stream is then mapped into the code page given by the string *codepage* and returned in the string *utf8_str*. Octave uses UTF-8 as its internal encoding. The string *codepage* must be an identifier of a valid code page. Examples for valid code pages are `"ISO-8859-1"`, `"Shift-JIS"`, or `"UTF-16"`. For a list of supported code pages, see https://www.gnu.org/software/libiconv. If *codepage* is omitted or empty, the system default codepage is used.

If *native_bytes* is a string vector, it is returned as is.

**See also:** [unicode2native], page 99.

### 5.4.2 Numerical Data and Strings

Apart from the string concatenation functions (see Section 5.3.2 [Concatenating Strings], page 81) which cast numerical data to the corresponding UTF-8 encoded characters, there are several functions that format numerical data as strings. `mat2str` and `num2str` convert real or complex matrices, while `int2str` converts integer matrices. `int2str` takes the real part of complex values and round fractional values to integer. A more flexible way to format numerical data as strings is the `sprintf` function (see Section 14.2.4 [Formatted Output], page 315, [sprintf], page 316).

*s* = mat2str (*x*)
*s* = mat2str (*x*, *n*)
*s* = mat2str (*x*, [*n1*, *n2*])
*s* = mat2str (..., "*class*")

Format real, complex, and logical matrices as strings.

The returned string may be used to reconstruct the original matrix by using the `eval` function.

The precision of the values is given by *n*. If *n* is a scalar then both real and imaginary parts of the matrix are printed to the same precision. Otherwise `n(1)` defines the precision of the real part and `n(2)` defines the precision of the imaginary part. The default for *n* is 15.

If the argument `"class"` is given then the class of *x* is included in the string in such a way that `eval` will result in the construction of a matrix of the same class.

```
mat2str (pi)
     ⇒ "3.14159265358979"

mat2str (pi, 5)
     ⇒ "3.1416"

mat2str ([ -1/3 + i/7; 1/3 - i/7 ], [4 2])
     ⇒ "[-0.3333+0.14i;0.3333-0.14i]"

mat2str ([ -1/3 +i/7; 1/3 -i/7 ], [4 2])
     ⇒ "[-0.3333+0i 0+0.14i;0.3333+0i -0-0.14i]"

mat2str (int16 ([1 -1]), "class")
     ⇒ "int16([1 -1])"

mat2str (logical (eye (2)))
     ⇒ "[true false;false true]"

isequal (x, eval (mat2str (x)))
     ⇒ 1
```

**See also:** [sprintf], page 316, [num2str], page 100, [int2str], page 101.

*str* = num2str (*x*)
*str* = num2str (*x*, *precision*)

*str* = num2str (*x*, *format*)

Convert a number (or array) to a string (or a character array).

The optional second argument may either give the number of significant digits (*precision*) to be used in the output or a format template string (*format*) as in `sprintf` (see Section 14.2.4 [Formatted Output], page 315). `num2str` can also process complex numbers.

Examples:

```
num2str (123.456)
  ⇒ 123.456

num2str (123.456, 4)
  ⇒ 123.5

s = num2str ([1, 1.34; 3, 3.56], "%5.1f")
  ⇒ s =
       1.0  1.3
       3.0  3.6
whos s
  ⇒ Variables in the current scope:
       Attr Name        Size                     Bytes  Class
       ==== ====        ====                     =====  =====
            s           2x8                         16  char
     Total is 16 elements using 16 bytes

num2str (1.234 + 27.3i)
  ⇒ 1.234+27.3i
```

The `num2str` function is not very flexible. For better control over the results, use `sprintf` (see Section 14.2.4 [Formatted Output], page 315).

Programming Notes:

For MATLAB compatibility, leading spaces are stripped before returning the string.

Integers larger than `flintmax` may not be displayed correctly.

For complex *x*, the format string may only contain one output conversion specification and nothing else. Otherwise, results will be unpredictable.

Any optional *format* specified by the programmer is used without modification. This is in contrast to MATLAB which tampers with the *format* based on internal heuristics.

**See also:** [sprintf], page 316, [int2str], page 101, [mat2str], page 100.

*str* = int2str (*n*)

Convert an integer (or array of integers) to a string (or a character array).

```
int2str (123)
  ⇒ 123

s = int2str ([1, 2, 3; 4, 5, 6])
  ⇒ s =
       1  2  3
       4  5  6

whos s
  ⇒ Variables in the current scope:
       Attr Name        Size                     Bytes  Class
       ==== ====        ====                     =====  =====
            s           2x7                         14  char
     Total is 14 elements using 14 bytes
```

This function is not very flexible. For better control over the results, use `sprintf` (see Section 14.2.4 [Formatted Output], page 315).

Programming Notes:

Non-integers are rounded to integers before display. Only the real part of complex numbers is displayed.

**See also:** [sprintf], page 316, [num2str], page 100, [mat2str], page 100.

*d* = `str2double` (*str*)

Convert a string to a real or complex number.

The string must be in one of the following formats where a and b are real numbers and the complex unit is `'i'` or `'j'`:

- a + bi
- a + b*i
- a + i*b
- bi + a
- b*i + a
- i*b + a

If present, a and/or b are of the form [+-]d[,.]d[[eE][+-]d] where the brackets indicate optional arguments and `'d'` indicates zero or more digits. The special input values `Inf`, `NaN`, and `NA` are also accepted.

*str* may be a character string, character matrix, or cell array. For character arrays the conversion is repeated for every row, and a double or complex array is returned. Empty rows in *s* are deleted and not returned in the numeric array. For cell arrays each character string element is processed and a double or complex array of the same dimensions as *str* is returned.

For unconvertible scalar or character string input `str2double` returns a NaN. Similarly, for character array input `str2double` returns a NaN for any row of *s* that could not be converted. For a cell array, `str2double` returns a NaN for any element of *s* for which conversion fails. Note that numeric elements in a mixed string/numeric cell array are not strings and the conversion will fail for these elements and return NaN.

Programming Note: `str2double` can replace `str2num`, is more efficient, and avoids the security risk of using `eval` on unknown data.

**See also:** [str2num], page 103.

`x = str2num (s)`
`[x, state] = str2num (s)`

Convert the string (or character array) *s* to a number (or an array).

Examples:

```
str2num ("3.141596")
      ⇒ 3.141596

str2num (["1, 2, 3"; "4, 5, 6"])
      ⇒ 1  2  3
        4  5  6
```

The optional second output, *state*, is logically true when the conversion is successful. If the conversion fails the numeric output, *x*, is empty and *state* is false.

**Caution:** As `str2num` uses the `eval` function to do the conversion, `str2num` will execute any code contained in the string *s*. Use `str2double` for a safer and faster conversion.

For cell array of strings use `str2double`.

**See also:** [str2double], page 102, [eval], page 187.

`d = bin2dec (str)`

Return the decimal number corresponding to the binary number represented by the string *str*.

For example:

```
bin2dec ("1110")
     ⇒ 14
```

Spaces are ignored during conversion and may be used to make the binary number more readable.

```
bin2dec ("1000 0001")
     ⇒ 129
```

If *str* is a string matrix, return a column vector with one converted number per row of *str*; Invalid rows evaluate to NaN.

If *str* is a cell array of strings, return a column vector with one converted number per cell element in *str*.

**See also:** [dec2bin], page 103, [base2dec], page 106, [hex2dec], page 104.

`bstr = dec2bin (d)`
`bstr = dec2bin (d, len)`

Return a string of ones and zeros representing the conversion of the integer *d* to a binary number.

If *d* is a matrix or cell array, return a string matrix with one row for each element in *d*, padded with leading zeros to the width of the largest value.

The optional second argument, *len*, specifies the minimum number of digits in the result.

For negative elements of *d*, return the binary value of the two's complement. The result is padded with leading ones to 8, 16, 32, or 64 bits as appropriate for the magnitude of the input. Positive input elements are padded with leading zeros to the same width.

Examples:

```
dec2bin (14)
     ⇒ "1110"

dec2bin (-14)
     ⇒ "11110010"
```

Programming tip: `dec2bin` discards any fractional part of the input. If you need the fractional part to be converted too, call `dec2base` with a nonzero number of decimal places. You can also use `fix` or `round` on fractional inputs to ensure predictable rounding behavior.

**See also:** [bin2dec], page 103, [dec2base], page 105, [dec2hex], page 104.

*hstr* = dec2hex (*d*)
*hstr* = dec2hex (*d*, *len*)

Return a string representing the conversion of the integer *d* to a hexadecimal (base16) number.

If *d* is negative, return the hexadecimal complement of *d*.

If *d* is a matrix or cell array, return a string matrix with one row for each element in *d*, padded with leading zeros to the width of the largest value.

The optional second argument, *len*, specifies the minimum number of digits in the result.

Examples:

```
dec2hex (2748)
     ⇒ "ABC"

dec2hex (-2)
     ⇒ "FE"
```

Programming tip: `dec2hex` discards any fractional part of the input. If you need the fractional part to be converted too, call `dec2base` with a nonzero number of decimal places. You can also use `fix` or `round` on fractional inputs to ensure predictable rounding behavior.

**See also:** [hex2dec], page 104, [dec2base], page 105, [dec2bin], page 103.

*d* = hex2dec (*str*)

Return the integer corresponding to the hexadecimal number represented by the string *str*.

For example:

```
hex2dec ("12B")
      ⇒ 299
hex2dec ("12b")
      ⇒ 299
```

If *str* is a string matrix, return a column vector with one converted number per row of *str*; Invalid rows evaluate to NaN.

If *str* is a cell array of strings, return a column vector with one converted number per cell element in *str*.

**See also:** [dec2hex], page 104, [base2dec], page 106, [bin2dec], page 103.

*str* = dec2base (*d*, *base*)
*str* = dec2base (*d*, *base*, *len*)
*str* = dec2base (*d*, *base*, *len*, *decimals*)

Return a string of symbols in base *base* corresponding to the value *d*.

```
dec2base (123, 3)
   ⇒ "11120"
```

If *d* is negative, then the result will represent *d* in complement notation. For example, negative binary numbers are in twos-complement, and analogously for other bases.

If *d* is a matrix or cell array, return a string matrix with one row per element in *d*, padded with leading zeros to the width of the largest value.

If *base* is a string then the characters of *base* are used as the symbols for the digits of *d*. Whitespace (spaces, tabs, newlines, , etc.) may not be used as a symbol.

```
dec2base (123, "aei")
   ⇒ "eeeia"
```

The optional third argument, *len*, specifies the minimum number of digits in the integer part of the result. If this is omitted, then `dec2base` uses enough digits to accommodate the input.

The optional fourth argument, *decimals*, specifies the number of digits to represent the fractional part of the input. If this is omitted, then it is set to zero, and `dec2base` returns an integer output for backward compatibility.

```
dec2base (100*pi, 16)
⇒ "13A"
dec2base (100*pi, 16, 4)
⇒ "013A"
dec2base (100*pi, 16, 4, 6)
⇒ "013A.28C59D"
dec2base (-100*pi, 16)
⇒ "EC6"
dec2base (-100*pi, 16, 4)
⇒ "FEC6"
dec2base (-100*pi, 16, 4, 6)
⇒ "FEC5.D73A63"
```

Programming tip: When passing negative inputs to `dec2base`, it is best to explicitly specify the length of the output required.

**See also:** [base2dec], page 106, [dec2bin], page 103, [dec2hex], page 104.

`d = base2dec (str, base)`

Convert *str* from a string of digits in base *base* to a decimal integer (base 10).

```
base2dec ("11120", 3)
   ⇒ 123
```

If *str* is a string matrix, return a column vector with one value per row of *str*. If a row contains invalid symbols then the corresponding value will be NaN.

If *str* is a cell array of strings, return a column vector with one value per cell element in *str*.

If *base* is a string, the characters of *base* are used as the symbols for the digits of *str*. Space (' ') may not be used as a symbol.

```
base2dec ("yyyzx", "xyz")
   ⇒ 123
```

**See also:** [dec2base], page 105, [bin2dec], page 103, [hex2dec], page 104.

`s = num2hex (n)`
`s = num2hex (n, "cell")`

Convert a numeric array to an array of hexadecimal strings.

For example:

```
num2hex ([-1, 1, e, Inf])
⇒ "bff0000000000000
    3ff0000000000000
    4005bf0a8b145769
    7ff0000000000000"
```

If the argument *n* is a single precision number or vector, the returned string has a length of 8. For example:

```
num2hex (single ([-1, 1, e, Inf]))
⇒ "bf800000
    3f800000
    402df854
    7f800000"
```

With the optional second argument `"cell"`, return a cell array of strings instead of a character array.

**See also:** [hex2num], page 106, [hex2dec], page 104, [dec2hex], page 104.

`n = hex2num (s)`
`n = hex2num (s, class)`

Typecast a hexadecimal character array or cell array of strings to an array of numbers.

By default, the input array is interpreted as a hexadecimal number representing a double precision value. If fewer than 16 characters are given the strings are right padded with `'0'` characters.

Given a string matrix, `hex2num` treats each row as a separate number.

```
hex2num (["4005bf0a8b145769"; "4024000000000000"])
   ⇒ [2.7183; 10.000]
```

The optional second argument *class* may be used to cause the input array to be interpreted as a different value type. Possible values are

| **Option** | **Characters** |
|---|---|
| `"int8"` | 2 |
| `"uint8"` | 2 |
| `"int16"` | 4 |
| `"uint16"` | 4 |
| `"int32"` | 8 |
| `"uint32"` | 8 |
| `"int64"` | 16 |
| `"uint64"` | 16 |
| `"char"` | 2 |
| `"single"` | 8 |
| `"double"` | 16 |

For example:

```
hex2num (["402df854"; "41200000"], "single")
   ⇒ [2.7183; 10.000]
```

**See also:** [num2hex], page 106, [hex2dec], page 104, [dec2hex], page 104.

`[a, ...] = strread (`*`str`*`)`
`[a, ...] = strread (`*`str`*`, `*`format`*`)`
`[a, ...] = strread (`*`str`*`, `*`format`*`, `*`format_repeat`*`)`
`[a, ...] = strread (`*`str`*`, `*`format`*`, `*`prop1`*`, `*`value1`*`, ...)`
`[a, ...] = strread (`*`str`*`, `*`format`*`, `*`format_repeat`*`, `*`prop1`*`, `*`value1`*`, ...)`

This function is obsolete. Use `textscan` instead.

Read data from a string.

The string *str* is split into words that are repeatedly matched to the specifiers in *format*. The first word is matched to the first specifier, the second to the second specifier and so forth. If there are more words than specifiers, the process is repeated until all words have been processed.

The string *format* describes how the words in *str* should be parsed. It may contain any combination of the following specifiers:

`%s` The word is parsed as a string.

`%f`
`%n` The word is parsed as a number and converted to double.

`%d`
`%u` The word is parsed as a number and converted to int32.

`%*`
`%*f`
`%*s` The word is skipped.

For %s and %d, %f, %n, %u and the associated %*s ... specifiers an optional width can be specified as %Ns, etc. where N is an integer > 1. For %f, format specifiers like %N.Mf are allowed.

`literals` In addition the format may contain literal character strings; these will be skipped during reading.

Parsed word corresponding to the first specifier are returned in the first output argument and likewise for the rest of the specifiers.

By default, *format* is `"%f"`, meaning that numbers are read from *str*. This will do if *str* contains only numeric fields.

For example, the string

```
str = "\
Bunny Bugs   5.5\n\
Duck Daffy  -7.5e-5\n\
Penguin Tux   6"
```

can be read using

```
[a, b, c] = strread (str, "%s %s %f");
```

Optional numeric argument *format_repeat* can be used for limiting the number of items read:

-1 (default) read all of the string until the end.

N Read N times *nargout* items. 0 (zero) is an acceptable value for *format_repeat*.

The behavior of `strread` can be changed via property-value pairs. The following properties are recognized:

`"commentstyle"`
: Parts of *str* are considered comments and will be skipped. *value* is the comment style and can be any of the following.

  - `"shell"` Everything from `#` characters to the nearest end-of-line is skipped.
  - `"c"` Everything between `/*` and `*/` is skipped.
  - `"c++"` Everything from `//` characters to the nearest end-of-line is skipped.
  - `"matlab"` Everything from `%` characters to the nearest end-of-line is skipped.
  - user-supplied. Two options: (1) One string, or 1x1 cell string: Skip everything to the right of it; (2) 2x1 cell string array: Everything between the left and right strings is skipped.

`"delimiter"`
: Any character in *value* will be used to split *str* into words (default value = any whitespace). Note that whitespace is implicitly added to the set of delimiter characters unless a `"%s"` format conversion specifier is supplied; see `"whitespace"` parameter below. The set of delimiter characters cannot be empty; if needed Octave substitutes a space as delimiter.

`"emptyvalue"`
: Value to return for empty numeric values in non-whitespace delimited data. The default is NaN. When the data type does not support NaN (int32 for example), then default is zero.

`"multipledelimsasone"`
: Treat a series of consecutive delimiters, without whitespace in between, as a single delimiter. Consecutive delimiter series need not be vertically `"aligned"`.

`"treatasempty"`
: Treat single occurrences (surrounded by delimiters or whitespace) of the string(s) in *value* as missing values.

`"returnonerror"`
: If *value* true (1, default), ignore read errors and return normally. If false (0), return an error.

`"whitespace"`
: Any character in *value* will be interpreted as whitespace and trimmed; the string defining whitespace must be enclosed in double quotes for proper processing of special characters like `"\t"`. In each data field, multiple consecutive whitespace characters are collapsed into one space and leading and trailing whitespace is removed. The default value for whitespace is `" \b\r\n\t"` (note the space). Whitespace is always added to the set of delimiter characters unless at least one `"%s"` format conversion specifier is supplied; in that case only whitespace explicitly specified in `"delimiter"` is retained as delimiter and removed from the set of whitespace characters. If whitespace characters are to be kept as-is (in e.g., strings), specify an empty value (i.e., `""`) for `"whitespace"`; obviously, whitespace cannot be a delimiter then.

When the number of words in *str* doesn't match an exact multiple of the number of format conversion specifiers, strread's behavior depends on the last character of *str*:

last character = `"\n"`
: Data columns are padded with empty fields or NaN so that all columns have equal length

last character is not `"\n"`
: Data columns are not padded; strread returns columns of unequal length

**See also:** [textscan], page 304, [sscanf], page 322.

### 5.4.3 JSON data encoding/decoding

JavaScript Object Notation, in short JSON, is a very common human readable and structured data format. GNU Octave supports encoding and decoding this format with the following two functions.

*JSON_txt* = `jsonencode` (*object*)
*JSON_txt* = `jsonencode` (..., "*ConvertInfAndNaN*", *TF*)
*JSON_txt* = `jsonencode` (..., "*PrettyPrint*", *TF*)
: Encode Octave data types into JSON text.

The input *object* is an Octave variable to encode.

The output *JSON_txt* is the JSON text that contains the result of encoding *object*.

If the value of the option `"ConvertInfAndNaN"` is true then `NaN`, `NA`, `-Inf`, and `Inf` values will be converted to `"null"` in the output. If it is false then they will remain as their original values. The default value for this option is true.

If the value of the option `"PrettyPrint"` is true, the output text will have indentations and line feeds. If it is false, the output will be condensed and written without whitespace. The default value for this option is false.

Programming Notes:

- Complex numbers are not supported.
- classdef objects are first converted to structs and then encoded.
- To preserve escape characters (e.g., `"\n"`), use single-quoted strings.
- Every character after the null character (`"\0"`) in a double-quoted string will be dropped during encoding.
- Encoding and decoding an array is not guaranteed to preserve the dimensions of the array. In particular, row vectors will be reshaped to column vectors.
- Encoding and decoding is not guaranteed to preserve the Octave data type because JSON supports fewer data types than Octave. For example, if you encode an `int8` and then decode it, you will get a `double`.

This table shows the conversions from Octave data types to JSON data types:

| **Octave data type** | **JSON data type** |
|---|---|
| logical scalar | Boolean |
| logical vector | Array of Boolean, reshaped to row vector |
| logical array | nested Array of Boolean |
| numeric scalar | Number |
| numeric vector | Array of Number, reshaped to row vector |
| numeric array | nested Array of Number |
| `NaN`, `NA`, `Inf`, `-Inf` when `"ConvertInfAndNaN" = true` | `"null"` |
| `NaN`, `NA`, `Inf`, `-Inf` when `"ConvertInfAndNaN" = false` | `"NaN"`, `"NaN"`, `"Infinity"`, `"-Infinity"` |
| empty array | `"[]"` |
| character vector | String |
| character array | Array of String |
| empty character array | `""` |
| cell scalar | Array |
| cell vector | Array, reshaped to row vector |
| cell array | Array, flattened to row vector |
| struct scalar | Object |
| struct vector | Array of Object, reshaped to row vector |
| struct array | nested Array of Object |
| classdef object | Object |

Examples:

```
jsonencode ([1, NaN; 3, 4])
⇒ [[1,null],[3,4]]

jsonencode ([1, NaN; 3, 4], "ConvertInfAndNaN", false)
⇒ [[1,NaN],[3,4]]

## Escape characters inside a single-quoted string
jsonencode ('\0\a\b\t\n\v\f\r')
⇒ "\\0\\a\\b\\t\\n\\v\\f\\r"

## Escape characters inside a double-quoted string
jsonencode ("\a\b\t\n\v\f\r")
⇒ "\u0007\b\t\n\u000B\f\r"

jsonencode ([true; false], "PrettyPrint", true)
⇒ ans = [
     true,
     false
   ]

jsonencode (['foo', 'bar'; 'foo', 'bar'])
⇒ ["foobar","foobar"]

jsonencode (struct ('a', Inf, 'b', [], 'c', struct ()))
⇒ {"a":null,"b":[],"c":{}}

jsonencode (struct ('structarray', struct ('a', {1; 3}, 'b', {2; 4})))
⇒ {"structarray":[{"a":1,"b":2},{"a":3,"b":4}]}

jsonencode ({'foo'; 'bar'; {'foo'; 'bar'}})
⇒ ["foo","bar",["foo","bar"]]

jsonencode (containers.Map({'foo'; 'bar'; 'baz'}, [1, 2, 3]))
⇒ {"bar":2,"baz":3,"foo":1}
```

**See also:** [jsondecode], page 111.

*object* = jsondecode (*JSON_txt*)
*object* = jsondecode (*...*, "*ReplacementStyle*", *rs*)
*object* = jsondecode (*...*, "*Prefix*", *pfx*)
*object* = jsondecode (*...*, "*makeValidName*", *TF*)

Decode text that is formatted in JSON.

The input *JSON_txt* is a string that contains JSON text.

The output *object* is an Octave object that contains the result of decoding *JSON_txt*.

For more information about the options `"ReplacementStyle"` and `"Prefix"`, see [matlab.lang.makeValidName], page 145.

If the value of the option `"makeValidName"` is false then names will not be changed by `matlab.lang.makeValidName` and the `"ReplacementStyle"` and `"Prefix"` options will be ignored.

NOTE: Decoding and encoding JSON text is not guaranteed to reproduce the original text as some names may be changed by `matlab.lang.makeValidName`.

This table shows the conversions from JSON data types to Octave data types:

| **JSON data type** | **Octave data type** |
|---|---|
| Boolean | scalar logical |
| Number | scalar double |
| String | vector of characters |
| Object | scalar struct (field names of the struct may be different from the keys of the JSON object due to `matlab_lang_makeValidName` |
| null, inside a numeric array | `NaN` |
| null, inside a non-numeric array | empty double array `[]` |
| Array, of different data types | cell array |
| Array, of Booleans | logical array |
| Array, of Numbers | double array |
| Array, of Strings | cell array of character vectors (`cellstr`) |
| Array of Objects, same field names | struct array |
| Array of Objects, different field names | cell array of scalar structs |

Examples:

```
jsondecode ('[1, 2, null, 3]')
    ⇒ ans =

      1
      2
    NaN
      3

jsondecode ('["foo", "bar", ["foo", "bar"]]')
    ⇒ ans =
       {
         [1,1] = foo
         [2,1] = bar
         [3,1] =
         {
           [1,1] = foo
           [2,1] = bar
         }

       }

jsondecode ('{"nu#m#ber": 7, "s#tr#ing": "hi"}', ...
            'ReplacementStyle', 'delete')
    ⇒ scalar structure containing the fields:

         number = 7
         string = hi
```

```
jsondecode ('{"nu#m#ber": 7, "s#tr#ing": "hi"}', ...
            'makeValidName', false)
    ⇒ scalar structure containing the fields:

        nu#m#ber = 7
        s#tr#ing = hi

jsondecode ('{"1": "one", "2": "two"}', 'Prefix', 'm_')
    ⇒ scalar structure containing the fields:

        m_1 = one
        m_2 = two
```

**See also:** [jsonencode], page 109, [matlab.lang.makeValidName], page 145.

## 5.5 Character Class Functions

Octave also provides the following character class test functions patterned after the functions in the standard C library. They all operate on string arrays and return matrices of zeros and ones. Elements that are nonzero indicate that the condition was true for the corresponding character in the string array. For example:

```
isalpha ("!Q@WERT^Y&")
    ⇒ [ 0, 1, 0, 1, 1, 1, 1, 0, 1, 0 ]
```

*tf* = `isalnum` (*s*)

Return a logical array which is true where the elements of *s* are letters or digits and false where they are not.

This is equivalent to (`isalpha (s) | isdigit (s)`).

**See also:** [isalpha], page 113, [isdigit], page 114, [ispunct], page 114, [isspace], page 114, [iscntrl], page 114.

*tf* = `isalpha` (*s*)

Return a logical array which is true where the elements of *s* are letters and false where they are not.

This is equivalent to (`islower (s) | isupper (s)`).

**See also:** [isdigit], page 114, [ispunct], page 114, [isspace], page 114, [iscntrl], page 114, [isalnum], page 113, [islower], page 113, [isupper], page 114.

*tf* = `isletter` (*s*)

Return a logical array which is true where the elements of *s* are letters and false where they are not.

This is an alias for the `isalpha` function.

**See also:** [isalpha], page 113, [isdigit], page 114, [ispunct], page 114, [isspace], page 114, [iscntrl], page 114, [isalnum], page 113.

*tf* = `islower` (*s*)

Return a logical array which is true where the elements of *s* are lowercase letters and false where they are not.

**See also:** [isupper], page 114, [isalpha], page 113, [isletter], page 113, [isalnum], page 113.

*tf* = `isupper` (*s*)

Return a logical array which is true where the elements of *s* are uppercase letters and false where they are not.

**See also:** [islower], page 113, [isalpha], page 113, [isletter], page 113, [isalnum], page 113.

*tf* = `isdigit` (*s*)

Return a logical array which is true where the elements of *s* are decimal digits (0-9) and false where they are not.

**See also:** [isxdigit], page 114, [isalpha], page 113, [isletter], page 113, [ispunct], page 114, [isspace], page 114, [iscntrl], page 114.

*tf* = `isxdigit` (*s*)

Return a logical array which is true where the elements of *s* are hexadecimal digits (0-9 and a-fA-F).

**See also:** [isdigit], page 114.

*tf* = `ispunct` (*s*)

Return a logical array which is true where the elements of *s* are punctuation characters and false where they are not.

**See also:** [isalpha], page 113, [isdigit], page 114, [isspace], page 114, [iscntrl], page 114.

*tf* = `isspace` (*s*)

Return a logical array which is true where the elements of *s* are whitespace characters (space, formfeed, newline, carriage return, tab, and vertical tab) and false where they are not.

**See also:** [iscntrl], page 114, [ispunct], page 114, [isalpha], page 113, [isdigit], page 114.

*tf* = `iscntrl` (*s*)

Return a logical array which is true where the elements of *s* are control characters and false where they are not.

**See also:** [ispunct], page 114, [isspace], page 114, [isalpha], page 113, [isdigit], page 114.

*tf* = `isgraph` (*s*)

Return a logical array which is true where the elements of *s* are printable characters (but not the space character) and false where they are not.

**See also:** [isprint], page 114.

*tf* = `isprint` (*s*)

Return a logical array which is true where the elements of *s* are printable characters (including the space character) and false where they are not.

**See also:** [isgraph], page 114.

*tf* = `isascii` (*s*)
: Return a logical array which is true where the elements of *s* are ASCII characters (in the range 0 to 127 decimal) and false where they are not.

*tf* = `isstrprop` (*str*, *prop*)
*tf* = `isstrprop` (*str*, *prop*, *'ForceCellOutput'*, *flag*)
: Test character string properties.

For example:

```
isstrprop ("abc123", "alpha")
⇒ [1, 1, 1, 0, 0, 0]
```

If *str* is a cell array, `isstrpop` is applied recursively to each element of the cell array.

Numeric arrays are converted to character strings.

The second argument *prop* must be one of

`"alpha"` True for characters that are alphabetic (letters).

`"alnum"`
`"alphanum"`
True for characters that are alphabetic or digits.

`"lower"` True for lowercase letters.

`"upper"` True for uppercase letters.

`"digit"` True for decimal digits (0-9).

`"xdigit"` True for hexadecimal digits (a-fA-F0-9).

`"space"`
`"wspace"` True for whitespace characters (space, formfeed, newline, carriage return, tab, vertical tab).

`"punct"` True for punctuation characters (printing characters except space or letter or digit).

`"cntrl"` True for control characters.

`"graph"`
`"graphic"`
True for printing characters except space.

`"print"` True for printing characters including space.

`"ascii"` True for characters that are in the range of ASCII encoding.

If the option `'ForceCellOutput'` is given and *flag* is true then a cell value is returned rather than a logical array.

**See also:** [isalpha], page 113, [isalnum], page 113, [islower], page 113, [isupper], page 114, [isdigit], page 114, [isxdigit], page 114, [isspace], page 114, [ispunct], page 114, [iscntrl], page 114, [isgraph], page 114, [isprint], page 114, [isascii], page 115.

# 6 Data Containers

Octave includes support for three different mechanisms to contain arbitrary data types in the same variable: Structures, which are C-like, and are indexed with named fields; containers.Map objects, which store data in key/value pairs; and cell arrays, where each element of the array can have a different data type and or shape. Multiple input arguments and return values of functions are organized as another data container, the comma-separated list.

## 6.1 Structures

Octave includes support for organizing data in structures. The current implementation uses an associative array with indices limited to strings, but the syntax is more like C-style structures.

### 6.1.1 Basic Usage and Examples

Here are some examples of using data structures in Octave.

Elements of structures can be of any value type. For example, the three expressions

```
x.a = 1;
x.b = [1, 2; 3, 4];
x.c = "string";
```

create a structure with three elements. The '`.`' character separates the structure name (in the example above `x`) from the field name and indicates to Octave that this variable is a structure. To print the value of the structure you can type its name, just as for any other variable:

```
x
     ⇒ x =

         scalar structure containing the fields:

           a =  1
           b =

              1   2
              3   4

           c = string
```

Note that Octave may print the elements in any order.

Structures may be copied just like any other variable:

```
y = x
     ⇒ y =

         scalar structure containing the fields:

           a =  1
           b =

              1   2
              3   4

           c = string
```

Since structures are themselves values, structure elements may reference other structures, as well. The following statement adds the field `d` to the structure `x`. The value of field `d` is itself a data structure containing the single field `a`, which has a value of 3.

```
x.d.a = 3;
x.d
     ⇒ ans =

         scalar structure containing the fields:

           a =  3

x
     ⇒ x =

         scalar structure containing the fields:

           a =  1
           b =

              1   2
              3   4

           c = string
           d =

             scalar structure containing the fields:

               a =  3
```

Note that when Octave prints the value of a structure that contains other structures, only a few levels are displayed. For example:

```
a.b.c.d.e = 1;
a
     ⇒ a =

         scalar structure containing the fields:

           b =

             scalar structure containing the fields:

               c =

                 scalar structure containing the fields:

                   d: 1x1 scalar struct
```

This prevents long and confusing output from large deeply nested structures. The number of levels to print for nested structures may be set with the function `struct_levels_to_print`, and the function `print_struct_array_contents` may be used to enable printing of the contents of structure arrays.

*val* = **struct_levels_to_print** ()
*old_val* = **struct_levels_to_print** (*new_val*)
*old_val* = **struct_levels_to_print** (*new_val*, "*local*")
Query or set the internal variable that specifies the number of structure levels to display.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [print_struct_array_contents], page 119.

*val* = **print_struct_array_contents** ()
*old_val* = **print_struct_array_contents** (*new_val*)
*old_val* = **print_struct_array_contents** (*new_val*, "*local*")
Query or set the internal variable that specifies whether to print struct array contents.

If true, values of struct array elements are printed. This variable does not affect scalar structures whose elements are always printed. In both cases, however, printing will be limited to the number of levels specified by *struct_levels_to_print*.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [struct_levels_to_print], page 119.

Functions can return structures. For example, the following function separates the real and complex parts of a matrix and stores them in two elements of the same structure variable `y`.

```
function y = f (x)
  y.re = real (x);
  y.im = imag (x);
endfunction
```

When called with a complex-valued argument, the function `f` returns the data structure containing the real and imaginary parts of the original function argument.

```
f (rand (2) + rand (2) * I)
     ⇒ ans =

         scalar structure containing the fields:

           re =

              0.040239  0.242160
              0.238081  0.402523

           im =

              0.26475  0.14828
              0.18436  0.83669
```

Function return lists can include structure elements, and they may be indexed like any other variable. For example:

```
[ x.u, x.s(2:3,2:3), x.v ] = svd ([1, 2; 3, 4]);
x

     ⇒ x =

         scalar structure containing the fields:

           u =

             -0.40455  -0.91451
             -0.91451   0.40455

           s =

              0.00000   0.00000   0.00000
              0.00000   5.46499   0.00000
              0.00000   0.00000   0.36597

           v =

             -0.57605   0.81742
             -0.81742  -0.57605
```

It is also possible to cycle through all the elements of a structure in a loop, using a special form of the `for` statement (see Section 10.5.1 [Looping Over Structure Elements], page 199).

### 6.1.2 Structure Arrays

A structure array is a particular instance of a structure, where each of the fields of the structure is represented by a cell array. Each of these cell arrays has the same dimensions. Conceptually, a structure array can also be seen as an array of structures with identical fields. An example of the creation of a structure array is

```
x(1).a = "string1";
x(2).a = "string2";
x(1).b = 1;
x(2).b = 2;
```

which creates a 1-by-2 structure array with two fields. Another way to create a structure array is with the `struct` function (see Section 6.1.3 [Creating Structures], page 122). As previously, to print the value of the structure array, you can type its name:

```
x
     ⇒ x =
       {
         1x2 struct array containing the fields:

           a
           b
       }
```

Individual elements of the structure array can be returned by indexing the variable like `x(1)`, which returns a structure with two fields:

```
x(1)
     ⇒ ans =
       {
         a = string1
         b =  1
       }
```

Furthermore, the structure array can return a comma-separated list of field values (see Section 6.4 [Comma-Separated Lists], page 141), if indexed by one of its own field names. For example:

```
x.a
     ⇒
       ans = string1
       ans = string2
```

Here is another example, using this comma-separated list on the left-hand side of an assignment:

```
[x.a] = deal ("new string1", "new string2");
 x(1).a
     ⇒ ans = new string1
 x(2).a
     ⇒ ans = new string2
```

Just as for numerical arrays, it is possible to use vectors as indices (see Section 8.1 [Index Expressions], page 159):

```
x(3:4) = x(1:2);
[x([1,3]).a] = deal ("other string1", "other string2");
x.a
     ⇒
        ans = other string1
        ans = new string2
        ans = other string2
        ans = new string2
```

The function `size` will return the size of the structure. For the example above

```
size (x)
     ⇒ ans =

          1   4
```

Elements can be deleted from a structure array in a similar manner to a numerical array, by assigning the elements to an empty matrix. For example

```
in = struct ("call1", {x, Inf, "last"},
             "call2", {x, Inf, "first"})
     ⇒ in =
        {
          1x3 struct array containing the fields:

            call1
            call2
        }

in(1) = [];
in.call1
     ⇒
       ans = Inf
       ans = last
```

### 6.1.3 Creating Structures

Besides the index operator `"."`, Octave can use dynamic naming `"(var)"` or the `struct` function to create structures. Dynamic naming uses the string value of a variable as the field name. For example:

```
a = "field2";
x.a = 1;
x.(a) = 2;
x
     ⇒ x =
        {
          a =  1
          field2 =  2
        }
```

Dynamic indexing also allows you to use arbitrary strings, not merely valid Octave identifiers (note that this does not work on MATLAB):

```
a = "long field with spaces (and funny char$)";
x.a = 1;
x.(a) = 2;
x
     ⇒ x =
        {
          a =  1
          long field with spaces (and funny char$) =  2
        }
```

The warning id `Octave:language-extension` can be enabled to warn about this usage. See [warning_ids], page 265.

More realistically, all of the functions that operate on strings can be used to build the correct field name before it is entered into the data structure.

```
names = ["Bill"; "Mary"; "John"];
ages  = [37; 26; 31];
for i = 1:rows (names)
  database.(names(i,:)) = ages(i);
endfor
database
     ⇒ database =
        {
          Bill =  37
          Mary =  26
          John =  31
        }
```

The third way to create structures is the `struct` command. `struct` takes pairs of arguments, where the first argument in the pair is the fieldname to include in the structure and the second is a scalar or cell array, representing the values to include in the structure or structure array. For example:

```
struct ("field1", 1, "field2", 2)
⇒ ans =
      {
        field1 =  1
        field2 =  2
      }
```

If the values passed to `struct` are a mix of scalar and cell arrays, then the scalar arguments are expanded to create a structure array with a consistent dimension. For example:

```
s = struct ("field1", {1, "one"}, "field2", {2, "two"},
        "field3", 3);
s.field1
     ⇒
        ans =  1
        ans = one

s.field2
     ⇒
        ans =  2
        ans = two

s.field3
     ⇒
        ans =  3
        ans =  3
```

If you want to create a struct which contains a cell array as an individual field, you must wrap it in another cell array as shown in the following example:

```
struct ("field1", {{1, "one"}}, "field2", 2)
     ⇒ ans =
        {
          field1 =

        {
          [1,1] =  1
          [1,2] = one
        }

          field2 =  2
        }
```

`s = struct ()`
`s = struct (`*field1*, *value1*, *field2*, *value2*, ...`)`
`s = struct (`*obj*`)`

Create a scalar or array structure and initialize its values.

The *field1*, *field2*, ... variables are strings specifying the names of the fields and the *value1*, *value2*, ... variables can be of any type.

If the values are cell arrays, create a structure array and initialize its values. The dimensions of each cell array of values must match. Singleton cells and non-cell values are repeated so that they fill the entire array. If the cells are empty, create an empty structure array with the specified field names.

If the argument is an object, return the underlying struct.

Observe that the syntax is optimized for struct **arrays**. Consider the following examples:

```
struct ("foo", 1)
  ⇒ scalar structure containing the fields:
    foo =  1

struct ("foo", {})
  ⇒ 0x0 struct array containing the fields:
    foo

struct ("foo", { {} })
  ⇒ scalar structure containing the fields:
    foo = {}(0x0)

struct ("foo", {1, 2, 3})
  ⇒ 1x3 struct array containing the fields:
    foo
```

The first case is an ordinary scalar struct—one field, one value. The second produces an empty struct array with one field and no values, since being passed an empty cell array of struct array values. When the value is a cell array containing a single entry, this becomes a scalar struct with that single entry as the value of the field. That single entry happens to be an empty cell array.

Finally, if the value is a nonscalar cell array, then `struct` produces a struct **array**.

**See also:** [cell2struct], page 141, [fieldnames], page 125, [getfield], page 127, [setfield], page 126, [rmfield], page 127, [isfield], page 126, [orderfields], page 127, [isstruct], page 125, [structfun], page 699.

The function `isstruct` can be used to test if an object is a structure or a structure array.

*tf* = `isstruct` (*x*)

Return true if *x* is a structure or a structure array.

**See also:** [ismatrix], page 69, [iscell], page 132, [isa], page 41.

### 6.1.4 Manipulating Structures

Other functions that can manipulate the fields of a structure are given below.

*n* = `numfields` (*s*)

Return the number of fields of the structure *s*.

**See also:** [fieldnames], page 125.

*names* = `fieldnames` (*struct*)
*names* = `fieldnames` (*obj*)
*names* = `fieldnames` (*javaobj*)
*names* = `fieldnames` ("*javaclassname*")

Return a cell array of strings with the names of the fields in the specified input.

When the input is a structure *struct*, the *names* are the elements of the structure.

When the input is an Octave object *obj*, the *names* are the public properties of the object.

When the input is a Java object *javaobj* or a string containing the name of a Java class *javaclassname*, the *names* are the public fields (data members) of the object or class.

**See also:** [numfields], page 125, [isfield], page 126, [orderfields], page 127, [struct], page 124, [properties], page 994.

`tf = isfield (x, "name")`
`tf = isfield (x, name)`

Return true if the *x* is a structure and it includes an element named *name*.

If *name* is a cell array of strings then a logical array of equal dimension is returned.

**See also:** [fieldnames], page 125.

`sout = setfield (s, field, val)`
`sout = setfield (s, sidx1, field1, fidx1, sidx2, field2, fidx2, ..., val)`

Return a *copy* of the structure *s* with the field member *field* set to the value *val*.

For example:

```
s = struct ();
s = setfield (s, "foo bar", 42);
```

This is equivalent to

```
s.("foo bar") = 42;
```

Note that ordinary structure syntax `s.foo bar = 42` cannot be used here, as the field name is not a valid Octave identifier because of the space character. Using arbitrary strings for field names is incompatible with MATLAB, and this usage will emit a warning if the warning ID `Octave:language-extension` is enabled. See [warning_ids], page 265.

With the second calling form, set a field of a structure array. The input *sidx* selects an element of the structure array, *field* specifies the field name of the selected element, and *fidx* selects which element of the field (in the case of an array or cell array). The *sidx*, *field*, and *fidx* inputs can be repeated to address nested structure array elements. The structure array index and field element index must be cell arrays while the field name must be a string.

For example:

```
s = struct ("baz", 42);
setfield (s, {1}, "foo", {1}, "bar", 54)
⇒
  ans =
    scalar structure containing the fields:
      baz =  42
      foo =
        scalar structure containing the fields:
          bar =  54
```

The example begins with an ordinary scalar structure to which a nested scalar structure is added. In all cases, if the structure index *sidx* is not specified it defaults to

1 (scalar structure). Thus, the example above could be written more concisely as `setfield (s, "foo", "bar", 54)`

Finally, an example with nested structure arrays:

```
sa.foo = 1;
sa = setfield (sa, {2}, "bar", {3}, "baz", {1, 4}, 5);
sa(2).bar(3)
⇒
  ans =
    scalar structure containing the fields:
      baz =  0   0   0   5
```

Here *sa* is a structure array whose field at elements 1 and 2 is in turn another structure array whose third element is a simple scalar structure. The terminal scalar structure has a field which contains a matrix value.

Note that the same result as in the above example could be achieved by:

```
sa.foo = 1;
sa(2).bar(3).baz(1,4) = 5
```

**See also:** [getfield], page 127, [rmfield], page 127, [orderfields], page 127, [isfield], page 126, [fieldnames], page 125, [isstruct], page 125, [struct], page 124.

`val = getfield (s, field)`
`val = getfield (s, sidx1, field1, fidx1, ...)`

Get the value of the field named *field* from a structure or nested structure *s*.

If *s* is a structure array then *sidx* selects an element of the structure array, *field* specifies the field name of the selected element, and *fidx* selects which element of the field (in the case of an array or cell array). For a more complete description of the syntax, see [setfield], page 126.

**See also:** [setfield], page 126, [rmfield], page 127, [orderfields], page 127, [isfield], page 126, [fieldnames], page 125, [isstruct], page 125, [struct], page 124.

`sout = rmfield (s, "f")`
`sout = rmfield (s, f)`

Return a *copy* of the structure (array) *s* with the field *f* removed.

If *f* is a cell array of strings or a character array, remove each of the named fields.

**See also:** [orderfields], page 127, [fieldnames], page 125, [isfield], page 126.

`sout = orderfields (s1)`
`sout = orderfields (s1, s2)`
`sout = orderfields (s1, {cellstr})`
`sout = orderfields (s1, p)`
`[sout, p] = orderfields (...)`

Return a *copy* of *s1* with fields arranged alphabetically, or as specified by the second input.

Given one input struct *s1*, arrange field names alphabetically.

If a second struct argument is given, arrange field names in *s1* as they appear in *s2*. The second argument may also specify the order in a cell array of strings *cellstr*. The second argument may also be a permutation vector.

The optional second output argument *p* is the permutation vector which converts the original name order to the new name order.

Examples:

```
s = struct ("d", 4, "b", 2, "a", 1, "c", 3);
t1 = orderfields (s)
  ⇒ t1 =
       scalar structure containing the fields:
         a =  1
         b =  2
         c =  3
         d =  4

t = struct ("d", {}, "c", {}, "b", {}, "a", {});
t2 = orderfields (s, t)
  ⇒ t2 =
       scalar structure containing the fields:
         d =  4
         c =  3
         b =  2
         a =  1

t3 = orderfields (s, [3, 2, 4, 1])
  ⇒ t3 =
       scalar structure containing the fields:
         a =  1
         b =  2
         c =  3
         d =  4

[t4, p] = orderfields (s, {"d", "c", "b", "a"})
  ⇒ t4 =
       scalar structure containing the fields:
         d =  4
         c =  3
         b =  2
         a =  1
    p =
       1
       4
       2
       3
```

**See also:** [fieldnames], page 125, [getfield], page 127, [setfield], page 126, [rmfield], page 127, [isfield], page 126, [isstruct], page 125, [struct], page 124.

*s* = substruct (*type*, *subs*, ...)

Create a subscript structure for use with `subsref` or `subsasgn`.

For example:

```
idx = substruct ("()", {3, ":"})
  ⇒ idx =
       scalar structure containing the fields:
         type = ()
         subs =
         {
           [1,1] =  3
           [1,2] = :
         }
x = [1, 2, 3;
     4, 5, 6;
     7, 8, 9];
subsref (x, idx)
  ⇒   7   8   9
```

Note: The keyword `end` cannot be used within `subsref` or `subsasgn` for indexing assignments.

**See also:** [subsref], page 978, [subsasgn], page 980.

### 6.1.5 Processing Data in Structures

The simplest way to process data in a structure is within a `for` loop (see Section 10.5.1 [Looping Over Structure Elements], page 199). A similar effect can be achieved with the `structfun` function, where a user defined function is applied to each field of the structure. See [structfun], page 699.

Alternatively, to process the data in a structure, the structure might be converted to another type of container before being treated.

`c = struct2cell (s)`

Create a new cell array from the objects stored in the struct object.

If $f$ is the number of fields in the structure, the resulting cell array will have a dimension vector corresponding to `[f size(s)]`. For example:

```
s = struct ("name", {"Peter", "Hannah", "Robert"},
            "age", {23, 16, 3});
c = struct2cell (s)
   ⇒ c = {2x1x3 Cell Array}
c(1,1,:)(:)
   ⇒
      {
        [1,1] = Peter
        [2,1] = Hannah
        [3,1] = Robert
      }
c(2,1,:)(:)
   ⇒
      {
        [1,1] = 23
        [2,1] = 16
        [3,1] = 3
      }
```

**See also:** [cell2struct], page 141, [namedargs2cell], page 130, [fieldnames], page 125.

`c = namedargs2cell (s)`

Create a cell array of field name/value pairs from a scalar structure.

Example:

```
s.Name = "Peter";
s.Height = 185;
s.Age = 42;

c = namedargs2cell (s)
  ⇒ { "Name", "Peter", "Height", 185, "Age", 42 }
```

**See also:** [struct2cell], page 129.

## 6.2 containers.Map

`m = containers.Map ()`
`m = containers.Map (keys, vals)`
`m = containers.Map (keys, vals, "UniformValues", is_uniform)`
`m = containers.Map ("KeyType", kt, "ValueType", vt)`

Create an object of the containers.Map class that stores a list of key/value pairs.

*keys* is an array of *unique* keys for the map. The keys can be numeric scalars or strings. The type for numeric keys may be one of `"double"`, `"single"`, `"int32"`, `"uint32"`, `"int64"`, or `"uint64"`. Other numeric or logical keys will be converted to `"double"`. A single string key may be entered as is. Multiple string keys are entered as a cell array of strings.

*vals* is an array of values for the map with the *same* number of elements as *keys*.

When called with no input arguments, a default map is created with strings as the key type and `"any"` as the value type.

The `"UniformValues"` option specifies whether the values of the map must be strictly of the same type. If *is_uniform* is true, any values which would be added to the map are first validated to ensure they are of the correct type.

When called with `"KeyType"` and `"ValueType"` arguments, create an empty map with the specified types. The inputs *kt* and *vt* are the types for the keys and values of the map, respectively. Allowed values for *kt* are `"char"`, `"double"`, `"single"`, `"int32"`, `"uint32"`, `"int64"`, `"uint64"`. Allowed values for *vt* are `"any"`, `"char"`, `"double"`, `"single"`, `"int32"`, `"uint32"`, `"int64"`, `"uint64"`, `"logical"`.

The return value *m* is an object of the containers.Map class.

**See also:** [struct], page 124.

## 6.3 Cell Arrays

It can be both necessary and convenient to store several variables of different size or type in one variable. A cell array is a container class able to do just that. In general cell arrays work just like *N*-dimensional arrays with the exception of the use of '`{`' and '`}`' as allocation and indexing operators.

### 6.3.1 Basic Usage of Cell Arrays

As an example, the following code creates a cell array containing a string and a 2-by-2 random matrix

```
c = {"a string", rand(2, 2)};
```

To access the elements of a cell array, it can be indexed with the { and } operators. Thus, the variable created in the previous example can be indexed like this:

```
c{1}
     ⇒ ans = a string
```

As with numerical arrays several elements of a cell array can be extracted by indexing with a vector of indexes

```
c{1:2}
     ⇒ ans = a string
     ⇒ ans =

               0.593993   0.627732
               0.377037   0.033643
```

The indexing operators can also be used to insert or overwrite elements of a cell array. The following code inserts the scalar 3 on the third place of the previously created cell array

```
c{3} = 3
     ⇒ c =

         {
           [1,1] = a string
           [1,2] =

              0.593993   0.627732
              0.377037   0.033643

           [1,3] =  3
         }
```

Details on indexing cell arrays are explained in Section 6.3.3 [Indexing Cell Arrays], page 137.

In general nested cell arrays are displayed hierarchically as in the previous example. In some circumstances it makes sense to reference them by their index, and this can be performed by the `celldisp` function.

`celldisp` (*c*)
`celldisp` (*c*, *name*)

Recursively display the contents of a cell array.

By default the values are displayed with the name of the variable *c*. However, this name can be replaced with the variable *name*. For example:

```
c = {1, 2, {31, 32}};
celldisp (c, "b")
   ⇒
      b{1} =
       1
      b{2} =
       2
      b{3}{1} =
       31
      b{3}{2} =
       32
```

**See also:** [disp], page 287.

To test if an object is a cell array, use the `iscell` function. For example:

```
iscell (c)
     ⇒ ans = 1

iscell (3)
     ⇒ ans = 0
```

*tf* = `iscell` (*x*)

Return true if *x* is a cell array object.

**See also:** [ismatrix], page 69, [isstruct], page 125, [iscellstr], page 140, [isa], page 41.

### 6.3.2 Creating Cell Arrays

The introductory example (see Section 6.3.1 [Basic Usage of Cell Arrays], page 131) showed how to create a cell array containing currently available variables. In many situations, however, it is useful to create a cell array and then fill it with data.

The `cell` function returns a cell array of a given size, containing empty matrices. This function is similar to the `zeros` function for creating new numerical arrays. The following example creates a 2-by-2 cell array containing empty matrices

```
c = cell (2,2)
     ⇒ c =

         {
           [1,1] = [](0x0)
           [2,1] = [](0x0)
           [1,2] = [](0x0)
           [2,2] = [](0x0)
         }
```

Just like numerical arrays, cell arrays can be multi-dimensional. The `cell` function accepts any number of positive integers to describe the size of the returned cell array. It is also possible to set the size of the cell array through a vector of positive integers. In the following example two cell arrays of equal size are created, and the size of the first one is displayed

```
c1 = cell (3, 4, 5);
c2 = cell ( [3, 4, 5] );
size (c1)
     ⇒ ans =
         3   4   5
```

As can be seen, the [size], page 48, function also works for cell arrays. As do other functions describing the size of an object, such as [length], page 48, [numel], page 47, [rows], page 47, and [columns], page 47.

*C* = `cell` (*n*)
*C* = `cell` (*m*, *n*)
*C* = `cell` (*m*, *n*, *k*, ...)
*C* = `cell` ([*m n* ...])

Create a new cell array object.

If invoked with a single scalar integer argument, return a square NxN cell array. If invoked with two or more scalar integer arguments, or a vector of integer values, return an array with the given dimensions.

**See also:** [cellstr], page 140, [mat2cell], page 135, [num2cell], page 133, [struct2cell], page 129.

As an alternative to creating empty cell arrays, and then filling them, it is possible to convert numerical arrays into cell arrays using the `num2cell`, `mat2cell` and `cellslices` functions.

*C* = num2cell (*A*)
*C* = num2cell (*A*, *dim*)

Convert the numeric matrix *A* to a cell array.

When no *dim* is specified, each element of *A* becomes a 1x1 element in the output *C*.

If *dim* is defined then individual elements of *C* contain all of the elements from *A* along the specified dimension. *dim* may also be a vector of dimensions with the same rule applied.

For example:

```
x = [1,2;3,4]
⇒
    1    2
    3    4

## each element of A becomes a 1x1 element of C
num2cell (x)
   ⇒
      {
        [1,1] =  1
        [2,1] =  3
        [1,2] =  2
        [2,2] =  4
      }
## all rows (dim 1) of A appear in each element of C
num2cell (x, 1)
   ⇒
      {
        [1,1] =
           1
           3
        [1,2] =
           2
           4
      }
## all columns (dim 2) of A appear in each element of C
num2cell (x, 2)
   ⇒
      {
        [1,1] =
           1   2
        [2,1] =
           3   4
      }
## all rows and cols appear in each element of C
## (hence, only 1 output)
num2cell (x, [1, 2])
   ⇒
```

```
{
  [1,1] =
     1   2
     3   4
}
```

**See also:** [mat2cell], page 135.

`C = mat2cell (A, dim1, dim2, ..., dimi, ..., dimn)`
`C = mat2cell (A, rowdim)`

Convert the matrix *A* to a cell array *C*.

Each dimension argument (*dim1*, *dim2*, etc.) is a vector of integers which specifies how to divide that dimension's elements amongst the new elements in the output *C*. The number of elements in the *i*-th dimension is `size (A, i)`. Because all elements in *A* must be partitioned, there is a requirement that `sum (dimi) == size (A, i)`. The size of the output cell *C* is numel (*dim1*) x numel (*dim2*) x ... x numel (*dimn*).

Given a single dimensional argument, *rowdim*, the output is divided into rows as specified. All other dimensions are not divided and thus all columns (dim 2), pages (dim 3), etc. appear in each output element.

Examples

```
x = reshape (1:12, [3, 4])'
⇒
    1    2    3
    4    5    6
    7    8    9
   10   11   12
```

```
## The 4 rows (dim1) are divided in to two cell elements
## with 2 rows each.
## The 3 cols (dim2) are divided in to three cell elements
## with 1 col each.
mat2cell (x, [2,2], [1,1,1])
⇒
{
  [1,1] =

     1
     4

  [2,1] =

     7
    10

  [1,2] =

     2
     5

  [2,2] =
     8
    11

  [1,3] =

     3
     6

  [2,3] =
     9
    12
}
```

```
## The 4 rows (dim1) are divided in to two cell elements
## with a 3/1 split.
## All columns appear in each output element.
mat2cell (x, [3,1])
⇒
{
  [1,1] =

     1   2   3
     4   5   6
     7   8   9

  [2,1] =

     10   11   12
}
```

**See also:** [num2cell], page 133, [cell2mat], page 141.

*sl* = cellslices (*x*, *lb*, *ub*, *dim*)

Given an array *x*, this function produces a cell array of slices from the array determined by the index vectors *lb*, *ub*, for lower and upper bounds, respectively.

In other words, it is equivalent to the following code:

```
n = length (lb);
sl = cell (1, n);
for i = 1:length (lb)
  sl{i} = x(:,...,lb(i):ub(i),...,:);
endfor
```

The position of the index is determined by *dim*. If not specified, slicing is done along the first non-singleton dimension.

**See also:** [cell2mat], page 141, [cellindexmat], page 139, [cellfun], page 697.

### 6.3.3 Indexing Cell Arrays

As shown in see Section 6.3.1 [Basic Usage of Cell Arrays], page 131, elements can be extracted from cell arrays using the '{' and '}' operators. If you want to extract or access subarrays which are still cell arrays, you need to use the '(' and ')' operators. The following example illustrates the difference:

```
c = {"1", "2", "3"; "x", "y", "z"; "4", "5", "6"};
c{2,3}
     ⇒ ans = z

c(2,3)
     ⇒ ans =
       {
         [1,1] = z
       }
```

So with '{}' you access elements of a cell array, while with '()' you access a sub array of a cell array.

Using the '(' and ')' operators, indexing works for cell arrays like for multi-dimensional arrays. As an example, all the rows of the first and third column of a cell array can be set to 0 with the following command:

```
c(:, [1, 3]) = {0}
     ⇒ =
        {
          [1,1] = 0
          [2,1] = 0
          [3,1] = 0
          [1,2] = 2
          [2,2] = y
          [3,2] = 5
          [1,3] = 0
          [2,3] = 0
          [3,3] = 0
        }
```

Note, that the above can also be achieved like this:

```
c(:, [1, 3]) = 0;
```

Here, the scalar '0' is automatically promoted to cell array '{0}' and then assigned to the subarray of c.

To give another example for indexing cell arrays with '()', you can exchange the first and the second row of a cell array as in the following command:

```
c = {1, 2, 3; 4, 5, 6};
c([1, 2], :) = c([2, 1], :)
     ⇒ =
        {
          [1,1] =  4
          [2,1] =  1
          [1,2] =  5
          [2,2] =  2
          [1,3] =  6
          [2,3] =  3
        }
```

Accessing multiple elements of a cell array with the '{' and '}' operators will result in a comma-separated list of all the requested elements (see Section 6.4 [Comma-Separated Lists], page 141). Using the '{' and '}' operators the first two rows in the above example can be swapped back like this:

```
[c{[1,2], :}] = deal (c{[2, 1], :})
     ⇒ =
        {
          [1,1] =  1
          [2,1] =  4
          [1,2] =  2
          [2,2] =  5
          [1,3] =  3
          [2,3] =  6
        }
```

As for struct arrays and numerical arrays, the empty matrix '`[]`' can be used to delete elements from a cell array:

```
x = {"1", "2"; "3", "4"};
x(1, :) = []
     ⇒ x =
        {
          [1,1] = 3
          [1,2] = 4
        }
```

The following example shows how to just remove the contents of cell array elements but not delete the space for them:

```
x = {"1", "2"; "3", "4"};
x(1, :) = {[]}
⇒ x =
      {
        [1,1] = [](0x0)
        [2,1] = 3
        [1,2] = [](0x0)
        [2,2] = 4
      }
```

The indexing operations operate on the cell array and not on the objects within the cell array. By contrast, `cellindexmat` applies matrix indexing to the objects within each cell array entry and returns the requested values.

`y = cellindexmat (x, varargin)`

Perform indexing of matrices in a cell array.

Given a cell array of matrices $x$, this function computes

```
Y = cell (size (X));
for i = 1:numel (X)
  Y{i} = X{i}(varargin{1}, varargin{2}, ..., varargin{N});
endfor
```

The indexing arguments may be scalar (`2`), arrays (`[1, 3]`), ranges (`1:3`), or the colon operator (`":"`). However, the indexing keyword `end` is not available.

**See also:** [cellslices], page 137, [cellfun], page 697.

### 6.3.4 Cell Arrays of Strings

One common use of cell arrays is to store multiple strings in the same variable. It is also possible to store multiple strings in a character matrix by letting each row be a string. This, however, introduces the problem that all strings must be of equal length. Therefore, it is recommended to use cell arrays to store multiple strings. For cases, where the character matrix representation is required for an operation, there are several functions that convert a cell array of strings to a character array and back. `char` and `strvcat` convert cell arrays to a character array (see Section 5.3.2 [Concatenating Strings], page 81), while the function `cellstr` converts a character array to a cell array of strings:

```
a = ["hello"; "world"];
c = cellstr (a)
     ⇒ c =
         {
           [1,1] = hello
           [2,1] = world
         }
```

*cstr* = cellstr (*strmat*)

Create a new cell array object from the elements of the string array *strmat*.

Each row of *strmat* becomes an element of *cstr*. Any trailing spaces in a row are deleted before conversion.

To convert back from a cellstr to a character array use `char`.

**See also:** [cell], page 133, [char], page 83.

One further advantage of using cell arrays to store multiple strings is that most functions for string manipulations included with Octave support this representation. As an example, it is possible to compare one string with many others using the `strcmp` function. If one of the arguments to this function is a string and the other is a cell array of strings, each element of the cell array will be compared to the string argument:

```
c = {"hello", "world"};
strcmp ("hello", c)
     ⇒ ans =
        1   0
```

The following string functions support cell arrays of strings: `char`, `strvcat`, `strcat` (see Section 5.3.2 [Concatenating Strings], page 81), `strcmp`, `strncmp`, `strcmpi`, `strncmpi` (see Section 5.3.4 [Searching in Strings], page 88), `str2double`, `deblank`, `strtrim`, `strtrunc`, `strfind`, `strmatch`, , `regexp`, `regexpi` (see Section 5.3 [String Operations], page 78) and `str2double` (see Section 5.4 [Converting Strings], page 99).

The function `iscellstr` can be used to test if an object is a cell array of strings.

*tf* = iscellstr (*cell*)

Return true if every element of the cell array *cell* is a character string.

**See also:** [ischar], page 77, [isstring], page 77.

### 6.3.5 Processing Data in Cell Arrays

Data that is stored in a cell array can be processed in several ways depending on the actual data. The simplest way to process that data is to iterate through it using one or more `for` loops. The same idea can be implemented more easily through the use of the `cellfun` function that calls a user-specified function on all elements of a cell array. See [cellfun], page 697.

An alternative is to convert the data to a different container, such as a matrix or a data structure. Depending on the data this is possible using the `cell2mat` and `cell2struct` functions.

`m = cell2mat (c)`
: Convert the cell array *c* into a matrix by concatenating all elements of *c* into a hyperrectangle.

  Elements of *c* must be numeric, logical, or char matrices; or cell arrays; or structs; and `cat` must be able to concatenate them together.

  **See also:** [mat2cell], page 135, [num2cell], page 133.

`S = cell2struct (cell, fields)`
`S = cell2struct (cell, fields, dim)`
: Convert *cell* to a structure.

  The number of fields in *fields* must match the number of elements in *cell* along dimension *dim*, that is `numel (fields) == size (cell, dim)`. If *dim* is omitted, a value of 1 is assumed.

```
S = cell2struct ({"Peter", "Hannah", "Robert";
                  185, 170, 168},
                 {"Name","Height"}, 1);
S(1)
   ⇒
      {
        Name   = Peter
        Height = 185
      }
```

  **See also:** [struct2cell], page 129, [cell2mat], page 141, [struct], page 124.

## 6.4 Comma-Separated Lists

Comma-separated lists[1] are the basic argument type to all Octave functions—both for input and return arguments. In the example

```
max (a, b)
```

'`a, b`' is a comma-separated list. Comma-separated lists can appear on both the right and left hand side of an assignment. For example

```
x = [1 0 1 0 0 1 1; 0 0 0 0 0 0 7];
[i, j] = find (x, 2, "last");
```

[1] Comma-separated lists are also sometimes referred to as *cs-lists*.

Here, '`x, 2, "last"`' is a comma-separated list constituting the input arguments of `find`. `find` returns a comma-separated list of output arguments which is assigned element by element to the comma-separated list '`i, j`'.

Another example of where comma-separated lists are used is in the creation of a new array with `[]` (see Section 4.1 [Matrices], page 52) or the creation of a cell array with `{}` (see Section 6.3.1 [Basic Usage of Cell Arrays], page 131). In the expressions

```
a = [1, 2, 3, 4];
c = {4, 5, 6, 7};
```

both '`1, 2, 3, 4`' and '`4, 5, 6, 7`' are comma-separated lists.

Comma-separated lists cannot be directly manipulated by the user. However, both structure arrays and cell arrays can be converted into comma-separated lists, and thus used in place of explicitly written comma-separated lists. This feature is useful in many ways, as will be shown in the following subsections.

### 6.4.1 Comma-Separated Lists Generated from Cell Arrays

As has been mentioned above (see Section 6.3.3 [Indexing Cell Arrays], page 137), elements of a cell array can be extracted into a comma-separated list with the `{` and `}` operators. By surrounding this list with `[` and `]`, it can be concatenated into an array. For example:

```
a = {1, [2, 3], 4, 5, 6};
b = [a{1:4}]
     ⇒ b =
         1   2   3   4   5
```

Similarly, it is possible to create a new cell array containing cell elements selected with `{}`. By surrounding the list with '`{`' and '`}`' a new cell array will be created, as the following example illustrates:

```
a = {1, rand(2, 2), "three"};
b = { a{ [1, 3] } }
     ⇒ b =
         {
           [1,1] =  1
           [1,2] = three
         }
```

Furthermore, cell elements (accessed by `{}`) can be passed directly to a function. The list of elements from the cell array will be passed as an argument list to a given function as if it is called with the elements as individual arguments. The two calls to `printf` in the following example are identical but the latter is simpler and can handle cell arrays of an arbitrary size:

```
c = {"GNU", "Octave", "is", "Free", "Software"};
printf ("%s ", c{1}, c{2}, c{3}, c{4}, c{5});
     ⊣ GNU Octave is Free Software
printf ("%s ", c{:});
     ⊣ GNU Octave is Free Software
```

If used on the left-hand side of an assignment, a comma-separated list generated with `{}` can be assigned to. An example is

```
in{1} = [10, 20, 30];
```

```
in{2} = inf;
in{3} = "last";
in{4} = "first";
out = cell (4, 1);
[out{1:3}] = in{1 : 3};
[out{4:6}] = in{[1, 2, 4]})
     ⇒ out =
        {
           [1,1] =

              10   20   30

           [2,1] = Inf
           [3,1] = last
           [4,1] =

              10   20   30

           [5,1] = Inf
           [6,1] = first
        }
```

### 6.4.2 Comma-Separated Lists Generated from Structure Arrays

Structure arrays can equally be used to create comma-separated lists. This is done by addressing one of the fields of a structure array. For example:

```
x = ceil (randn (10, 1));
in = struct ("call1", {x, 3, "last"},
             "call2", {x, inf, "first"});
out = struct ("call1", cell (2, 1), "call2", cell (2, 1));
[out.call1] = find (in.call1);
[out.call2] = find (in.call2);
```

# 7 Variables

Variables let you give names to values and refer to them later. You have already seen variables in many of the examples. The name of a variable must be a sequence of letters, digits and underscores, but it may not begin with a digit. Octave does not enforce a limit on the length of variable names, but it is seldom useful to have variables with names longer than about 30 characters. The following are all valid variable names

```
x
x15
__foo_bar_baz__
fucnrdthsucngtagdjb
```

However, names like `__foo_bar_baz__` that begin and end with two underscores are understood to be reserved for internal use by Octave. You should not use them in code you write, except to access Octave's documented internal variables and built-in symbolic constants.

Case is significant in variable names. The symbols `a` and `A` are distinct variables.

A variable name is a valid expression by itself. It represents the variable's current value. Variables are given new values with *assignment operators* and *increment operators*. See Section 8.6 [Assignment Expressions], page 179.

There is one automatically created variable with a special meaning. The `ans` variable always contains the result of the last computation, where the output wasn't assigned to any variable. The code `a = cos (pi)` will assign the value -1 to the variable `a`, but will not change the value of `ans`. However, the code `cos (pi)` will set the value of `ans` to -1.

Variables in Octave do not have fixed types, so it is possible to first store a numeric value in a variable and then to later use the same name to hold a string value in the same program. Variables may not be used before they have been given a value. Doing so results in an error.

`ans` [Automatic Variable]

The most recently computed result that was not explicitly assigned to a variable.

For example, after the expression

```
3^2 + 4^2
```

is evaluated, the value returned by `ans` is 25.

*tf* = `isvarname (`*name*`)`

Return true if *name* is a valid variable name.

A valid variable name is composed of letters, digits, and underscores (`"_"`), and the first character must not be a digit.

**See also:** [iskeyword], page 1193, [exist], page 152, [who], page 150.

*varname* = `matlab.lang.makeValidName (`*str*`)`
*varname* = `matlab.lang.makeValidName (...,` `"ReplacementStyle"`, *rs*`)`
*varname* = `matlab.lang.makeValidName (...,` `"Prefix"`, *pfx*`)`
[*varname*, *ismodified*] = `matlab.lang.makeValidName (...)`

Create valid variable name *varname* from *str*.

The input *str* must be a string or a cell array of strings. The output *varname* will be of the same type.

A valid variable name is a sequence of letters, digits, and underscores that does not begin with a digit.

The `"ReplacementStyle"` option specifies how invalid characters are handled. Acceptable values are

`"underscore"` (default)
: Replace all invalid characters with an underscore (`"_"`).

`"delete"`
: Remove any invalid character.

`"hex"`
: Replace all invalid characters with their hexadecimal representation.

Whitespace characters are always removed **prior** to the application of the `"ReplacementStyle"`. Lowercase letters following a whitespace will be changed to uppercase.

The `"Prefix"` option specifies the string *pfx* to add as a prefix to the input if it begins with a digit. *pfx* must be a valid variable name itself. The default prefix is `"x"`.

The optional output *ismodified* is a logical array indicating whether the respective element in *str* was a valid name or not.

**See also:** [iskeyword], page 1193, [isvarname], page 145, [matlab.lang.makeUniqueStrings], page 146.

*uniqstr* = matlab.lang.makeUniqueStrings (*str*)
*uniqstr* = matlab.lang.makeUniqueStrings (*str*, *ex*)
*uniqstr* = matlab.lang.makeUniqueStrings (*str*, *ex*, *maxlength*)
[*uniqstr*, *ismodified*] = matlab.lang.makeUniqueStrings (...)

Construct a list of unique strings from a list of strings.

The input *str* must be a string or a cell array of strings. The output *uniqstr* will be of the same type.

The algorithm makes two strings unique by appending an underscore (`"_"` and a numeric count to the second string.

If *ex* is a string or a cell array of strings, *uniqstr* will contain elements that are unique between themselves and with respect to *ex*.

If *ex* is an index array or a logical array for *str* then it selects the subset of *str* that are made unique. Unselected elements are not modified.

The optional input *maxlength* specifies the maximum length of any string in *uniqstr*. If an input string cannot be made unique without exceeding *maxlength* an error is emitted.

The optional output *ismodified* is a logical array indicating whether each element in *str* was modified to make it unique.

**See also:** [unique], page 871, [matlab.lang.makeValidName], page 145.

*n* = namelengthmax ()

Return the MATLAB compatible maximum variable name length.

Octave is capable of storing strings up to $2^{31} - 1$ in length. However for MATLAB compatibility all variable, function, and structure field names should be shorter than the length returned by `namelengthmax`. In particular, variables stored to a MATLAB file format (`*.mat`) will have their names truncated to this length.

## 7.1 Global Variables

See keyword: [global], page 1193,

A *global* variable is one that may be accessed anywhere within Octave. This is in contrast to a local variable which can only be accessed outside of its current context if it is passed explicitly, such as by including it as a parameter when calling a function (`fcn (local_var1, local_var2)`).

A variable is declared global by using a `global` declaration statement. The following statements are all global declarations.

```
global a
global a b
global c = 2
global d = 3 e f = 5
```

Note that the `global` qualifier extends only to the next end-of-statement indicator which could be a comma (‘`,`’), semicolon (‘`;`’), or newline (‘`'\n'`’). For example, the following code declares one global variable, *a*, and one local variable *b* to which the value 1 is assigned.

```
global a, b = 1
```

A global variable may only be initialized once in a `global` statement. For example, after executing the following code

```
global gvar = 1
global gvar = 2
```

the value of the global variable `gvar` is 1, not 2. Issuing a ‘`clear gvar`’ command does not change the above behavior, but ‘`clear all`’ does.

It is necessary declare a variable as global within a function body in order to access the one universal variable. For example,

```
global x
function f ()
  x = 1;
endfunction
f ()
```

does *not* set the value of the global variable `x` to 1. Instead, a local variable, with name `x`, is created and assigned the value of 1. In order to change the value of the global variable `x`, you must also declare it to be global within the function body, like this

```
function f ()
  global x;
  x = 1;
endfunction
```

Passing a global variable in a function parameter list will make a local copy and *not* modify the global value. For example, given the function

```
function f (x)
  x = 0
endfunction
```

and the definition of `x` as a global variable at the top level,

```
global x = 13
```

the expression

```
f (x)
```

will display the value of `x` from inside the function as 0, but the value of `x` at the top level remains unchanged, because the function works with a *copy* of its argument.

Programming Note: While global variables occasionally are the right solution to a coding problem, modern best practice discourages their use. Code which relies on global variables may behave unpredictably between different users and can be difficult to debug. This is because global variables can introduce systemic changes so that localizing a bug to a particular function, or to a particular loop within a function, becomes difficult.

*tf* = isglobal (*name*)

Return true if *name* is a globally visible variable.

For example:

```
global x
isglobal ("x")
   ⇒ 1
```

**See also:** [isvarname], page 145, [exist], page 152.

## 7.2 Persistent Variables

See keyword: [persistent], page 1193,

A variable that has been declared *persistent* within a function will retain its contents in memory between subsequent calls to the same function. The difference between persistent variables and global variables is that persistent variables are local in scope to a particular function and are not visible elsewhere.

The following example uses a persistent variable to create a function that prints the number of times it has been called.

```
function count_calls ()
  persistent calls = 0;
  printf ("'count_calls' has been called %d times\n",
          ++calls);
endfunction

for i = 1:3
  count_calls ();
endfor

⊣ 'count_calls' has been called 1 times
⊣ 'count_calls' has been called 2 times
⊣ 'count_calls' has been called 3 times
```

As the example shows, a variable may be declared persistent using a `persistent` declaration statement. The following statements are all persistent declarations.

```
persistent a
persistent a b
persistent c = 2
persistent d = 3 e f = 5
```

The behavior of persistent variables is equivalent to the behavior of static variables in C.

One restriction for persistent variables is, that neither input nor output arguments of a function can be persistent:

```
function y = foo ()
  persistent y = 0;  # Not allowed!
endfunction

foo ()
⊣ error: can't make function parameter y persistent
```

Like global variables, a persistent variable may only be initialized once. For example, after executing the following code

```
persistent pvar = 1
persistent pvar = 2
```

the value of the persistent variable `pvar` is 1, not 2.

If a persistent variable is declared but not initialized to a specific value, it will contain an empty matrix. So, it is also possible to initialize a persistent variable by checking whether it is empty, as the following example illustrates.

```
function count_calls ()
  persistent calls;
  if (isempty (calls))
    calls = 0;
  endif
  printf ("'count_calls' has been called %d times\n",
          ++calls);
endfunction
```

This implementation behaves in exactly the same way as the previous implementation of `count_calls`.

The value of a persistent variable is kept in memory until it is explicitly cleared. Assuming that the implementation of `count_calls` is saved on disk, we get the following behavior.

```
for i = 1:2
  count_calls ();
endfor
⊣ 'count_calls' has been called 1 times
⊣ 'count_calls' has been called 2 times

clear
```

```
for i = 1:2
  count_calls ();
endfor
⊣ 'count_calls' has been called 3 times
⊣ 'count_calls' has been called 4 times

clear all
for i = 1:2
  count_calls ();
endfor
⊣ 'count_calls' has been called 1 times
⊣ 'count_calls' has been called 2 times

clear count_calls
for i = 1:2
  count_calls ();
endfor
⊣ 'count_calls' has been called 1 times
⊣ 'count_calls' has been called 2 times
```

That is, the persistent variable is only removed from memory when the function containing the variable is removed. Note that if the function definition is typed directly into the Octave prompt, the persistent variable will be cleared by a simple `clear` command as the entire function definition will be removed from memory. If you do not want a persistent variable to be removed from memory even if the function is cleared, you should use the `mlock` function (see Section 11.10.6 [Function Locking], page 236).

## 7.3 Status of Variables

When creating simple one-shot programs it can be very convenient to see which variables are available at the prompt. The function `who` and its siblings `whos` and `whos_line_format` will show different information about what is in memory, as the following shows.

```
str = "A random string";
who
 ⊣ Variables in the current scope:
 ⊣
 ⊣ ans  str
```

`who`
`who` *pattern ...*
`who` *option pattern ...*
`C = who (...)`

List currently defined variables matching the given patterns.

Valid pattern syntax is the same as described for the `clear` command. If no patterns are supplied, all variables are listed.

By default, only variables visible in the local scope are displayed.

The following are valid options, but may not be combined.

`global` List variables in the global scope rather than the current scope.

`-regexp` The patterns are considered to be regular expressions when matching the variables to display. The same pattern syntax accepted by the `regexp` function is used.

`-file` The next argument is treated as a filename. All variables found within the specified file are listed. No patterns are accepted when reading variables from a file.

If called as a function, return a cell array of defined variable names matching the given patterns.

**See also:** [whos], page 151, [isglobal], page 148, [isvarname], page 145, [exist], page 152, [regexp], page 95.

`whos`
`whos` *pattern ...*
`whos` *option pattern ...*
`S = whos ("`*pattern*`", ...)`

Provide detailed information on currently defined variables matching the given patterns.

Options and pattern syntax are the same as for the `who` command.

Extended information about each variable is summarized in a table with the following default entries.

Attr Attributes of the listed variable. Possible attributes are:

- blank Variable in local scope
- `c` Variable of complex type.
- `f` Formal parameter (function argument).
- `g` Variable with global scope.
- `p` Persistent variable.

Name The name of the variable.

Size The logical size of the variable. A scalar is 1x1, a vector is 1xN or Nx1, a 2-D matrix is MxN.

Bytes The amount of memory currently used to store the variable.

Class The class of the variable. Examples include double, single, char, uint16, cell, and struct.

The table can be customized to display more or less information through the function `whos_line_format`.

If `whos` is called as a function, return a struct array of defined variable names matching the given patterns. Fields in the structure describing each variable are: name, size, bytes, class, global, sparse, complex, nesting, persistent.

**See also:** [who], page 150, [whos_line_format], page 152.

*val* = whos_line_format ()
*old_val* = whos_line_format (*new_val*)
*old_val* = whos_line_format (*new_val*, "*local*")

Query or set the format string used by the command `whos`.

A full format string is:

```
%[modifier]<command>[:width[:left-min[:balance]]];
```

The following command sequences are available:

`%a` Prints attributes of variables (c=complex, s=sparse, f=formal parameter, g=global, p=persistent).

`%b` Prints number of bytes occupied by variables.

`%c` Prints class names of variables.

`%e` Prints elements held by variables.

`%n` Prints variable names.

`%s` Prints dimensions of variables.

`%t` Prints type names of variables.

Every command may also have an alignment modifier:

`l` Left alignment.

`r` Right alignment (default).

`c` Column-aligned (only applicable to command %s).

The `width` parameter is a positive integer specifying the minimum number of columns used for printing. No maximum is needed as the field will auto-expand as required.

The parameters `left-min` and `balance` are only available when the column-aligned modifier is used with the command '`%s`'. `balance` specifies the column number within the field width which will be aligned between entries. Numbering starts from 0 which indicates the leftmost column. `left-min` specifies the minimum field width to the left of the specified balance column.

The default format is:

```
"  %la:5; %ln:6; %cs:16:6:1;  %rb:12;  %lc:-1;\n"
```

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [whos], page 151.

Instead of displaying which variables are in memory, it is possible to determine if a given variable is available. That way it is possible to alter the behavior of a program depending on the existence of a variable. The following example illustrates this.

```
if (! exist ("meaning", "var"))
  disp ("The program has no 'meaning'");
endif
```

`c = exist (name)`
`c = exist (name, type)`

Check for the existence of *name* as a variable, function, file, directory, or class.

The return code *c* is one of

| | |
|---|---|
| 1 | *name* is a variable. |
| 2 | *name* is an absolute filename, an ordinary file in Octave's `path`, or (after appending '`.m`') a function file in Octave's `path`. |
| 3 | *name* is a '`.oct`' or '`.mex`' file in Octave's `path`. |
| 5 | *name* is a built-in function. |
| 7 | *name* is a directory. |
| 8 | *name* is a classdef class. |
| 103 | *name* is a function not associated with a file (entered on the command line). |
| 0 | *name* does not exist. |

If the optional argument *type* is supplied, check only for symbols of the specified type. Valid types are

| | |
|---|---|
| `"var"` | Check only for variables. |
| `"builtin"` | Check only for built-in functions. |
| `"dir"` | Check only for directories. |
| `"file"` | Check only for files and directories. |
| `"class"` | Check only for classdef classes. |

If no type is given, and there are multiple possible matches for name, `exist` will return a code according to the following priority list: variable, built-in function, oct-file, directory, file, class.

`exist` returns 2 if a regular file called *name* is present in Octave's search path. For information about other types of files not on the search path use some combination of the functions `file_in_path` and `stat` instead.

Programming Note: If *name* is implemented by a buggy .oct/.mex file, calling *exist* may cause Octave to crash. To maintain high performance, Octave trusts .oct/.mex files instead of sandboxing them.

**See also:** [file_in_loadpath], page 229, [file_in_path], page 1045, [dir_in_loadpath], page 230, [stat], page 1041.

Usually Octave will manage the memory, but sometimes it can be practical to remove variables from memory manually. This is usually needed when working with large variables that fill a substantial part of the memory. On a computer that uses the IEEE 754 floating point format, the following program allocates a matrix that requires around 128 MB memory.

```
large_matrix = zeros (4000, 4000);
```

Since having this variable in memory might slow down other computations, it can be necessary to remove it manually from memory. The `clear` or `clearvars` functions do this.

`clear`
`clear` *`pattern ...`*
`clear` *`options pattern ...`*

Delete the names matching the given *pattern*s thereby freeing memory.

The *pattern* may contain the following special characters:

`?` Match any single character.

`*` Match zero or more characters.

`[` *`list`* `]` Match the list of characters specified by *list*. If the first character is `!` or `^`, match all characters except those specified by *list*. For example, the pattern `[a-zA-Z]` will match all lowercase and uppercase alphabetic characters. On Windows, square brackets are matched literally and are not used to group characters.

For example, the command

```
clear foo b*r
```

clears the name `foo` and all names that begin with the letter '`b`' and end with the letter '`r`'.

If `clear` is called without any arguments, all user-defined variables are cleared from the current workspace (i.e., local variables). Any global variables present will no longer be visible in the current workspace, but they will continue to exist in the global workspace. Functions are unaffected by this form of `clear`.

The following options are available in both long and short form

`all, -all, -a`
: Clear all local and global user-defined variables, and all functions from the symbol table.

`-exclusive, -x`
: Clear variables that do **not** match the following pattern.

`functions, -functions, -f`
: Clear function names from the function symbol table. Persistent variables will be re-initialized to their default value unless the function has been locked in memory with `mlock`.

`global, -global, -g`
: Clear global variable names.

`variables, -variables, -v`
: Clear local variable names.

`classes, -classes, -c`
: Clear the class structure table and all objects.

`-regexp, -r`
: The *pattern* arguments are treated as regular expressions and any matches will be cleared.

With the exception of `-exclusive` and `-regexp`, all long options can be used without the dash as well. Note that, aside from `-exclusive`, only one other option may appear. All options must appear before any patterns.

Programming Notes: The command `clear` *`name`* only clears the variable *name* when both a variable and a (shadowed) function named *name* are currently defined. For example, suppose you have defined a function `foo`, and then hidden it by performing the assignment `foo = 2`. Executing the command `clear foo` once will clear the variable definition and restore the definition of `foo` as a function. Executing `clear foo` a second time will clear the function definition.

When a local variable name, which is linked to a global variable, is cleared only the local copy of the variable is removed. The global copy is untouched and can be restored with `global` *`global_varname`*. Conversely, `clear -g` *`global_varname`* will remove both the local and global variables.

**See also:** [clearvars], page 155, [who], page 150, [whos], page 151, [exist], page 152, [mlock], page 237.

`clearvars`
`clearvars` *`pattern ...`*
`clearvars` *-regexp* *`pattern ...`*
`clearvars` *`...`* *-except* *`pattern ...`*
`clearvars` *`...`* *-except -regexp* *`pattern ...`*
`clearvars` *-global ...*

Delete the variables matching the given *pattern*s from memory.

The *pattern* may contain the following special characters:

`?` Match any single character.

`*` Match zero or more characters.

`[ list ]` Match the list of characters specified by *list*. If the first character is `!` or `^`, match all characters except those specified by *list*. For example, the pattern `[a-zA-Z]` will match all lowercase and uppercase alphabetic characters.

If the `-regexp` option is given then subsequent patterns are treated as regular expressions and any matches will be cleared.

If the `-except` option is given then subsequent patterns select variables that will **not** be cleared.

If the `-global` option is given then all patterns will be applied to global variables rather than local variables.

When called with no arguments, `clearvars` deletes all local variables.

Example Code:

Clear all variables starting with `'x'` and the specific variable `"foobar"`

```
clearvars x* foobar
```

Clear the specific variable `"foobar"` and use regular expressions to clear all variables starting with `'x'` or `'y'`.

```
clearvars foobar -regexp ^x ^y
```

Clear all variables except for `"foobar"`

```
clearvars -except foobar
```

Clear all variables beginning with `"foo"`, except for those ending in `"bar"`

```
clearvars foo* -except -regexp bar$
```

**See also:** [clear], page 154, [who], page 150, [whos], page 151, [exist], page 152.

`pack ()`

Consolidate workspace memory in MATLAB.

This function is provided for compatibility, but does nothing in Octave.

**See also:** [clear], page 154.

Information about a function or variable such as its location in the file system can also be acquired from within Octave. This is usually only useful during development of programs, and not within a program.

`type` *`name ...`*
`type` *`-q name ...`*
`text = type ("`*`name`*`", ...)`

Display the contents of *name* which may be a file, function (m-file), variable, operator, or keyword.

`type` normally prepends a header line describing the category of *name* such as function or variable; The `-q` option suppresses this behavior.

If no output variable is used the contents are displayed on screen. Otherwise, a cell array of strings is returned, where each element corresponds to the contents of each requested function.

`which` *`name ...`*
`[`*`str`*`, ...] = which (`*`'name'`*`, ...)`

Display the type of each *name*.

If *name* is defined from a function file, the full name of the file is also displayed.

**See also:** [help], page 21, [lookfor], page 22.

`what`
`what` *`dir`*
`w = what (`*`dir`*`)`

List the Octave specific files in directory *dir*.

If *dir* is not specified then the current directory is used.

If a return argument is requested, the files found are returned in the structure *w*. The structure contains the following fields:

| | |
|---|---|
| path | Full path to directory *dir* |
| m | Cell array of m-files |
| mat | Cell array of mat files |
| mex | Cell array of mex files |

| | |
|---|---|
| oct | Cell array of oct files |
| mdl | Cell array of mdl files |
| slx | Cell array of slx files |
| p | Cell array of p-files |
| classes | Cell array of class directories (`@classname/`) |
| packages | Cell array of package directories (`+pkgname/`) |

Compatibility Note: Octave does not support mdl, slx, and p files. `what` will always return an empty list for these categories.

**See also:** [which], page 156, [ls], page 1069, [exist], page 152.

# 8 Expressions

Expressions are the basic building block of statements in Octave. An expression evaluates to a value, which you can print, test, store in a variable, pass to a function, or assign a new value to a variable with an assignment operator.

An expression can serve as a statement on its own. Most other kinds of statements contain one or more expressions which specify data to be operated on. As in other languages, expressions in Octave include variables, array references, constants, and function calls, as well as combinations of these with various operators.

## 8.1 Index Expressions

An *index expression* allows you to reference or extract selected elements of a vector, a matrix (2-D), or a higher-dimensional array. Arrays may be indexed in one of three ways: [Component Indexing], page 159, [Linear Indexing], page 161, and [Logical Indexing], page 161.

### Component Indexing

Component indices may be scalars, vectors, ranges, or the special operator ':', which selects entire rows, columns, or higher-dimensional slices.

Component index expression consists of a set of parentheses enclosing $M$ expressions separated by commas. Each individual index value, or component, is used for the respective dimension of the object that it is applied to. In other words, the first index component is used for the first dimension (rows) of the object, the second index component is used for the second dimension (columns) of the object, and so on. The number of index components $M$ defines the dimensionality of the index expression. An index with two components would be referred to as a 2-D index because it has two dimensions.

In the simplest case, 1) all components are scalars, and 2) the dimensionality of the index expression $M$ is equal to the dimensionality of the object it is applied to. For example:

```
A = reshape (1:8, 2, 2, 2)  # Create 3-D array
A =

ans(:,:,1) =

   1   3
   2   4

ans(:,:,2) =

   5   7
   6   8

A(2, 1, 2)   # second row, first column of second slice
             # in third dimension: ans = 6
```

The size of the returned object in a specific dimension is equal to the number of elements in the corresponding component of the index expression. When all components are scalars,

the result is a single output value. However, if any component is a vector or range then the returned values are the Cartesian product of the indices in the respective dimensions. For example:

```
A([1, 2], 1, 2) ≡ [A(1,1,2); A(2,1,2)]
⇒
ans =
   5
   6
```

The total number of returned values is the product of the number of elements returned for each index component. In the example above, the total is 2*1*1 = 2 elements.

Notice that the size of the returned object in a given dimension is equal to the number of elements in the index expression for that dimension. In the code above, the first index component (`[1, 2]`) was specified as a row vector, but its shape is unimportant. The important fact is that the component specified two values, and hence the result must have a size of two in the first dimension; and because the first dimension corresponds to rows, the overall result is a column vector.

```
A(1, [2, 1, 1], 1)    # result is a row vector: ans = [3, 1, 1]
A(ones (2, 2), 1, 1)  # result is a column vector: ans = [1; 1; 1; 1]
```

The first line demonstrates again that the size of the output in a given dimension is equal to the number of elements in the respective indexing component. In this case, the output has three elements in the second dimension (which corresponds to columns), so the result is a row vector. The example also shows how repeating entries in the index expression can be used to replicate elements in the output. The last example further proves that the shape of the indexing component is irrelevant, it is only the number of elements (2x2 = 4) which is important.

The above rules apply whenever the dimensionality of the index expression is greater than one ($M > 1$). However, for one-dimensional index expressions special rules apply and the shape of the output **is** determined by the shape of the indexing component. For example:

```
A([1, 2])  # result is a row vector: ans = [1, 2]
A([1; 2])  # result is a column vector: ans = [1; 2]
```

The shape rules for $A(P)$ are:

- When at least one of $A$ or $P$ has two or more dimensions, then $A(P)$ takes the shape of $P$. This happens when at least one of the variables is a 2-D matrix or an N-D array.
- When both $A$ and $P$ are 1-D vectors, then $A(P)$ takes the shape of $A$ itself. In particular, when $A$ is a row vector, then $A(P)$ is also a row vector irrespective of $P$'s shape. The case when $A$ is a column vector is analogous.

A colon (':') may be used as an index component to select all of the elements in a specified dimension. Given the matrix,

```
A = [1, 2; 3, 4]
```

all of the following expressions are equivalent and select the first row of the matrix.

```
A(1, [1, 2])  # row 1, columns 1 and 2
A(1, 1:2)     # row 1, columns in range 1-2
A(1, :)       # row 1, all columns
```

When a colon is used in the special case of 1-D indexing the result is always a column vector. Creating column vectors with a colon index is a very frequently encountered code idiom and is faster and generally clearer than calling `reshape` for this case.

```
A(:)    # result is column vector: ans = [1; 2; 3; 4]
A(:)'   # result is row vector: ans = [1, 2, 3, 4]
```

In index expressions the keyword `end` automatically refers to the last entry for a particular dimension. This magic index can also be used in ranges and typically eliminates the needs to call `size` or `length` to gather array bounds before indexing. For example:

```
A(1:end/2)        # first half of A => [1, 2]
A(end + 1) = 5;   # append element
A(end) = [];      # delete element
A(1:2:end)        # odd elements of A => [1, 3]
A(2:2:end)        # even elements of A => [2, 4]
A(end:-1:1)       # reversal of A => [4, 3, 2, 1]
```

For more information see the ["end" keyword], page 1194.

## Linear Indexing

It is permissible to use a 1-D index with a multi-dimensional object. This is also called *linear indexing*. In this case, the elements of the multi-dimensional array are taken in column-first order like in Fortran. That is, the columns of the array are imagined to be stacked on top of each other to form a column vector and then the single linear index is applied to this vector.

```
A = [1, 2, 3; 4, 5, 6; 7, 8, 9];

A(4)    # linear index of 4th element in 2-D array: ans = 2
A(3:5)  # result has shape of index component: ans = [7, 2, 5]
A([1, 2, 2, 1]) # result includes repeated elements: ans = [1, 4, 4, 1]
```

## Logical Indexing

Logical values can also be used to index matrices and cell arrays. When indexing with a logical array the result will be a vector containing the values corresponding to `true` parts of the logical array. The following examples illustrates this.

```
data = [ 1, 2; 3, 4 ];
idx = [true, false; false true];
data(idx)
     ⇒ ans = [ 1; 4 ]

idx = (data <= 2);
data(idx)
     ⇒ ans = [ 1; 2 ]
```

Instead of creating the `idx` array it is possible to replace `data(idx)` with `data( data <= 2 )` in the above code.

While the size of the logical index expressions is usually the same as that of the array being indexed, this is not a necessary condition. If the logical index is a different size than

the array, then elements from the array are matched against elements from the logical index based on their linear order, just as with linear indexing.

```
data = [ 1, 2, 3; 4, 5, 6 ];
idx = [ true, false, false, true ];
data(idx)
     ⇒ ans = [ 1 5 ]  # idx selected the 1st and 4th position elements
```

If the logical index is larger than then array an out of bounds error will occur if any true values attempt to select a linear position larger than the number of elements in the array.

```
idx = [ true, true, false; false, true, false; true; false; false ];
data(idx)
     ⇒ ans = [ 1; 2; 5; 3 ] # returns positions 1, 3, 4, 5 in a column

idx = [ true, true, false; false, true, false; true; false; true ];
data(idx)
     ⇒ error: a(9): out of bound 6 (dimensions are 2x3)
```

False elements of a logical index outside the array size are ignored, but when the ninth element of the logical index is true it triggers an error as it tries to select a nonexistent 9th element of the array.

### 8.1.1 Advanced Indexing

#### Chained indexing

Octave permits the use of repeated (chained) index expressions to extract subsets of an array in a single command without the need to use intermediate variables. This can make it easier to write code with either complicated indexing operations or using multiple indexing methods. The following example shows two equivalent index extraction operations:

```
A = reshape (1:16, 4, 4);
B = A(2:4, 2:3);
C = B(3:5);
D = C( [ true, false, true ] )
     ⇒ D = [ 8, 11 ]

D = A(2:4, 2:3)(3:5)([ true, false, true ])
     ⇒ D = [ 8, 11 ]
```

Chained indexing will necessarily be slower than a single index expression producing the same results, but is usually more computationally efficient than performing multiple discrete indexing operations with intermediate variable assignments.

Note that chained indexing is only compatible with right-hand expressions and can not be used on the left-hand side of assignment operations.

#### Component to linear index conversion

When it is necessary to extract subsets of entries out of an array whose indices cannot be written as a Cartesian product of components, linear indexing together with the function `sub2ind` can be used. For example:

```
A = reshape (1:8, 2, 2, 2)  # Create 3-D array
A =

ans(:,:,1) =

   1   3
   2   4

ans(:,:,2) =

   5   7
   6   8

A(sub2ind (size (A), [1, 2, 1], [1, 1, 2], [1, 2, 1]))
   ⇒ ans = [A(1, 1, 1), A(2, 1, 2), A(1, 2, 1)]
```

*ind* = sub2ind (*dims*, *i*, *j*)
*ind* = sub2ind (*dims*, *s1*, *s2*, ..., *sN*)

Convert subscripts to linear indices.

The input *dims* is a dimension vector where each element is the size of the array in the respective dimension (see [size], page 48). The remaining inputs are scalars or vectors of subscripts to be converted.

The output vector *ind* contains the converted linear indices.

Background: Array elements can be specified either by a linear index which starts at 1 and runs through the number of elements in the array, or they may be specified with subscripts for the row, column, page, etc. The functions `ind2sub` and `sub2ind` interconvert between the two forms.

The linear index traverses dimension 1 (rows), then dimension 2 (columns), then dimension 3 (pages), etc. until it has numbered all of the elements. Consider the following 3-by-3 matrices:

```
[(1,1), (1,2), (1,3)]     [1, 4, 7]
[(2,1), (2,2), (2,3)] ==> [2, 5, 8]
[(3,1), (3,2), (3,3)]     [3, 6, 9]
```

The left matrix contains the subscript tuples for each matrix element. The right matrix shows the linear indices for the same matrix.

The following example shows how to convert the two-dimensional indices `(2,1)` and `(2,3)` of a 3-by-3 matrix to linear indices with a single call to `sub2ind`.

```
s1 = [2, 2];
s2 = [1, 3];
ind = sub2ind ([3, 3], s1, s2)
    ⇒ ind =  2   8
```

**See also:** [ind2sub], page 163, [size], page 48.

[*s1*, *s2*, ..., *sN*] = ind2sub (*dims*, *ind*)

Convert linear indices to subscripts.

The input *dims* is a dimension vector where each element is the size of the array in the respective dimension (see [size], page 48). The second input *ind* contains linear indices to be converted.

The outputs *s1*, ..., *sN* contain the converted subscripts.

Background: Array elements can be specified either by a linear index which starts at 1 and runs through the number of elements in the array, or they may be specified with subscripts for the row, column, page, etc. The functions `ind2sub` and `sub2ind` interconvert between the two forms.

The linear index traverses dimension 1 (rows), then dimension 2 (columns), then dimension 3 (pages), etc. until it has numbered all of the elements. Consider the following 3-by-3 matrices:

```
[1, 4, 7]     [(1,1), (1,2), (1,3)]
[2, 5, 8] ==> [(2,1), (2,2), (2,3)]
[3, 6, 9]     [(3,1), (3,2), (3,3)]
```

The left matrix contains the linear indices for each matrix element. The right matrix shows the subscript tuples for the same matrix.

The following example shows how to convert the linear indices `2` and `8` to appropriate subscripts of a 3-by-3 matrix.

```
ind = [2, 8];
[r, c] = ind2sub ([3, 3], ind)
    ⇒ r =  2   2
    ⇒ c =  1   3
```

If the number of output subscripts exceeds the number of dimensions, the exceeded dimensions are set to `1`. On the other hand, if fewer subscripts than dimensions are provided, the exceeding dimensions are merged into the final requested dimension. For clarity, consider the following examples:

```
ind  = [2, 8];
dims = [3, 3];
## same as dims = [3, 3, 1]
[r, c, s] = ind2sub (dims, ind)
    ⇒ r =  2   2
    ⇒ c =  1   3
    ⇒ s =  1   1
## same as dims = [9]
r = ind2sub (dims, ind)
    ⇒ r =  2   8
```

**See also:** [sub2ind], page 163, [size], page 48.

*tf* = `isindex` (*ind*)
*tf* = `isindex` (*ind*, *n*)

Return true if *ind* is a valid index.

Valid indices are either positive integers (although possibly of real data type), or logical arrays.

If present, *n* specifies the maximum extent of the dimension to be indexed. When possible the internal result is cached so that subsequent indexing using *ind* will not perform the check again.

Implementation Note: Strings are first converted to double values before the checks for valid indices are made. Unless a string contains the NULL character "\0", it will always be a valid index.

## Component count not equal to dimensionality

An array with 'nd' dimensions can be indexed by an index expression which has from 1 to 'nd' components. For the ordinary and most common case, the number of components 'M' matches the number of dimensions 'nd'. In this case the ordinary indexing rules apply and each component corresponds to the respective dimension of the array.

However, if the number of indexing components exceeds the number of dimensions (M > nd) then the excess components must all be singletons (1). Moreover, if M < nd, the behavior is equivalent to reshaping the input object so as to merge the trailing nd - M dimensions into the last index dimension M. Thus, the result will have the dimensionality of the index expression, and not the original object. This is the case whenever dimensionality of the index is greater than one (M > 1), so that the special rules for linear indexing are not applied. This is easiest to understand with an example:

```
A = reshape (1:8, 2, 2, 2)  # Create 3-D array
A =

ans(:,:,1) =

   1   3
   2   4

ans(:,:,2) =

   5   7
   6   8

## 2-D indexing causes third dimension to be merged into second dimension.
## Equivalent array for indexing, Atmp, is now 2x4.
Atmp = reshape (A, 2, 4)
Atmp =

   1   3   5   7
   2   4   6   8


A(2,1)   # Reshape to 2x4 matrix, second entry of first column: ans = 2
A(2,4)   # Reshape to 2x4 matrix, second entry of fourth column: ans = 8
A(:,:)   # Reshape to 2x4 matrix, select all rows & columns, ans = Atmp
```

Note here the elegant use of the double colon to replace the call to the reshape function.

## Array replication

Another advanced use of linear indexing is to create arrays filled with a single value. This can be done by using an index of ones on a scalar value. The result is an object with the dimensions of the index expression and every element equal to the original scalar. For example, the following statements

```
a = 13;
a(ones (1, 4))
```

produce a row vector whose four elements are all equal to 13.

Similarly, by indexing a scalar with two vectors of ones it is possible to create a matrix. The following statements

```
a = 13;
a(ones (1, 2), ones (1, 3))
```

create a 2x3 matrix with all elements equal to 13. This could also have been written as

```
13(ones (2, 3))
```

It is more efficient to use indexing rather than the code construction `scalar * ones (M, N, ...)` because it avoids the unnecessary multiplication operation. Moreover, multiplication may not be defined for the object to be replicated whereas indexing an array is always defined. The following code shows how to create a 2x3 cell array from a base unit which is not itself a scalar.

```
{"Hello"}(ones (2, 3))
```

It should be noted that `ones (1, n)` (a row vector of ones) results in a range object (with zero increment). A range is stored internally as a starting value, increment, end value, and total number of values; hence, it is more efficient for storage than a vector or matrix of ones whenever the number of elements is greater than 4. In particular, when '`r`' is a row vector, the expressions

```
r(ones (1, n), :)
r(ones (n, 1), :)
```

will produce identical results, but the first one will be significantly faster, at least for '`r`' and '`n`' large enough. In the first case the index is held in compressed form as a range which allows Octave to choose a more efficient algorithm to handle the expression.

A general recommendation for users unfamiliar with these techniques is to use the function `repmat` for replicating smaller arrays into bigger ones, which uses such tricks.

## Indexing for performance enhancement

A second use of indexing is to speed up code. Indexing is a fast operation and judicious use of it can reduce the requirement for looping over individual array elements, which is a slow operation.

Consider the following example which creates a 10-element row vector $a$ containing the values $a_i = \sqrt{i}$.

```
for i = 1:10
  a(i) = sqrt (i);
endfor
```

It is quite inefficient to create a vector using a loop like this. In this case, it would have been much more efficient to use the expression

```
a = sqrt (1:10);
```

which avoids the loop entirely.

In cases where a loop cannot be avoided, or a number of values must be combined to form a larger matrix, it is generally faster to set the size of the matrix first (pre-allocate

storage), and then insert elements using indexing commands. For example, given a matrix `a`,

```
[nr, nc] = size (a);
x = zeros (nr, n * nc);
for i = 1:n
  x(:,(i-1)*nc+1:i*nc) = a;
endfor
```

is considerably faster than

```
x = a;
for i = 1:n-1
  x = [x, a];
endfor
```

because Octave does not have to repeatedly resize the intermediate result.

For more performance improvement suggestions see Chapter 19 [Vectorization and Faster Code Execution], page 689.

## 8.2 Calling Functions

A *function* is a name for a particular calculation. Because it has a name, you can ask for it by name at any point in the program. For example, the function `sqrt` computes the square root of a number.

A fixed set of functions are *built-in*, which means they are available in every Octave program. The `sqrt` function is one of these. In addition, you can define your own functions. See Chapter 11 [Functions and Scripts], page 205, for information about how to do this.

The way to use a function is with a *function call* expression, which consists of the function name followed by a list of *arguments* in parentheses. The arguments are expressions which give the raw materials for the calculation that the function will do. When there is more than one argument, they are separated by commas. If there are no arguments, you can omit the parentheses, but it is a good idea to include them anyway, to clearly indicate that a function call was intended. Here are some examples:

```
sqrt (x^2 + y^2)      # One argument
ones (n, m)           # Two arguments
rand ()               # No arguments
```

Each function expects a particular number of arguments. For example, the `sqrt` function must be called with a single argument, the number to take the square root of:

`sqrt (`*`argument`*`)`

Some of the built-in functions take a variable number of arguments, depending on the particular usage, and their behavior is different depending on the number of arguments supplied.

Like every other expression, the function call has a value, which is computed by the function based on the arguments you give it. In this example, the value of `sqrt (`*`argument`*`)` is the square root of the argument. A function can also have side effects, such as assigning the values of certain variables or doing input or output operations.

Unlike most languages, functions in Octave may return multiple values. For example, the following statement

```
[u, s, v] = svd (a)
```

computes the singular value decomposition of the matrix `a` and assigns the three result matrices to `u`, `s`, and `v`.

The left side of a multiple assignment expression is itself a list of expressions, that is, a list of variable names potentially qualified by index expressions. See also Section 8.1 [Index Expressions], page 159, and Section 8.6 [Assignment Expressions], page 179.

### 8.2.1 Call by Value

In Octave, unlike Fortran, function arguments are passed by value, which means that each argument in a function call is evaluated and assigned to a temporary location in memory before being passed to the function. There is currently no way to specify that a function parameter should be passed by reference instead of by value. This means that it is impossible to directly alter the value of a function parameter in the calling function. It can only change the local copy within the function body. For example, the function

```
function f (x, n)
  while (n-- > 0)
    disp (x);
  endwhile
endfunction
```

displays the value of the first argument *n* times. In this function, the variable *n* is used as a temporary variable without having to worry that its value might also change in the calling function. Call by value is also useful because it is always possible to pass constants for any function parameter without first having to determine that the function will not attempt to modify the parameter.

The caller may use a variable as the expression for the argument, but the called function does not know this: it only knows what value the argument had. For example, given a function called as

```
foo = "bar";
fcn (foo)
```

you should not think of the argument as being "the variable `foo`." Instead, think of the argument as the string value, `"bar"`.

Even though Octave uses pass-by-value semantics for function arguments, values are not copied unnecessarily. For example,

```
x = rand (1000);
f (x);
```

does not actually force two 1000 by 1000 element matrices to exist *unless* the function `f` modifies the value of its argument. Then Octave must create a copy to avoid changing the value outside the scope of the function `f`, or attempting (and probably failing!) to modify the value of a constant or the value of a temporary result.

### 8.2.2 Recursion

With some restrictions[1], recursive function calls are allowed. A *recursive function* is one which calls itself, either directly or indirectly. For example, here is an inefficient[2] way to compute the factorial of a given integer:

```
function retval = fact (n)
  if (n > 0)
    retval = n * fact (n-1);
  else
    retval = 1;
  endif
endfunction
```

This function is recursive because it calls itself directly. It eventually terminates because each time it calls itself, it uses an argument that is one less than was used for the previous call. Once the argument is no longer greater than zero, it does not call itself, and the recursion ends.

The function `max_recursion_depth` may be used to specify a limit to the recursion depth and prevents Octave from recursing infinitely. Similarly, the function `max_stack_depth` may be used to specify limit to the depth of function calls, whether recursive or not. These limits help prevent stack overflow on the computer Octave is running on, so that instead of exiting with a signal, the interpreter will throw an error and return to the command prompt.

*val* = max_recursion_depth ()
*old_val* = max_recursion_depth (*new_val*)
*old_val* = max_recursion_depth (*new_val*, "*local*")

Query or set the internal limit on the number of times a function may be called recursively.

If the limit is exceeded, an error message is printed and control returns to the top level.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [max_stack_depth], page 169.

*val* = max_stack_depth ()
*old_val* = max_stack_depth (*new_val*)
*old_val* = max_stack_depth (*new_val*, "*local*")

Query or set the internal limit on the number of times a function may be called recursively.

---

[1] Some of Octave's functions are implemented in terms of functions that cannot be called recursively. For example, the ODE solver `lsode` is ultimately implemented in a Fortran subroutine that cannot be called recursively, so `lsode` should not be called either directly or indirectly from within the user-supplied function that `lsode` requires. Doing so will result in an error.

[2] It would be much better to use `prod (1:n)`, or `gamma (n+1)` instead, after first checking to ensure that the value `n` is actually a positive integer.

If the limit is exceeded, an error message is printed and control returns to the top level.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [max_recursion_depth], page 169.

### 8.2.3 Access via Handle

A function may be abstracted and referenced via a function handle acquired using the special operator '`@`'. For example,

```
f = @plus;
f (2, 2)
⇒  4
```

is equivalent to calling `plus (2, 2)` directly. Beyond abstraction for general programming, function handles find use in callback methods for figures and graphics by adding listeners to properties or assigning pre-existing actions, such as in the following example:

```
function mydeletefcn (h, ~, msg)
  printf (msg);
endfunction
sombrero;
set (gcf, "deletefcn", {@mydeletefcn, "Bye!\n"});
close;
```

The above will print `"Bye!"` to the terminal upon the closing (deleting) of the figure. There are many graphics property actions for which a callback function may be assigned, including, `buttondownfcn`, `windowscrollwheelfcn`, `createfcn`, `deletefcn`, `keypressfcn`, etc.

Note that the '`@`' character also plays a role in defining class functions, i.e., methods, but not as a syntactical element. Rather it begins a directory name containing methods for a class that shares the directory name sans the '`@`' character. See Chapter 34 [Object Oriented Programming], page 973.

## 8.3 Arithmetic Operators

The following arithmetic operators are available, and work on scalars and matrices. The element-by-element operators and functions broadcast (see Section 19.2 [Broadcasting], page 691).

`x + y` Addition (always works element by element). If both operands are matrices, the number of rows and columns must both agree, or they must be broadcastable to the same shape.

`x - y` Subtraction (always works element by element). If both operands are matrices, the number of rows and columns of both must agree, or they must be broadcastable to the same shape.

`x * y` Matrix multiplication. The number of columns of *x* must agree with the number of rows of *y*.

*x* .* *y* Element-by-element multiplication. If both operands are matrices, the number of rows and columns must both agree, or they must be broadcastable to the same shape.

*x* / *y* Right division. This is conceptually equivalent to the expression

```
(inv (y') * x')'
```

but it is computed without forming the inverse of *y'*.

If the system is not square, or if the coefficient matrix is singular, a minimum norm solution is computed.

*x* ./ *y* Element-by-element right division.

*x* \ *y* Left division. This is conceptually equivalent to the expression

```
inv (x) * y
```

but it is computed without forming the inverse of *x*.

If the system is not square, or if the coefficient matrix is singular, a minimum norm solution is computed.

*x* .\ *y* Element-by-element left division. Each element of *y* is divided by each corresponding element of *x*.

*x* ^ *y* Power operator. If *x* and *y* are both scalars, this operator returns *x* raised to the power *y*. If *x* is a scalar and *y* is a square matrix, the result is computed using an eigenvalue expansion. If *x* is a square matrix, the result is computed by repeated multiplication if *y* is an integer, and by an eigenvalue expansion if *y* is not an integer. An error results if both *x* and *y* are matrices.

The implementation of this operator needs to be improved.

*x* .^ *y* Element-by-element power operator. If both operands are matrices, the number of rows and columns must both agree, or they must be broadcastable to the same shape. If several complex results are possible, the one with smallest non-negative argument (angle) is taken. This rule may return a complex root even when a real root is also possible. Use `realpow`, `realsqrt`, `cbrt`, or `nthroot` if a real result is preferred.

-*x* Negation.

+*x* Unary plus. This operator has no effect on the operand.

*x*' Complex conjugate transpose. For real arguments, this operator is the same as the transpose operator. For complex arguments, this operator is equivalent to the expression

```
conj (x.')
```

*x*.' Transpose.

Note that because Octave's element-by-element operators begin with a '`.`', there is a possible ambiguity for statements like

```
1./m
```

because the period could be interpreted either as part of the constant or as part of the operator. To resolve this conflict, Octave treats the expression as if you had typed

```
(1) ./ m
```

and not

```
(1.) / m
```

Although this is inconsistent with the normal behavior of Octave's lexer, which usually prefers to break the input into tokens by preferring the longest possible match at any given point, it is more useful in this case.

Note also that some combinations of binary operators and whitespace can create apparent ambiguities with the Command Syntax form of calling functions. See Section 11.13 [Command Syntax and Function Syntax], page 250, for a description of how Octave treats that syntax.

`B = ctranspose (A)`

Return the complex conjugate transpose of $A$.

This function and `A'` are equivalent.

**See also:** [transpose], page 174.

`y = pagectranspose (A)`

Return the page-wise complex conjugate transpose of the N-dimensional array $A$.

This is equivalent to `A(:,:, k)'` for each page $k$.

**See also:** [pagetranspose], page 174, [ctranspose], page 172, [permute], page 572.

`C = ldivide (A, B)`

Return the element-by-element left division of $A$ and $B$.

This function and `A .\ B` are equivalent.

**See also:** [rdivide], page 173, [mldivide], page 172, [times], page 173, [plus], page 173.

`C = minus (A, B)`

This function and `A - B` are equivalent.

**See also:** [plus], page 173, [uminus], page 174.

`C = mldivide (A, B)`

Return the matrix left division of $A$ and $B$.

This function and `A \ B` are equivalent.

If the system is not square, or if the coefficient matrix is singular, a minimum norm solution is computed.

**See also:** [mrdivide], page 173, [ldivide], page 172, [rdivide], page 173, [linsolve], page 652.

`C = mpower (A, B)`

Return the matrix power operation of $A$ raised to the $B$ power.

This function and `A ^ B` are equivalent.

**See also:** [power], page 173, [mtimes], page 173, [plus], page 173, [minus], page 172.

`C = mrdivide (A, B)`

Return the matrix right division of $A$ and $B$.

This function and `A / B` are equivalent.

If the system is not square, or if the coefficient matrix is singular, a minimum norm solution is computed.

**See also:** [mldivide], page 172, [rdivide], page 173, [plus], page 173, [minus], page 172.

`C = mtimes (A, B)`
`C = mtimes (A1, A2, ...)`

Return the matrix multiplication product of inputs.

This function and `A * B` are equivalent. If more arguments are given, the multiplication is applied cumulatively from left to right:

```
(...((A1 * A2) * A3) * ...)
```

**See also:** [times], page 173, [plus], page 173, [minus], page 172, [rdivide], page 173, [mrdivide], page 173, [mldivide], page 172, [mpower], page 172, [tensorprod], page 675.

`C = plus (A, B)`
`C = plus (A1, A2, ...)`

This function and `A + B` are equivalent.

If more arguments are given, the summation is applied cumulatively from left to right:

```
(...((A1 + A2) + A3) + ...)
```

**See also:** [minus], page 172, [uplus], page 174.

`C = power (A, B)`

Return the element-by-element operation of $A$ raised to the $B$ power.

This function and `A .^ B` are equivalent.

If several complex results are possible, returns the one with smallest non-negative argument (angle). Use `realpow`, `realsqrt`, `cbrt`, or `nthroot` if a real result is preferred.

**See also:** [mpower], page 172, [realpow], page 604, [realsqrt], page 604, [cbrt], page 604, [nthroot], page 604.

`C = rdivide (A, B)`

Return the element-by-element right division of $A$ and $B$.

This function and `A ./ B` are equivalent.

**See also:** [ldivide], page 172, [mrdivide], page 173, [times], page 173, [plus], page 173.

`C = times (A, B)`
`C = times (A1, A2, ...)`

Return the element-by-element multiplication product of inputs.

This function and `A .* B` are equivalent. If more arguments are given, the multiplication is applied cumulatively from left to right:

```
(...((A1 .* A2) .* A3) .* ...)
```

**See also:** [mtimes], page 173, [rdivide], page 173.

`B = transpose (A)`
: Return the transpose of $A$.

  This function and `A.'` are equivalent.

  **See also:** [ctranspose], page 172.

`B = pagetranspose (A)`
: Return the page-wise transpose of the N-dimensional array $A$.

  This is equivalent to `A(:,:, k).'` for each page $k$.

  **See also:** [pagectranspose], page 172, [transpose], page 174, [permute], page 572.

`B = uminus (A)`
: This function and `- A` are equivalent.

  **See also:** [uplus], page 174, [minus], page 172.

`B = uplus (A)`
: This function and `+ A` are equivalent.

  **See also:** [uminus], page 174, [plus], page 173.

## 8.4 Comparison Operators

*Comparison operators* compare numeric values for relationships such as equality. They are written using *relational operators*.

All of Octave's comparison operators return a value of 1 if the comparison is true, or 0 if it is false. For matrix values, they all work on an element-by-element basis. Broadcasting rules apply. See Section 19.2 [Broadcasting], page 691. For example:

```
[1, 2; 3, 4] == [1, 3; 2, 4]
     ⇒  1  0
        0  1
```

According to broadcasting rules, if one operand is a scalar and the other is a matrix, the scalar is compared to each element of the matrix in turn, and the result is the same size as the matrix.

`x < y`
: True if $x$ is less than $y$.

`x <= y`
: True if $x$ is less than or equal to $y$.

`x == y`
: True if $x$ is equal to $y$.

`x >= y`
: True if $x$ is greater than or equal to $y$.

`x > y`
: True if $x$ is greater than $y$.

`x != y`
`x ~= y`
: True if $x$ is not equal to $y$.

For complex numbers, the following ordering is defined: *z1* `<` *z2* if and only if

```
abs (z1) < abs (z2)
|| (abs (z1) == abs (z2) && arg (z1) < arg (z2))
```

This is consistent with the ordering used by *max*, *min* and *sort*, but is not consistent with MATLAB, which only compares the real parts.

String comparisons may also be performed with the `strcmp` function, not with the comparison operators listed above. See Chapter 5 [Strings], page 75.

*TF* = `eq` (*A*, *B*)
Return true if the two inputs are equal.

This function is equivalent to `A == B`.

**See also:** [ne], page 175, [isequal], page 175, [le], page 175, [ge], page 175, [gt], page 175, [ne], page 175, [lt], page 175.

*TF* = `ge` (*A*, *B*)
This function is equivalent to `A >= B`.

**See also:** [le], page 175, [eq], page 175, [gt], page 175, [ne], page 175, [lt], page 175.

*TF* = `gt` (*A*, *B*)
This function is equivalent to `A > B`.

**See also:** [le], page 175, [eq], page 175, [ge], page 175, [ne], page 175, [lt], page 175.

*tf* = `isequal` (*x1*, *x2*, ...)
Return true if all of *x1*, *x2*, ... are equal.

**See also:** [isequaln], page 175.

*tf* = `isequaln` (*x1*, *x2*, ...)
Return true if all of *x1*, *x2*, ... are equal under the additional assumption that NaN == NaN (no comparison of NaN placeholders in dataset).

**See also:** [isequal], page 175.

*TF* = `le` (*A*, *B*)
This function is equivalent to `A <= B`.

**See also:** [eq], page 175, [ge], page 175, [gt], page 175, [ne], page 175, [lt], page 175.

*TF* = `lt` (*A*, *B*)
This function is equivalent to `A < B`.

**See also:** [le], page 175, [eq], page 175, [ge], page 175, [gt], page 175, [ne], page 175.

*TF* = `ne` (*A*, *B*)
Return true if the two inputs are not equal.

This function is equivalent to `A != B`.

**See also:** [eq], page 175, [isequal], page 175, [le], page 175, [ge], page 175, [lt], page 175.

## 8.5 Boolean Expressions

### 8.5.1 Element-by-element Boolean Operators

An *element-by-element boolean expression* is a combination of comparison expressions using the boolean operators "or" ('`|`'), "and" ('`&`'), and "not" ('`!`'), along with parentheses to control nesting. The truth of the boolean expression is computed by combining the truth values of the corresponding elements of the component expressions. A value is considered to be false if it is zero, and true otherwise.

Element-by-element boolean expressions can be used wherever comparison expressions can be used. They can be used in `if` and `while` statements. However, a matrix value used as the condition in an `if` or `while` statement is only true if *all* of its elements are nonzero.

Like comparison operations, each element of an element-by-element boolean expression also has a numeric value (1 if true, 0 if false) that comes into play if the result of the boolean expression is stored in a variable, or used in arithmetic.

Here are descriptions of the three element-by-element boolean operators.

*boolean1* `&` *boolean2*
: Elements of the result are true if both corresponding elements of *boolean1* and *boolean2* are true.

*boolean1* `|` *boolean2*
: Elements of the result are true if either of the corresponding elements of *boolean1* or *boolean2* is true.

`!` *boolean*
`~` *boolean*
: Each element of the result is true if the corresponding element of *boolean* is false.

These operators work on an element-by-element basis. For example, the expression

```
[1, 0; 0, 1] & [1, 0; 2, 3]
```

returns a two by two identity matrix.

For the binary operators, broadcasting rules apply. See Section 19.2 [Broadcasting], page 691. In particular, if one of the operands is a scalar and the other a matrix, the operator is applied to the scalar and each element of the matrix.

For the binary element-by-element boolean operators, both subexpressions *boolean1* and *boolean2* are evaluated before computing the result. This can make a difference when the expressions have side effects. For example, in the expression

```
a & b++
```

the value of the variable *b* is incremented even if the variable *a* is zero.

This behavior is necessary for the boolean operators to work as described for matrix-valued operands.

*TF* = `and` (*x*, *y*)
*TF* = `and` (*x1*, *x2*, ...)
: Return the logical AND of *x* and *y*.

This function is equivalent to the operator syntax `x & y`. If more than two arguments are given, the logical AND is applied cumulatively from left to right:

```
(...((x1 & x2) & x3) & ...)
```

**See also:** [or], page 177, [not], page 177, [xor], page 566.

`z = not (x)`

Return the logical NOT of *x*.

This function is equivalent to the operator syntax `! x`.

**See also:** [and], page 176, [or], page 177, [xor], page 566.

`TF = or (x, y)`
`TF = or (x1, x2, ...)`

Return the logical OR of *x* and *y*.

This function is equivalent to the operator syntax `x | y`. If more than two arguments are given, the logical OR is applied cumulatively from left to right:

```
(...((x1 | x2) | x3) | ...)
```

**See also:** [and], page 176, [not], page 177, [xor], page 566.

### 8.5.2 Short-circuit Boolean Operators

Combined with the implicit conversion to scalar values in `if` and `while` conditions, Octave's element-by-element boolean operators are often sufficient for performing most logical operations. However, it is sometimes desirable to stop evaluating a boolean expression as soon as the overall truth value can be determined. Octave's *short-circuit* boolean operators work this way.

*boolean1* `&&` *boolean2*

The expression *boolean1* is evaluated and converted to a scalar using the equivalent of the operation `all (`*boolean1*`(:))`. If *boolean1* is not a logical value, it is considered true if its value is nonzero, and false if its value is zero. If *boolean1* is an array, it is considered true only if it is non-empty and all elements are nonzero. If *boolean1* evaluates to false, the result of the overall expression is false. If it is true, the expression *boolean2* is evaluated in the same way as *boolean1*. If it is true, the result of the overall expression is true. Otherwise the result of the overall expression is false.

**Warning:** the one exception to the equivalence with evaluating `all (`*boolean1*`(:))` is when `boolean1` an the empty array. For MATLAB compatibility, the truth value of an empty array is always `false` so `[] && true` evaluates to `false` even though `all ([])` is `true`.

*boolean1* `||` *boolean2*

The expression *boolean1* is evaluated and converted to a scalar using the equivalent of the operation `all (`*boolean1*`(:))`. If *boolean1* is not a logical value, it is considered true if its value is nonzero, and false if its value is zero. If *boolean1* is an array, it is considered true only if it is non-empty and all elements are nonzero. If *boolean1* evaluates to true, the result of the overall expression is true. If it is false, the expression *boolean2* is evaluated in the same way as

*boolean1*. If it is true, the result of the overall expression is true. Otherwise the result of the overall expression is false.

**Warning:** the truth value of an empty matrix is always `false`, see the previous list item for details.

The fact that both operands may not be evaluated before determining the overall truth value of the expression can be important. For example, in the expression

```
a && b++
```

the value of the variable *b* is only incremented if the variable *a* is nonzero.

This can be used to write somewhat more concise code. For example, it is possible write

```
function f (a, b, c)
  if (nargin > 2 && ischar (c))
    ...
```

instead of having to use two `if` statements to avoid attempting to evaluate an argument that doesn't exist. For example, without the short-circuit feature, it would be necessary to write

```
function f (a, b, c)
  if (nargin > 2)
    if (ischar (c))
      ...
```

Writing

```
function f (a, b, c)
  if (nargin > 2 & ischar (c))
    ...
```

would result in an error if `f` were called with one or two arguments because Octave would be forced to try to evaluate both of the operands for the operator '`&`'.

MATLAB has special behavior that allows the operators '`&`' and '`|`' to short-circuit when used in the truth expression for `if` and `while` statements. Octave behaves the same way for compatibility, however, the use of the '`&`' and '`|`' operators in this way is strongly discouraged and a warning will be issued. Instead, you should use the '`&&`' and '`||`' operators that always have short-circuit behavior.

Finally, the ternary operator (?:) is not supported in Octave. If short-circuiting is not important, it can be replaced by the `ifelse` function.

*M* = `merge` (*mask*, *tval*, *fval*)
*M* = `ifelse` (*mask*, *tval*, *fval*)

Merge elements of *true_val* and *false_val*, depending on the value of *mask*.

If *mask* is a logical scalar, the other two arguments can be arbitrary values. Otherwise, *mask* must be a logical array, and *tval*, *fval* should be arrays of matching class, or cell arrays. In the scalar mask case, *tval* is returned if *mask* is true, otherwise *fval* is returned.

In the array mask case, both *tval* and *fval* must be either scalars or arrays with dimensions equal to *mask*. The result is constructed as follows:

```
result(mask) = tval(mask);
result(! mask) = fval(! mask);
```

*mask* can also be arbitrary numeric type, in which case it is first converted to logical.

Programming Note: `ifelse` is an alias for `merge` and can be used interchangeably.

**See also:** [logical], page 66, [diff], page 567.

## 8.6 Assignment Expressions

An *assignment* is an expression that stores a new value into a variable. For example, the following expression assigns the value 1 to the variable `z`:

```
z = 1
```

After this expression is executed, the variable `z` has the value 1. Whatever old value `z` had before the assignment is forgotten. The '`=`' sign is called an *assignment operator*.

Assignments can store string values also. For example, the following expression would store the value `"this food is good"` in the variable `message`:

```
thing = "food"
predicate = "good"
message = [ "this " , thing , " is " , predicate ]
⇒ "this food is good"
```

(This also illustrates concatenation of strings.)

Most operators (addition, concatenation, and so on) have no effect except to compute a value. If you ignore the value, you might as well not use the operator. An assignment operator is different. It does produce a value, but even if you ignore the value, the assignment still makes itself felt through the alteration of the variable. We call this a *side effect*.

The left-hand operand of an assignment need not be a variable (see Chapter 7 [Variables], page 145). It can also be an element of a matrix (see Section 8.1 [Index Expressions], page 159) or a list of return values (see Section 8.2 [Calling Functions], page 167). These are all called *lvalues*, which means they can appear on the left-hand side of an assignment operator. The right-hand operand may be any expression. It produces the new value which the assignment stores in the specified variable, matrix element, or list of return values.

It is important to note that variables do *not* have permanent types. The type of a variable is simply the type of whatever value it happens to hold at the moment. In the following program fragment, the variable `foo` has a numeric value at first, and a string value later on:

```
>> foo = 1
foo = 1
>> foo = "bar"
foo = bar
```

When the second assignment gives `foo` a string value, the fact that it previously had a numeric value is forgotten.

Assignment of a scalar to an indexed matrix sets all of the elements that are referenced by the indices to the scalar value. For example, if `a` is a matrix with at least two columns,

```
a(:, 2) = 5
```

sets all the elements in the second column of `a` to 5.

When an assignment sets the value of a vector, matrix, or array element at a position or dimension outside of that variable's current size, the array size will be increased to accommodate the new values:

```
>> a = [1, 2, 3]
a = 1 2 3
>> a(4) = 4
a = 1 2 3 4
>> a(2, :) = [5, 6, 7, 8]
a =
   1   2   3   4
   5   6   7   8
```

Attempting to increase the size of an array such that the desired output size is ambiguous will result in an error:

```
>> a(9) = 10
⊣ error: Invalid resizing operation or ambiguous assignment to an
out-of-bounds array element
```

This is because adding the 9th element creates an ambiguity in the desired array position for the value 10, each possibility requiring a different array size expansion to accommodate the assignment.

Assignments may be made with fewer specified elements than would be required to fill the newly expanded array as long as the assignment is unambiguous. In these cases the array will be automatically padded with *null* values:

```
>> a = [1, 2]
a =   1   2
>> a(4) = 5
a =   1   2   0   5
>> a(3, :) = [6, 7, 8, 9]
a =
   1   2   0   5
   0   0   0   0
   6   7   8   9
>> a(4, 5) = 10
a =
    1    2    0    5    0
    0    0    0    0    0
    6    7    8    9    0
    0    0    0    0   10
```

For all built-in types, the *null* value will be appropriate to that object type.

Numeric arrays:

```
>> a = int32 ([1, 2])
a = 1, 2
>> a(4) = 5
a = 1 2 0 5
```

Logical arrays:

```
>> a = [true, false, true]
a = 1 0 1
>> d(5) = true
d = 1 0 1 0 1
```

Character arrays:

```
>> a = "abc"
a = abc
>> a(5) = "d"
a = abcd
>> double (a)
ans = 97 98 99 0 100
```

Cell arrays:

```
>> e = {1, "foo", [3, 4]};
>> e(5) = "bar"
e =
{
  [1,1] = 1
  [1,2] = foo
  [1,3] =

     3   4

  [1,4] = [](0x0)
  [1,5] = bar
}
```

Struct arrays:

```
>> a = struct("foo",1,"bar",2);
>> a(3) = struct("foo",3,"bar",9)
a =

  1x3 struct array containing the fields:

    foo
    bar

>> a.foo
ans = 1
ans = [](0x0)
ans = 3
>> a.bar
ans = 2
ans = [](0x0)
ans = 9
```

Note that Octave currently is unable to concatenate arbitrary object types into arrays. Such behavior must be explicitly defined within the object class or attempts at concatenation will result in an error. See Chapter 34 [Object Oriented Programming], page 973.

Assigning an empty matrix '`[]`' works in most cases to allow you to delete rows or columns of matrices and vectors. See Section 4.1.1 [Empty Matrices], page 56. For example, given a 4 by 5 matrix $A$, the assignment

```
A (3, :) = []
```

deletes the third row of $A$, and the assignment

```
A (:, 1:2:5) = []
```

deletes the first, third, and fifth columns.

Deleting part of an array object will necessarily resize the object. When the deletion allows for consistent size reduction across a dimension, e.g., one element of a vector, or one row or column of a matrix, the size along that dimension will be reduced while preserving dimensionality. If, however, dimensionality cannot be maintained, the object will be reshaped into a vector following column-wise element ordering:

```
>> a = [1, 2, 3, 4; 5, 6, 7, 8]
a =
   1   2   3   4
   5   6   7   8
>> a(:, 3) = []
a =
   1   2   4
   5   6   8
>> a(4) = []
a = 1 5 2 4 8
```

An assignment is an expression, so it has a value. Thus, `z = 1` as an expression has the value 1. One consequence of this is that you can write multiple assignments together:

```
x = y = z = 0
```

stores the value 0 in all three variables. It does this because the value of `z = 0`, which is 0, is stored into `y`, and then the value of `y = z = 0`, which is 0, is stored into `x`.

This is also true of assignments to lists of values, so the following is a valid expression

```
[a, b, c] = [u, s, v] = svd (a)
```

that is exactly equivalent to

```
[u, s, v] = svd (a)
a = u
b = s
c = v
```

In expressions like this, the number of values in each part of the expression need not match. For example, the expression

```
[a, b] = [u, s, v] = svd (a)
```

is equivalent to

```
[u, s, v] = svd (a)
a = u
b = s
```

The number of values on the left side of the expression can, however, not exceed the number of values on the right side. For example, the following will produce an error.

```
[a, b, c, d] = [u, s, v] = svd (a);
⊣ error: element number 4 undefined in return list
```

The symbol ~ may be used as a placeholder in the list of lvalues, indicating that the corresponding return value should be ignored and not stored anywhere:

```
[~, s, v] = svd (a);
```

This is cleaner and more memory efficient than using a dummy variable. The `nargout` value for the right-hand side expression is not affected. If the assignment is used as an expression, the return value is a comma-separated list with the ignored values dropped.

A very common programming pattern is to increment an existing variable with a given value, like this

```
a = a + 2;
```

This can be written in a clearer and more condensed form using the `+=` operator

```
a += 2;
```

Similar operators also exist for subtraction (`-=`), multiplication (`*=`), and division (`/=`). An expression of the form

*expr1 op= expr2*

is evaluated as

*expr1* = (*expr1*) *op* (*expr2*)

where *op* can be either `+`, `-`, `*`, or `/`, as long as *expr2* is a simple expression with no side effects. If *expr2* also contains an assignment operator, then this expression is evaluated as

*temp* = *expr2*
*expr1* = (*expr1*) *op temp*

where *temp* is a placeholder temporary value storing the computed result of evaluating *expr2*. So, the expression

```
a *= b+1
```

is evaluated as

```
a = a * (b+1)
```

and *not*

```
a = a * b + 1
```

You can use an assignment anywhere an expression is called for. For example, it is valid to write `x != (y = 1)` to set `y` to 1 and then test whether `x` equals 1. But this style tends to make programs hard to read. Except in a one-shot program, you should rewrite it to get rid of such nesting of assignments. This is never very hard.

## 8.7 Increment Operators

*Increment operators* increase or decrease the value of a variable by 1. The operator to increment a variable is written as '`++`'. It may be used to increment a variable either before or after taking its value.

For example, to pre-increment the variable *x*, you would write `++x`. This would add one to *x* and then return the new value of *x* as the result of the expression. It is exactly the same as the expression `x = x + 1`.

To post-increment a variable *x*, you would write `x++`. This adds one to the variable *x*, but returns the value that *x* had prior to incrementing it. For example, if *x* is equal to 2, the result of the expression `x++` is 2, and the new value of *x* is 3.

For matrix and vector arguments, the increment and decrement operators work on each element of the operand.

The increment and decrement operators must "hug" their corresponding variable. That means, no white spaces are allowed between these operators and the variable they affect.

Here is a list of all the increment and decrement expressions.

`++x` This expression increments the variable *x*. The value of the expression is the *new* value of *x*. It is equivalent to the expression `x = x + 1`.

`--x` This expression decrements the variable *x*. The value of the expression is the *new* value of *x*. It is equivalent to the expression `x = x - 1`.

`x++` This expression causes the variable *x* to be incremented. The value of the expression is the *old* value of *x*.

`x--` This expression causes the variable *x* to be decremented. The value of the expression is the *old* value of *x*.

## 8.8 Operator Precedence

*Operator precedence* determines how operators are grouped, when different operators appear close by in one expression. For example, '`*`' has higher precedence than '`+`'. Thus, the expression `a + b * c` means to multiply `b` and `c`, and then add `a` to the product (i.e., `a + (b * c)`).

You can overrule the precedence of the operators by using parentheses. You can think of the precedence rules as saying where the parentheses are assumed if you do not write parentheses yourself. In fact, it is wise to use parentheses whenever you have an unusual combination of operators, because other people who read the program may not remember what the precedence is in this case. You might forget as well, and then you too could make a mistake. Explicit parentheses will help prevent any such mistake.

When operators of equal precedence are used together, the leftmost operator groups first, except for the assignment operators, which group in the opposite order. Thus, the expression `a - b + c` groups as `(a - b) + c`, but the expression `a = b = c` groups as `a = (b = c)`.

The precedence of prefix unary operators is important when another operator follows the operand. For example, `-x^2` means `-(x^2)`, because '`-`' has lower precedence than '`^`'.

Here is a table of the operators in Octave, in order of decreasing precedence. Unless noted, all operators group left to right.

`function call and array indexing, cell array indexing, and structure element indexing`
: '`()`' '`{}`' '`.`'

postfix increment, and postfix decrement
: ‘++’ ‘--’

  These operators group right to left.

transpose and exponentiation
: ‘'’ ‘.'’ ‘^’ ‘.^’

unary plus, unary minus, prefix increment, prefix decrement, and logical "not"
: ‘+’ ‘-’ ‘++’ ‘--’ ‘~’ ‘!’

multiply and divide
: ‘*’ ‘/’ ‘\’ ‘.\’ ‘.*’ ‘./’

add, subtract
: ‘+’ ‘-’

colon
: ‘:’

relational
: ‘<’ ‘<=’ ‘==’ ‘>=’ ‘>’ ‘!=’ ‘~=’

element-wise "and"
: ‘&’

element-wise "or"
: ‘|’

logical "and"
: ‘&&’

logical "or"
: ‘||’

assignment
: ‘=’ ‘+=’ ‘-=’ ‘*=’ ‘/=’ ‘\=’ ‘^=’ ‘.*=’ ‘./=’ ‘.\=’ ‘.^=’ ‘|=’ ‘&=’

  These operators group right to left.

# 9 Evaluation

Normally, you evaluate expressions simply by typing them at the Octave prompt, or by asking Octave to interpret commands that you have saved in a file.

Sometimes, you may find it necessary to evaluate an expression that has been computed and stored in a string, which is exactly what the `eval` function lets you do.

`eval (`*try*`)`
`eval (`*try*`, `*catch*`)`
`[`*var1*`, ...] = eval (...)`

Parse the string *try* and evaluate it as if it were an Octave program.

If execution fails, evaluate the optional string *catch*.

The string *try* is evaluated in the current context, so any results remain available after `eval` returns.

The following example creates the variable *A* with the approximate value of pi (3.1416) in the current workspace.

```
eval ('A = acos (-1);');
```

If an error occurs during the evaluation of *try* then the *catch* string is evaluated, as the following example shows:

```
eval ('error ("This is a bad example");',
      'printf ("This error occurred:\n%s\n", lasterr ());');
     ⊣ This error occurred:
       This is a bad example
```

Rather than create variables as part of the code string *try*, it is clearer and slightly faster to assign the results of evaluation to an output variable(s). The first example can be re-written as

```
A = eval ('acos (-1);');
```

Programming Note: if you are only using `eval` as an error-capturing mechanism, rather than for the execution of arbitrary code strings, consider using `try/catch` blocks or `unwind_protect/unwind_protect_cleanup` blocks instead. These techniques have higher performance and don't introduce the security considerations that the evaluation of arbitrary code does.

**See also:** [evalin], page 191, [evalc], page 187, [assignin], page 191, [feval], page 189, [try], page 1197, [unwind_protect], page 1197.

The `evalc` function additionally captures any console output produced by the evaluated expression.

*s* `= evalc (`*try*`)`
*s* `= evalc (`*try*`, `*catch*`)`
`[~, `*var1*`, ...] = evalc (...)`

Parse and evaluate the string *try* as if it were an Octave program, while capturing the output into the return variable *s*.

If execution fails, evaluate the optional string *catch*.

This function behaves like `eval`, but any output or warning messages which would normally be written to the console are captured and returned in the string *s*.

If the first output *s* is ignored with `~` then the actual results of the code evaluation (not the string capture) will be assigned to output variables *var1*, *var2*, etc.

Example 1:

```
s = evalc ("t = 42"), t
  ⇒ s = t =  42

  ⇒ t =  42
```

Example 2:

```
[~, p] = evalc ("pi")
  ⇒ p = 3.1416
```

Programming Note: The `diary` is disabled during the execution of this function. When `system` is used, any output produced by external programs is *not* captured, unless their output is captured by the `system` function itself.

**See also:** [eval], page 187, [diary], page 35.

## 9.1 Calling a Function by its Name

The `feval` function allows you to call a function from a string containing its name. This is useful when writing a function that needs to call user-supplied functions. The `feval` function takes the name of the function to call as its first argument, and the remaining arguments are given to the function.

The following example is a simple-minded function using `feval` that finds the root of a user-supplied function of one variable using Newton's method.

```
function result = newtroot (fname, x)

# usage: newtroot (fname, x)
#
#   fname : a string naming a function f(x).
#   x     : initial guess

  delta = tol = sqrt (eps);
  maxit = 200;
  fx = feval (fname, x);
  for i = 1:maxit
    if (abs (fx) < tol)
      result = x;
      return;
    else
      fx_new = feval (fname, x + delta);
      deriv = (fx_new - fx) / delta;
      x = x - fx / deriv;
      fx = fx_new;
    endif
```

```
  endfor

  result = x;

endfunction
```

Note that this is only meant to be an example of calling user-supplied functions and should not be taken too seriously. In addition to using a more robust algorithm, any serious code would check the number and type of all the arguments, ensure that the supplied function really was a function, etc. See Section 4.8 [Predicates for Numeric Objects], page 68, for a list of predicates for numeric objects, and see Section 7.3 [Status of Variables], page 150, for a description of the `exist` function.

`[`*y1*`, `*y2*`, ...] = feval (`*'fcn'*`, `*x1*`, `*x2*`, ...)`
`[`*y1*`, `*y2*`, ...] = feval (@`*fcn*`, `*x1*`, `*x2*`, ...)`

Evaluate the function *fcn* with inputs *x1*, *x2*, ...

The function *fcn* may be specified by name in a string or given as a function handle. Any arguments after the first are passed as inputs to the named function. For example,

```
feval ("acos", -1)
     ⇒ 3.1416
```

calls the function `acos` with the argument '`-1`'.

The function `feval` can also be used with function handles of any sort (see Section 11.12.1 [Function Handles], page 248). Historically, `feval` was the only way to call user-supplied functions in strings, but function handles are now preferred due to the cleaner syntax they offer. For example,

```
f = @exp;
feval (f, 1)
    ⇒ 2.7183
f (1)
    ⇒ 2.7183
```

are equivalent ways to call the function referred to by *f*. If it cannot be predicted beforehand whether *f* is a function handle, function name in a string, or inline function then `feval` can be used instead.

**See also:** [builtin], page 235, [eval], page 187, [evalin], page 191.

A similar function `run` exists for calling user script files, that are not necessarily on the user path

`run `*script*
`run ("`*script*`")`

Run *script* in the current workspace.

Scripts which reside in directories specified in Octave's load path, and which end with the extension `.m`, can be run simply by typing their name. For scripts not located on the load path, use `run`.

The filename *script* can be a bare, fully qualified, or relative filename and with or without a file extension. If no extension is specified, Octave will first search for a script with the `.m` extension before falling back to the script name without an extension.

Implementation Note: If *script* includes a path component, then `run` first changes the working directory to the directory where *script* is found. Next, the script is executed. Finally, `run` returns to the original working directory *unless script* has specifically changed directories.

**See also:** [path], page 229, [addpath], page 227, [source], page 239.

## 9.2 Evaluation in a Different Context

Before you evaluate an expression you need to substitute the values of the variables used in the expression. These are stored in the symbol table. Whenever the interpreter starts a new function it saves the current symbol table and creates a new one, initializing it with the list of function parameters and a couple of predefined variables such as `nargin`. Expressions inside the function use the new symbol table.

Sometimes you want to write a function so that when you call it, it modifies variables in your own context. This allows you to use a pass-by-name style of function, which is similar to using a pointer in programming languages such as C.

Consider how you might write `save` and `load` as m-files. For example:

```
function create_data
  x = linspace (0, 10, 10);
  y = sin (x);
  save mydata x y
endfunction
```

With `evalin`, you could write `save` as follows:

```
function save (file, name1, name2)
  f = open_save_file (file);
  save_var (f, name1, evalin ("caller", name1));
  save_var (f, name2, evalin ("caller", name2));
endfunction
```

Here, '`caller`' is the `create_data` function and `name1` is the string `"x"`, which evaluates simply as the value of `x`.

You later want to load the values back from `mydata` in a different context:

```
function process_data
  load mydata
  ... do work ...
endfunction
```

With `assignin`, you could write `load` as follows:

```
function load (file)
  f = open_load_file (file);
  [name, val] = load_var (f);
  assignin ("caller", name, val);
  [name, val] = load_var (f);
  assignin ("caller", name, val);
endfunction
```

Here, '`caller`' is the `process_data` function.

You can set and use variables at the command prompt using the context 'base' rather than 'caller'.

These functions are rarely used in practice. One example is the `fail ('code', 'pattern')` function which evaluates 'code' in the caller's context and checks that the error message it produces matches the given pattern. Other examples such as `save` and `load` are written in C++ where all Octave variables are in the 'caller' context and `evalin` is not needed.

`evalin (`*context*`, `*try*`)`
`evalin (`*context*`, `*try*`, `*catch*`)`
`[`*var1*`, ...] = evalin (...)`

Like `eval`, except that the expressions are evaluated in the context *context*, which may be either `"caller"` or `"base"`.

**See also:** [eval], page 187, [assignin], page 191.

`assignin (`*context*`, `*varname*`, `*value*`)`

Assign *value* to *varname* in context *context*, which may be either `"base"` or `"caller"`.

**See also:** [evalin], page 191.

# 10 Statements

Statements may be a simple constant expression or a complicated list of nested loops and conditional statements.

*Control statements* such as `if`, `while`, and so on control the flow of execution in Octave programs. All the control statements start with special keywords such as `if` and `while`, to distinguish them from simple expressions. Many control statements contain other statements; for example, the `if` statement contains another statement which may or may not be executed.

Each control statement has a corresponding *end* statement that marks the end of the control statement. For example, the keyword `endif` marks the end of an `if` statement, and `endwhile` marks the end of a `while` statement. You can use the keyword `end` anywhere a more specific end keyword is expected, but using the more specific keywords is preferred because if you use them, Octave is able to provide better diagnostics for mismatched or missing end tokens.

The list of statements contained between keywords like `if` or `while` and the corresponding end statement is called the *body* of a control statement.

## 10.1 The if Statement

The `if` statement is Octave's decision-making statement. There are three basic forms of an `if` statement. In its simplest form, it looks like this:

```
if (condition)
  then-body
endif
```

*condition* is an expression that controls what the rest of the statement will do. The *then-body* is executed only if *condition* is true.

The condition in an `if` statement is considered true if its value is nonzero, and false if its value is zero. If the value of the conditional expression in an `if` statement is a vector or a matrix, it is considered true only if it is non-empty and *all* of the elements are nonzero. The conceptually equivalent code when *condition* is a matrix is shown below.

```
if (matrix) ≡ if (all (matrix(:)))
```

The second form of an if statement looks like this:

```
if (condition)
  then-body
else
  else-body
endif
```

If *condition* is true, *then-body* is executed; otherwise, *else-body* is executed.

Here is an example:

```
if (rem (x, 2) == 0)
  printf ("x is even\n");
else
  printf ("x is odd\n");
endif
```

In this example, if the expression `rem (x, 2) == 0` is true (that is, the value of `x` is divisible by 2), then the first `printf` statement is evaluated, otherwise the second `printf` statement is evaluated.

The third and most general form of the `if` statement allows multiple decisions to be combined in a single statement. It looks like this:

```
if (condition)
  then-body
elseif (condition)
  elseif-body
else
  else-body
endif
```

Any number of `elseif` clauses may appear. Each condition is tested in turn, and if one is found to be true, its corresponding *body* is executed. If none of the conditions are true and the `else` clause is present, its body is executed. Only one `else` clause may appear, and it must be the last part of the statement.

In the following example, if the first condition is true (that is, the value of `x` is divisible by 2), then the first `printf` statement is executed. If it is false, then the second condition is tested, and if it is true (that is, the value of `x` is divisible by 3), then the second `printf` statement is executed. Otherwise, the third `printf` statement is performed.

```
if (rem (x, 2) == 0)
  printf ("x is even\n");
elseif (rem (x, 3) == 0)
  printf ("x is odd and divisible by 3\n");
else
  printf ("x is odd\n");
endif
```

Note that the `elseif` keyword must not be spelled `else if`, as is allowed in Fortran. If it is, the space between the `else` and `if` will tell Octave to treat this as a new `if` statement within another `if` statement's `else` clause. For example, if you write

```
if (c1)
  body-1
else if (c2)
  body-2
endif
```

Octave will expect additional input to complete the first `if` statement. If you are using Octave interactively, it will continue to prompt you for additional input. If Octave is reading this input from a file, it may complain about missing or mismatched `end` statements, or, if you have not used the more specific `end` statements (`endif`, `endfor`, etc.), it may simply produce incorrect results, without producing any warning messages.

It is much easier to see the error if we rewrite the statements above like this,

```
if (c1)
  body-1
else
  if (c2)
    body-2
  endif
```

using the indentation to show how Octave groups the statements. See Chapter 11 [Functions and Scripts], page 205.

## 10.2 The switch Statement

It is very common to take different actions depending on the value of one variable. This is possible using the `if` statement in the following way

```
if (X == 1)
  do_something ();
elseif (X == 2)
  do_something_else ();
else
  do_something_completely_different ();
endif
```

This kind of code can however be very cumbersome to both write and maintain. To overcome this problem Octave supports the `switch` statement. Using this statement, the above example becomes

```
switch (X)
  case 1
    do_something ();
  case 2
    do_something_else ();
  otherwise
    do_something_completely_different ();
endswitch
```

This code makes the repetitive structure of the problem more explicit, making the code easier to read, and hence maintain. Also, if the variable `X` should change its name, only one line would need changing compared to one line per case when `if` statements are used.

The general form of the `switch` statement is

```
switch (expression)
  case label
    command_list
  case label
    command_list
  ...

  otherwise
    command_list
endswitch
```

where *label* can be any expression. However, duplicate *label* values are not detected, and only the *command_list* corresponding to the first match will be executed. For the `switch` statement to be meaningful at least one `case` *`label command_list`* clause must be present, while the `otherwise` *`command_list`* clause is optional.

If *label* is a cell array the corresponding *command_list* is executed if *any* of the elements of the cell array match *expression*. As an example, the following program will print '`Variable is either 6 or 7`'.

```
A = 7;
switch (A)
  case { 6, 7 }
    printf ("variable is either 6 or 7\n");
  otherwise
    printf ("variable is neither 6 nor 7\n");
endswitch
```

As with all other specific `end` keywords, `endswitch` may be replaced by `end`, but you can get better diagnostics if you use the specific forms.

One advantage of using the `switch` statement compared to using `if` statements is that the *label*s can be strings. If an `if` statement is used it is *not* possible to write

```
if (X == "a string") # This is NOT valid
```

since a character-to-character comparison between `X` and the string will be made instead of evaluating if the strings are equal. This special-case is handled by the `switch` statement, and it is possible to write programs that look like this

```
switch (X)
  case "a string"
    do_something
  ...
endswitch
```

### 10.2.1 Notes for the C Programmer

The `switch` statement is also available in the widely used C programming language. There are, however, some differences between the statement in Octave and C

- Cases are exclusive, so they don't 'fall through' as do the cases in the `switch` statement of the C language.
- The *command_list* elements are not optional. Making the list optional would have meant requiring a separator between the label and the command list. Otherwise, things like

  ```
  switch (foo)
    case (1) -2
    ...
  ```

  would produce surprising results, as would

  ```
  switch (foo)
    case (1)
    case (2)
      doit ();
    ...
  ```

particularly for C programmers. If `doit()` should be executed if *foo* is either `1` or `2`, the above code should be written with a cell array like this

```
switch (foo)
  case { 1, 2 }
    doit ();
  ...
```

## 10.3 The while Statement

In programming, a *loop* means a part of a program that is (or at least can be) executed two or more times in succession.

The `while` statement is the simplest looping statement in Octave. It repeatedly executes a statement as long as a condition is true. As with the condition in an `if` statement, the condition in a `while` statement is considered true if its value is nonzero, and false if its value is zero. If the value of the conditional expression in a `while` statement is a vector or a matrix, it is considered true only if it is non-empty and *all* of the elements are nonzero.

Octave's `while` statement looks like this:

```
while (condition)
  body
endwhile
```

Here *body* is a statement or list of statements that we call the *body* of the loop, and *condition* is an expression that controls how long the loop keeps running.

The first thing the `while` statement does is test *condition*. If *condition* is true, it executes the statement *body*. After *body* has been executed, *condition* is tested again, and if it is still true, *body* is executed again. This process repeats until *condition* is no longer true. If *condition* is initially false, the body of the loop is never executed.

This example creates a variable `fib` that contains the first ten elements of the Fibonacci sequence.

```
fib = ones (1, 10);
i = 3;
while (i <= 10)
  fib (i) = fib (i-1) + fib (i-2);
  i++;
endwhile
```

Here the body of the loop contains two statements.

The loop works like this: first, the value of `i` is set to 3. Then, the `while` tests whether `i` is less than or equal to 10. This is the case when `i` equals 3, so the value of the `i`-th element of `fib` is set to the sum of the previous two values in the sequence. Then the `i++` increments the value of `i` and the loop repeats. The loop terminates when `i` reaches 11.

A newline is not required between the condition and the body; but using one makes the program clearer unless the body is very simple.

## 10.4 The do-until Statement

The `do-until` statement is similar to the `while` statement, except that it repeatedly executes a statement until a condition becomes true, and the test of the condition is at the end of the loop, so the body of the loop is always executed at least once. As with the condition in an `if` statement, the condition in a `do-until` statement is considered true if its value is nonzero, and false if its value is zero. If the value of the conditional expression in a `do-until` statement is a vector or a matrix, it is considered true only if it is non-empty and *all* of the elements are nonzero.

Octave's `do-until` statement looks like this:

```
do
  body
until (condition)
```

Here *body* is a statement or list of statements that we call the *body* of the loop, and *condition* is an expression that controls how long the loop keeps running.

This example creates a variable `fib` that contains the first ten elements of the Fibonacci sequence.

```
fib = ones (1, 10);
i = 2;
do
  i++;
  fib (i) = fib (i-1) + fib (i-2);
until (i == 10)
```

A newline is not required between the `do` keyword and the body; but using one makes the program clearer unless the body is very simple.

## 10.5 The for Statement

The `for` statement makes it more convenient to count iterations of a loop. The general form of the `for` statement looks like this:

```
for var = expression
  body
endfor
```

where *body* stands for any statement or list of statements, *expression* is any valid expression, and *var* may take several forms. Usually it is a simple variable name or an indexed variable. If the value of *expression* is a structure, *var* may also be a vector with two elements. See Section 10.5.1 [Looping Over Structure Elements], page 199, below.

The assignment expression in the `for` statement works a bit differently than Octave's normal assignment statement. Instead of assigning the complete result of the expression, it assigns each column of the expression to *var* in turn. If *expression* is a range, a row vector, or a scalar, the value of *var* will be a scalar each time the loop body is executed. If *var* is a column vector or a matrix, *var* will be a column vector each time the loop body is executed.

The following example shows another way to create a vector containing the first ten elements of the Fibonacci sequence, this time using the `for` statement:

```
fib = ones (1, 10);
for i = 3:10
  fib(i) = fib(i-1) + fib(i-2);
endfor
```

This code works by first evaluating the expression `3:10`, to produce a range of values from 3 to 10 inclusive. Then the variable `i` is assigned the first element of the range and the body of the loop is executed once. When the end of the loop body is reached, the next value in the range is assigned to the variable `i`, and the loop body is executed again. This process continues until there are no more elements to assign.

Within Octave is it also possible to iterate over matrices or cell arrays using the `for` statement. For example consider

```
disp ("Loop over a matrix")
for i = [1,3;2,4]
  i
endfor
disp ("Loop over a cell array")
for i = {1,"two";"three",4}
  i
endfor
```

In this case the variable `i` takes on the value of the columns of the matrix or cell matrix. So the first loop iterates twice, producing two column vectors `[1;2]`, followed by `[3;4]`, and likewise for the loop over the cell array. This can be extended to loops over multi-dimensional arrays. For example:

```
a = [1,3;2,4]; c = cat (3, a, 2*a);
for i = c
  i
endfor
```

In the above case, the multi-dimensional matrix *c* is reshaped to a two-dimensional matrix as `reshape (c, rows (c), prod (size (c)(2:end)))` and then the same behavior as a loop over a two-dimensional matrix is produced.

Although it is possible to rewrite all `for` loops as `while` loops, the Octave language has both statements because often a `for` loop is both less work to type and more natural to think of. Counting the number of iterations is very common in loops and it can be easier to think of this counting as part of looping rather than as something to do inside the loop.

### 10.5.1 Looping Over Structure Elements

A special form of the `for` statement allows you to loop over all the elements of a structure:

```
for [ val, key ] = expression
  body
endfor
```

In this form of the `for` statement, the value of *expression* must be a structure. If it is, *key* and *val* are set to the name of the element and the corresponding value in turn, until there are no more elements. For example:

```
x.a = 1
x.b = [1, 2; 3, 4]
x.c = "string"
for [val, key] = x
  key
  val
endfor
```

```
⊣ key = a
⊣ val = 1
⊣ key = b
⊣ val =
⊣
⊣   1  2
⊣   3  4
⊣
⊣ key = c
⊣ val = string
```

The elements are not accessed in any particular order. If you need to cycle through the list in a particular way, you will have to use the function `fieldnames` and sort the list yourself.

## 10.6 The break Statement

The `break` statement jumps out of the innermost `while`, `do-until`, or `for` loop that encloses it. The `break` statement may only be used within the body of a loop. The following example finds the smallest divisor of a given integer, and also identifies prime numbers:

```
num = 103;
div = 2;
while (div*div <= num)
  if (rem (num, div) == 0)
    break;
  endif
  div++;
endwhile
if (rem (num, div) == 0)
  printf ("Smallest divisor of %d is %d\n", num, div)
else
  printf ("%d is prime\n", num);
endif
```

When the remainder is zero in the first `while` statement, Octave immediately *breaks out* of the loop. This means that Octave proceeds immediately to the statement following the loop and continues processing. (This is very different from the `exit` statement which stops the entire Octave program.)

Here is another program equivalent to the previous one. It illustrates how the *condition* of a `while` statement could just as well be replaced with a `break` inside an `if`:

```
num = 103;
div = 2;
while (1)
  if (rem (num, div) == 0)
    printf ("Smallest divisor of %d is %d\n", num, div);
    break;
  endif
  div++;
  if (div*div > num)
    printf ("%d is prime\n", num);
    break;
  endif
endwhile
```

## 10.7 The continue Statement

The `continue` statement, like `break`, is used only inside `while`, `do-until`, or `for` loops. It skips over the rest of the loop body, causing the next cycle around the loop to begin immediately. Contrast this with `break`, which jumps out of the loop altogether. Here is an example:

```
# print elements of a vector of random
# integers that are even.

# first, create a row vector of 10 random
# integers with values between 0 and 100:

vec = round (rand (1, 10) * 100);

# print what we're interested in:

for x = vec
  if (rem (x, 2) != 0)
    continue;
  endif
  printf ("%d\n", x);
endfor
```

If one of the elements of *vec* is an odd number, this example skips the print statement for that element, and continues back to the first statement in the loop.

This is not a practical example of the `continue` statement, but it should give you a clear understanding of how it works. Normally, one would probably write the loop like this:

```
for x = vec
  if (rem (x, 2) == 0)
    printf ("%d\n", x);
  endif
endfor
```

## 10.8 The unwind_protect Statement

Octave supports a limited form of exception handling modeled after the unwind-protect form of Lisp.

The general form of an `unwind_protect` block looks like this:

```
unwind_protect
  body
unwind_protect_cleanup
  cleanup
end_unwind_protect
```

where *body* and *cleanup* are both optional and may contain any Octave expressions or commands. The statements in *cleanup* are guaranteed to be executed regardless of how control exits *body*.

This is useful to protect temporary changes to global variables from possible errors. For example, the following code will always restore the original value of the global variable `frobnosticate` even if an error occurs in the first part of the `unwind_protect` block.

```
save_frobnosticate = frobnosticate;
unwind_protect
  frobnosticate = true;
  ...
unwind_protect_cleanup
  frobnosticate = save_frobnosticate;
end_unwind_protect
```

Without `unwind_protect`, the value of *frobnosticate* would not be restored if an error occurs while evaluating the first part of the `unwind_protect` block because evaluation would stop at the point of the error and the statement to restore the value would not be executed.

In addition to unwind_protect, Octave supports another form of exception handling, the `try` block.

## 10.9 The try Statement

The original form of a `try` block looks like this:

```
try
  body
catch
  cleanup
end_try_catch
```

where *body* and *cleanup* are both optional and may contain any Octave expressions or commands. The statements in *cleanup* are only executed if an error occurs in *body*.

No warnings or error messages are printed while *body* is executing. If an error does occur during the execution of *body*, *cleanup* can use the functions `lasterr` or `lasterror` to access the text of the message that would have been printed, as well as its identifier. The alternative form,

```
try
  body
catch err
  cleanup
end_try_catch
```

will automatically store the output of `lasterror` in the structure *err*. See Chapter 12 [Errors and Warnings], page 255, for more information about the `lasterr` and `lasterror` functions.

## 10.10 Continuation Lines

In the Octave language, most statements end with a newline character and you must tell Octave to ignore the newline character in order to continue a statement from one line to the next. Lines that end with the characters `...` are joined with the following line before they are divided into tokens by Octave's parser. For example, the lines

```
x = long_variable_name ...
    + longer_variable_name ...
    - 42
```

form a single statement.

Any text between the continuation marker and the newline character is ignored. For example, the statement

```
x = long_variable_name ...    # comment one
    + longer_variable_name ...comment two
    - 42                      # last comment
```

is equivalent to the one shown above.

Inside double-quoted string constants, the character `\` has to be used as continuation marker. The `\` must appear at the end of the line just before the newline character:

```
s = "This text starts in the first line \
and is continued in the second line."
```

Input that occurs inside parentheses can be continued to the next line without having to use a continuation marker. For example, it is possible to write statements like

```
if (fine_dining_destination == on_a_boat
    || fine_dining_destination == on_a_train)
  seuss (i, will, not, eat, them, sam, i, am, i,
         will, not, eat, green, eggs, and, ham);
endif
```

without having to add to the clutter with continuation markers.

# 11 Functions and Scripts

Complicated Octave programs can often be simplified by defining functions. Functions can be defined directly on the command line during interactive Octave sessions, or in external files, and can be called just like built-in functions.

## 11.1 Introduction to Function and Script Files

There are seven different things covered in this section.

1. Typing in a function at the command prompt.
2. Storing a group of commands in a file — called a script file.
3. Storing a function in a file—called a function file.
4. Subfunctions in function files.
5. Multiple functions in one script file.
6. Private functions.
7. Nested functions.

Both function files and script files end with an extension of .m, for MATLAB compatibility. If you want more than one independent functions in a file, it must be a script file (see Section 11.11 [Script Files], page 238), and to use these functions you must execute the script file before you can use the functions that are in the script file.

## 11.2 Defining Functions

In its simplest form, the definition of a function named *name* looks like this:

```
function name
  body
endfunction
```

A valid function name is like a valid variable name: a sequence of letters, digits and underscores, not starting with a digit. Functions share the same pool of names as variables.

The function *body* consists of Octave statements. It is the most important part of the definition, because it says what the function should actually *do*.

For example, here is a function that, when executed, will ring the bell on your terminal (assuming that it is possible to do so):

```
function wakeup
  printf ("\a");
endfunction
```

The `printf` statement (see Chapter 14 [Input and Output], page 287) simply tells Octave to print the string `"\a"`. The special character '`\a`' stands for the alert character (ASCII 7). See Chapter 5 [Strings], page 75.

Once this function is defined, you can ask Octave to evaluate it by typing the name of the function.

Normally, you will want to pass some information to the functions you define. The syntax for passing parameters to a function in Octave is

```
function name (arg-list)
  body
endfunction
```

where *arg-list* is a comma-separated list of the function's arguments. When the function is called, the argument names are used to hold the argument values given in the call. The list of arguments may be empty, in which case this form is equivalent to the one shown above.

To print a message along with ringing the bell, you might modify the `wakeup` to look like this:

```
function wakeup (message)
  printf ("\a%s\n", message);
endfunction
```

Calling this function using a statement like this

```
wakeup ("Rise and shine!");
```

will cause Octave to ring your terminal's bell and print the message '`Rise and shine!`', followed by a newline character (the '`\n`' in the first argument to the `printf` statement).

In most cases, you will also want to get some information back from the functions you define. Here is the syntax for writing a function that returns a single value:

```
function ret-var = name (arg-list)
  body
endfunction
```

The symbol *ret-var* is the name of the variable that will hold the value to be returned by the function. This variable must be defined before the end of the function body in order for the function to return a value.

Variables used in the body of a function are local to the function. Variables named in *arg-list* and *ret-var* are also local to the function. See Section 7.1 [Global Variables], page 147, for information about how to access global variables inside a function.

For example, here is a function that computes the average of the elements of a vector:

```
function retval = avg (v)
  retval = sum (v) / length (v);
endfunction
```

If we had written `avg` like this instead,

```
function retval = avg (v)
  if (isvector (v))
    retval = sum (v) / length (v);
  endif
endfunction
```

and then called the function with a matrix instead of a vector as the argument, Octave would have printed an error message like this:

```
error: value on right hand side of assignment is undefined
```

because the body of the `if` statement was never executed, and `retval` was never defined. To prevent obscure errors like this, it is a good idea to always make sure that the return

variables will always have values, and to produce meaningful error messages when problems are encountered. For example, `avg` could have been written like this:

```
function retval = avg (v)
  retval = 0;
  if (isvector (v))
    retval = sum (v) / length (v);
  else
    error ("avg: expecting vector argument");
  endif
endfunction
```

There is still one remaining problem with this function. What if it is called without an argument? Without additional error checking, Octave will probably print an error message that won't really help you track down the source of the error. To allow you to catch errors like this, Octave provides each function with an automatic variable called `nargin`. Each time a function is called, `nargin` is automatically initialized to the number of arguments that have actually been passed to the function. For example, we might rewrite the `avg` function like this:

```
function retval = avg (v)
  retval = 0;
  if (nargin != 1)
    usage ("avg (vector)");
  endif
  if (isvector (v))
    retval = sum (v) / length (v);
  else
    error ("avg: expecting vector argument");
  endif
endfunction
```

Octave automatically reports an error for functions written in .m file code if they are called with more arguments than expected. Octave does not automatically report an error if a function is called with too few arguments, since functions in general may have default arguments, but any attempt to use a variable that has not been given a value will result in an error. Functions can check the arguments they are called with to avoid such problems and to provide more context-specific error messages.

*n* = nargin ()
*n* = nargin (*fcn*)

Report the number of input arguments to a function.

Called from within a function, return the number of arguments passed to the function. At the top level, return the number of command line arguments passed to Octave.

If called with the optional argument *fcn*—a function name or handle—return the declared number of arguments that the function can accept.

If the last argument to *fcn* is *varargin* the returned value is negative. For example, the function `union` for sets is declared as

```
function [y, ia, ib] = union (a, b, varargin)

and

nargin ("union")
⇒ -3
```

Programming Note: `nargin` does not work on compiled functions (`.oct` files) such as built-in or dynamically loaded functions.

**See also:** [nargout], page 211, [narginchk], page 217, [varargin], page 213, [inputname], page 208.

*namestr* = inputname (*n*)
*namestr* = inputname (*n*, *ids_only*)

Return the name of the *n*-th argument to the calling function.

If the argument is not a simple variable name, return an empty string. Examples which will return `""` are numbers (`5.1`), expressions (*y*`/2`), and cell or structure indexing (*c*`{1}` or *s*.*field*).

`inputname` is only useful within a function. When used at the command line or within a script it always returns an empty string.

By default, return an empty string if the *n*-th argument is not a valid variable name. If the optional argument *ids_only* is false, return the text of the argument even if it is not a valid variable name. This is an Octave extension that allows the programmer to view exactly how the function was invoked even when the inputs are complex expressions.

**See also:** [nargin], page 207, [narginchk], page 217.

*val* = silent_functions ()
*old_val* = silent_functions (*new_val*)
*old_val* = silent_functions (*new_val*, "*local*")

Query or set the internal variable that controls whether internal output from a function is suppressed.

If this option is disabled, Octave will display the results produced by evaluating expressions within a function body that are not terminated with a semicolon.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

## 11.3 Returning from a Function

The body of a user-defined function can contain a `return` statement. This statement returns control to the rest of the Octave program. It looks like this:

```
return
```

Unlike the `return` statement in C, Octave's `return` statement cannot be used to return a value from a function. Instead, you must assign values to the list of return variables that are part of the `function` statement. The `return` statement simply makes it easier to exit a function from a deeply nested loop or conditional statement.

Here is an example of a function that checks to see if any elements of a vector are nonzero.

```
function retval = any_nonzero (v)
  retval = 0;
  for i = 1:length (v)
    if (v (i) != 0)
      retval = 1;
      return;
    endif
  endfor
  printf ("no nonzero elements found\n");
endfunction
```

Note that this function could not have been written using the `break` statement to exit the loop once a nonzero value is found without adding extra logic to avoid printing the message if the vector does contain a nonzero element.

`return`

When Octave encounters the keyword `return` inside a function or script, it returns control to the caller immediately. At the top level, the return statement is ignored. A `return` statement is assumed at the end of every function definition.

## 11.4 Multiple Return Values

Unlike many other computer languages, Octave allows you to define functions that return more than one value. The syntax for defining functions that return multiple values is

```
function [ret-list] = name (arg-list)
  body
endfunction
```

where *name*, *arg-list*, and *body* have the same meaning as before, and *ret-list* is a comma-separated list of variable names that will hold the values returned from the function. The list of return values must have at least one element. If *ret-list* has only one element, this form of the `function` statement is equivalent to the form described in the previous section.

Here is an example of a function that returns two values, the maximum element of a vector and the index of its first occurrence in the vector.

```
function [max, idx] = vmax (v)
  idx = 1;
  max = v (idx);
  for i = 2:length (v)
    if (v (i) > max)
      max = v (i);
      idx = i;
    endif
  endfor
endfunction
```

In this particular case, the two values could have been returned as elements of a single array, but that is not always possible or convenient. The values to be returned may not

have compatible dimensions, and it is often desirable to give the individual return values distinct names.

It is possible to use the `nthargout` function to obtain only some of the return values or several at once in a cell array. See Section 3.1.5 [Cell Array Objects], page 46.

*arg* = nthargout (*n*, *fcn*, ...)
*arg* = nthargout (*n*, *ntot*, *fcn*, ...)

Return the *n*th output argument of the function specified by the function handle or string *fcn*.

Any additional arguments are passed directly to *fcn*. The total number of arguments to call *fcn* with can be passed in *ntot*; by default *ntot* is *n*. The input *n* can also be a vector of indices of the output, in which case the output will be a cell array of the requested output arguments.

The intended use of `nthargout` is to avoid intermediate variables. For example, when finding the indices of the maximum entry of a matrix, the following two compositions of `nthargout`

```
m = magic (5);
cell2mat (nthargout ([1, 2], @ind2sub, size (m),
                     nthargout (2, @max, m(:))))
⇒ 5   3
```

are completely equivalent to the following lines:

```
m = magic (5);
[~, idx] = max (M(:));
[i, j] = ind2sub (size (m), idx);
[i, j]
⇒ 5   3
```

It can also be helpful to have all output arguments collected in a single cell array as the following code demonstrates:

```
USV = nthargout ([1:3], @svd, hilb (5));
```

Programming Note: Equivalent functionality to the `nthargout` function is frequently possible by using the character '`~`' in code to ignore outputs. This functionality is implemented by the Octave interpreter and is even more efficient than using this function.

Equivalent Code:

```
idx = nthargout (2, @max, rand (100, 1));
≡
[~, idx] = max (rand (100, 1));
```

**See also:** [nargin], page 207, [nargout], page 211, [varargin], page 213, [varargout], page 212, [isargout], page 215.

In addition to setting `nargin` each time a function is called, Octave also automatically initializes `nargout` to the number of values that are expected to be returned. This allows you to write functions that behave differently depending on the number of values that the user of the function has requested. The implicit assignment to the built-in variable `ans` does not figure in the count of output arguments, so the value of `nargout` may be zero.

The `svd` and `hist` functions are examples of built-in functions that behave differently depending on the value of `nargout`. For example, `hist` will draw a histogram when called with no output variables, but if called with outputs it will return the frequency counts and/or bin centers without creating a plot.

It is possible to write functions that only set some return values. For example, calling the function

```
function [x, y, z] = f ()
  x = 1;
  z = 2;
endfunction
```

as

```
[a, b, c] = f ()
```

produces:

```
a = 1

b = [](0x0)

c = 2
```

along with a warning.

*n* = `nargout ()`
*n* = `nargout (`*fcn*`)`

Report the number of output arguments from a function.

Called from within a function, return the number of values the caller expects to receive. At the top level, `nargout` with no argument is undefined and will produce an error.

If called with the optional argument *fcn*—a function name or handle—return the number of declared output values that the function can produce.

If the final output argument is *varargout* the returned value is negative.

For example,

```
f ()
```

will cause `nargout` to return 0 inside the function `f` and

```
[s, t] = f ()
```

will cause `nargout` to return 2 inside the function `f`.

In the second usage,

```
nargout (@histc)   # or nargout ("histc") using a string input
```

will return 2, because `histc` has two outputs, whereas

```
nargout (@imread)
```

will return -2, because `imread` has two outputs and the second is *varargout*.

Programming Note. `nargout` does not work for built-in functions and returns -1 for all anonymous functions.

**See also:** [nargin], page 207, [varargout], page 212, [isargout], page 215, [nthargout], page 210.

## 11.5 Variable-length Return Lists

It is possible to return a variable number of output arguments from a function using a syntax that's similar to the one used with the special `varargin` parameter name. To let a function return a variable number of output arguments the special output parameter name `varargout` is used. As with `varargin`, `varargout` is a cell array that will contain the requested output arguments.

As an example the following function sets the first output argument to 1, the second to 2, and so on.

```
function varargout = one_to_n ()
  for i = 1:nargout
    varargout{i} = i;
  endfor
endfunction
```

When called this function returns values like this

```
[a, b, c] = one_to_n ()
     ⇒ a =  1
     ⇒ b =  2
     ⇒ c =  3
```

If `varargin` (`varargout`) does not appear as the last element of the input (output) parameter list, then it is not special, and is handled the same as any other parameter name.

`[`*r1*`, `*r2*`, ..., `*rn*`] = deal (`*a*`)`
`[`*r1*`, `*r2*`, ..., `*rn*`] = deal (`*a1*`, `*a2*`, ..., `*an*`)`

Copy the input parameters into the corresponding output parameters.

If only a single input parameter is supplied, its value is copied to each of the outputs.

For example,

```
[a, b, c] = deal (x, y, z);
```

is equivalent to

```
a = x;
b = y;
c = z;
```

and

```
[a, b, c] = deal (x);
```

is equivalent to

```
a = b = c = x;
```

Programming Note: `deal` is often used with comma-separated lists derived from cell arrays or structures. This is unnecessary as the interpreter can perform the same action without the overhead of a function call. For example:

```
c = {[1 2], "Three", 4};
[x, y, z] = c{:}
⇒
   x =

      1   2

   y = Three
   z =  4
```

**See also:** [cell2struct], page 141, [struct2cell], page 129, [repmat], page 582.

## 11.6 Variable-length Argument Lists

Sometimes the number of input arguments is not known when the function is defined. As an example think of a function that returns the smallest of all its input arguments. For example:

```
a = smallest (1, 2, 3);
b = smallest (1, 2, 3, 4);
```

In this example both `a` and `b` would be 1. One way to write the `smallest` function is

```
function val = smallest (arg1, arg2, arg3, arg4, arg5)
  body
endfunction
```

and then use the value of `nargin` to determine which of the input arguments should be considered. The problem with this approach is that it can only handle a limited number of input arguments.

If the special parameter name `varargin` appears at the end of a function parameter list it indicates that the function takes a variable number of input arguments. Using `varargin` the function looks like this

```
function val = smallest (varargin)
  body
endfunction
```

In the function body the input arguments can be accessed through the variable `varargin`. This variable is a cell array containing all the input arguments. See Section 6.3 [Cell Arrays], page 131, for details on working with cell arrays. The `smallest` function can now be defined like this

```
function val = smallest (varargin)
  val = min ([varargin{:}]);
endfunction
```

This implementation handles any number of input arguments, but it's also a very simple solution to the problem.

A slightly more complex example of `varargin` is a function `print_arguments` that prints all input arguments. Such a function can be defined like this

```
function print_arguments (varargin)
  for i = 1:length (varargin)
    printf ("Input argument %d: ", i);
    disp (varargin{i});
  endfor
endfunction
```

This function produces output like this

```
print_arguments (1, "two", 3);
     ⊣ Input argument 1:  1
     ⊣ Input argument 2: two
     ⊣ Input argument 3:  3
```

`[`*reg*`, `*prop*`] = parseparams (`*params*`)`
`[`*reg*`, `*var1*`, ...] = parseparams (`*params*`, `*name1*`, `*default1*`, ...)`

Return in *reg* the cell elements of *param* up to the first string element and in *prop* all remaining elements beginning with the first string element.

For example:

```
[reg, prop] = parseparams ({1, 2, "linewidth", 10})
reg =
{
  [1,1] = 1
  [1,2] = 2
}
prop =
{
  [1,1] = linewidth
  [1,2] = 10
}
```

The parseparams function may be used to separate regular numeric arguments from additional arguments given as property/value pairs of the *varargin* cell array.

In the second form of the call, available options are specified directly with their default values given as name-value pairs. If *params* do not form name-value pairs, or if an option occurs that does not match any of the available options, an error occurs.

When called from an m-file function, the error is prefixed with the name of the caller function.

The matching of options is case-insensitive.

**See also:** [varargin], page 213, [inputParser], page 222.

## 11.7 Ignoring Arguments

In the formal argument list, it is possible to use the dummy placeholder `~` instead of a name. This indicates that the corresponding argument value should be ignored and not stored to any variable.

```
function val = pick2nd (~, arg2)
  val = arg2;
endfunction
```

The value of `nargin` is not affected by using this declaration.

Return arguments can also be ignored using the same syntax. For example, the sort function returns both the sorted values, and an index vector for the original input which will result in a sorted output. Ignoring the second output is simple—don't request more than one output. But ignoring the first, and calculating just the second output, requires the use of the ~ placeholder.

```
x = [2, 3, 1];
[s, i] = sort (x)
⇒
s =

   1   2   3

i =

   3   1   2

[~, i] = sort (x)
⇒
i =

   3   1   2
```

When using the ~ placeholder, commas—not whitespace—must be used to separate output arguments. Otherwise, the interpreter will view ~ as the logical not operator.

```
[~ i] = sort (x)
parse error:

  invalid left hand side of assignment
```

Functions may take advantage of ignored outputs to reduce the number of calculations performed. To do so, use the `isargout` function to query whether the output argument is wanted. For example:

```
function [out1, out2] = long_function (x, y, z)
  if (isargout (1))
    ## Long calculation
    ...
    out1 = result;
  endif
  ...
endfunction
```

*tf* = isargout (*k*)

Within a function, return a logical value indicating whether the argument *k* will be assigned to a variable on output.

If the result is false, the argument has been ignored during the function call through the use of the tilde (~) special output argument. Functions can use `isargout` to avoid performing unnecessary calculations for outputs which are unwanted.

If $k$ is outside the range `1:max (nargout)`, the function returns false. $k$ can also be an array, in which case the function works element-by-element and a logical array is returned. At the top level, `isargout` returns an error.

**See also:** [nargout], page 211, [varargout], page 212, [nthargout], page 210.

## 11.8 Default Arguments

Since Octave supports variable number of input arguments, it is very useful to assign default values to some input arguments. When an input argument is declared in the argument list it is possible to assign a default value to the argument like this

```
function name (arg1 = val1, ...)
  body
endfunction
```

If no value is assigned to *arg1* by the user, it will have the value *val1*.

As an example, the following function implements a variant of the classic "Hello, World" program.

```
function hello (who = "World")
  printf ("Hello, %s!\n", who);
endfunction
```

When called without an input argument the function prints the following

```
hello ();
     ⊣ Hello, World!
```

and when it's called with an input argument it prints the following

```
hello ("Beautiful World of Free Software");
     ⊣ Hello, Beautiful World of Free Software!
```

Sometimes it is useful to explicitly tell Octave to use the default value of an input argument. This can be done writing a ':' as the value of the input argument when calling the function.

```
hello (:);
     ⊣ Hello, World!
```

## 11.9 Validating Arguments

Octave is a weakly typed programming language. Thus it is possible to call a function with arguments, that probably cause errors or might have undesirable side effects. For example calling a string processing function with a huge sparse matrix.

It is good practice at the head of a function to verify that it has been called correctly. Octave offers several functions for this purpose.

### 11.9.1 Validating the number of Arguments

In Octave the following idiom is seen frequently at the beginning of a function definition:

```
if (nargin < min_#_inputs || nargin > max_#_inputs)
  print_usage ();
endif
```

which stops the function execution and prints a message about the correct way to call the function whenever the number of inputs is wrong.

Similar error checking is provided by `narginchk` and `nargoutchk`.

`narginchk` (*minargs*, *maxargs*)

Check for correct number of input arguments.

Generate an error message if the number of arguments in the calling function is outside the range *minargs* and *maxargs*. Otherwise, do nothing.

Both *minargs* and *maxargs* must be scalar numeric values. Zero, Inf, and negative values are all allowed, and *minargs* and *maxargs* may be equal.

Note that this function evaluates `nargin` on the caller.

**See also:** [nargoutchk], page 217, [error], page 255, [nargout], page 211, [nargin], page 207.

`nargoutchk` (*minargs*, *maxargs*)
*msgstr* = `nargoutchk` (*minargs*, *maxargs*, *nargs*)
*msgstr* = `nargoutchk` (*minargs*, *maxargs*, *nargs*, "*string*")
*msgstruct* = `nargoutchk` (*minargs*, *maxargs*, *nargs*, "*struct*")

Check for correct number of output arguments.

In the first form, return an error if the number of arguments is not between *minargs* and *maxargs*. Otherwise, do nothing. Note that this function evaluates the value of `nargout` on the caller so its value must have not been tampered with.

Both *minargs* and *maxargs* must be numeric scalars. Zero, Inf, and negative are all valid, and they can have the same value.

For backwards compatibility, the other forms return an appropriate error message string (or structure) if the number of outputs requested is invalid.

This is useful for checking to that the number of output arguments supplied to a function is within an acceptable range.

**See also:** [narginchk], page 217, [error], page 255, [nargout], page 211, [nargin], page 207.

### 11.9.2 Validating the type of Arguments

Besides the number of arguments, inputs can be checked for various properties. `validatestring` is used for string arguments and `validateattributes` for numeric arguments.

*validstr* = `validatestring` (*str*, *strarray*)
*validstr* = `validatestring` (*str*, *strarray*, *funcname*)
*validstr* = `validatestring` (*str*, *strarray*, *funcname*, *varname*)
*validstr* = `validatestring` (*...*, *position*)

Verify that *str* is an element, or substring of an element, in *strarray*.

When *str* is a character string to be tested, and *strarray* is a cell array of strings of valid values, then *validstr* will be the validated form of *str* where validation is defined as *str* being a member or substring of *validstr*. This is useful for both verifying and expanding short options, such as `"r"`, to their longer forms, such as `"red"`. If *str* is

a substring of *validstr*, and there are multiple matches, the shortest match will be returned if all matches are substrings of each other. Otherwise, an error will be raised because the expansion of *str* is ambiguous. All comparisons are case insensitive.

The additional inputs *funcname*, *varname*, and *position* are optional and will make any generated validation error message more specific.

Examples:

```
validatestring ("r", {"red", "green", "blue"})
⇒ "red"

validatestring ("b", {"red", "green", "blue", "black"})
⇒ error: validatestring: multiple unique matches were found for 'b':
   blue, black
```

**See also:** [strcmp], page 88, [strcmpi], page 89, [validateattributes], page 218, [inputParser], page 222.

`validateattributes` (*A*, *classes*, *attributes*)
`validateattributes` (*A*, *classes*, *attributes*, *arg_idx*)
`validateattributes` (*A*, *classes*, *attributes*, *func_name*)
`validateattributes` (*A*, *classes*, *attributes*, *func_name*, *arg_name*)
`validateattributes` (*A*, *classes*, *attributes*, *func_name*, *arg_name*, *arg_idx*)

Check validity of input argument.

Confirms that the argument *A* is valid by belonging to one of *classes*, and holding all of the *attributes*. If it does not, an error is thrown, with a message formatted accordingly. The error message can be made further complete by the function name *fun_name*, the argument name *arg_name*, and its position in the input *arg_idx*.

*classes* must be a cell array of strings (an empty cell array is allowed) with the name of classes (remember that a class name is case sensitive). In addition to the class name, the following categories names are also valid:

`"float"` Floating point value comprising classes `"double"` and `"single"`.

`"integer"`
: Integer value comprising classes (u)int8, (u)int16, (u)int32, (u)int64.

`"numeric"`
: Numeric value comprising either a floating point or integer value.

*attributes* must be a cell array with names of checks for *A*. Some of them require an additional value to be supplied right after the name (see details for each below).

`"<="` All values are less than or equal to the following value in *attributes*.

`"<"` All values are less than the following value in *attributes*.

`">="` All values are greater than or equal to the following value in *attributes*.

`">"` All values are greater than the following value in *attributes*.

`"2d"` A 2-dimensional matrix. Note that vectors and empty matrices have 2 dimensions, one of them being of length 1, or both length 0.

`"3d"` Has no more than 3 dimensions. A 2-dimensional matrix is a 3-D matrix whose 3rd dimension is of length 1.

`"binary"` All values are either 1 or 0.

`"column"` Values are arranged in a single column.

`"decreasing"`
No value is *NaN*, and each is less than the preceding one.

`"diag"` Value is a diagonal matrix.

`"even"` All values are even numbers.

`"finite"` All values are finite.

`"increasing"`
No value is *NaN*, and each is greater than the preceding one.

`"integer"`
All values are integer. This is different than using `isinteger` which only checks its an integer type. This checks that each value in *A* is an integer value, i.e., it has no decimal part.

`"ncols"` Has exactly as many columns as the next value in *attributes*.

`"ndims"` Has exactly as many dimensions as the next value in *attributes*.

`"nondecreasing"`
No value is *NaN*, and each is greater than or equal to the preceding one.

`"nonempty"`
It is not empty.

`"nonincreasing"`
No value is *NaN*, and each is less than or equal to the preceding one.

`"nonnan"` No value is a `NaN`.

`"nonnegative"`
All values are non negative.

`"nonsparse"`
It is not a sparse matrix.

`"nonzero"`
No value is zero.

`"nrows"` Has exactly as many rows as the next value in *attributes*.

`"numel"` Has exactly as many elements as the next value in *attributes*.

`"odd"` All values are odd numbers.

`"positive"`
All values are positive.

`"real"` It is a non-complex matrix.

`"row"` Values are arranged in a single row.

`"scalar"` It is a scalar.

`"size"` Its size has length equal to the values of the next in *attributes*. The next value must is an array with the length for each dimension. To ignore the check for a certain dimension, the value of `NaN` can be used.

`"square"` Is a square matrix.

`"vector"` Values are arranged in a single vector (column or vector).

**See also:** [isa], page 41, [validatestring], page 217, [inputParser], page 222.

As alternatives to `validateattributes` there are several shorter convenience functions to check for individual properties.

`mustBeFinite (x)`

Require that input *x* is finite.

Raise an error if any element of the input *x* is not finite, as determined by `isfinite (x)`.

**See also:** [mustBeNonNan], page 221, [isfinite], page 567.

`mustBeGreaterThan (x, c)`

Require that input *x* is greater than *c*.

Raise an error if any element of the input *x* is not greater than *c*, as determined by *x* `>` *c*.

**See also:** [mustBeGreaterThanOrEqual], page 220, [mustBeLessThan], page 220, [gt], page 175.

`mustBeGreaterThanOrEqual (x, c)`

Require that input *x* is greater than or equal to *c*.

Raise an error if any element of the input *x* is not greater than or equal to *c*, as determined by *x* `>=` *c*.

**See also:** [mustBeGreaterThan], page 220, [mustBeLessThanOrEqual], page 221, [ge], page 175.

`mustBeInteger (x)`

Require that input *x* is integer-valued (but not necessarily integer-typed).

Raise an error if any element of the input *x* is not a finite, real, integer-valued numeric value, as determined by various checks.

**See also:** [mustBeNumeric], page 222.

`mustBeLessThan (x, c)`

Require that input *x* is less than *c*.

Raise an error if any element of the input *x* is not less than *c*, as determined by *x* `<` *c*.

**See also:** [mustBeLessThanOrEqual], page 221, [mustBeGreaterThan], page 220, [lt], page 175.

`mustBeLessThanOrEqual (x, c)`

Require that input is less than or equal to a given value.

Raise an error if any element of the input *x* is not less than or equal to *c*, as determined by `x <= c`.

**See also:** [mustBeLessThan], page 220, [mustBeGreaterThanOrEqual], page 220, [le], page 175.

`mustBeMember (x, valid)`

Require that input *x* is a member of a set of given valid values.

Raise an error if any element of the input *x* is not a member of the set *valid*, as determined by `ismember (x)`.

Programming Note: char inputs may behave strangely because of the interaction between chars and cellstrings when calling `ismember` on them. But it will probably "do what you mean" if you just use it naturally. To guarantee operation, convert all char arrays to cellstrings with `cellstr`.

**See also:** [mustBeNonempty], page 221, [ismember], page 875.

`mustBeNegative (x)`

Require that input *x* is negative.

Raise an error if any element of the input *x* is not negative, as determined by `x < 0`.

**See also:** [mustBeNonnegative], page 221.

`mustBeNonempty (x)`

Require that input *x* is nonempty.

Raise an error if the input *x* is empty, as determined by `isempty (x)`.

**See also:** [mustBeMember], page 221, [mustBeNonzero], page 222, [isempty], page 49.

`mustBeNonNan (x)`

Require that input *x* is non-`NaN`.

Raise an error if any element of the input *x* is `NaN`, as determined by `isnan (x)`.

**See also:** [mustBeFinite], page 220, [mustBeNonempty], page 221, [isnan], page 567.

`mustBeNonnegative (x)`

Require that input *x* is not negative.

Raise an error if any element of the input *x* is negative, as determined by `x >= 0`.

**See also:** [mustBeNonzero], page 222, [mustBePositive], page 222.

`mustBeNonpositive (x)`

Require that input *x* is not positive.

Raise an error if any element of the input *x* is positive, as determined by `x <= 0`.

**See also:** [mustBeNegative], page 221, [mustBeNonzero], page 222.

`mustBeNonsparse (x)`

Require that input *x* is not sparse.

Raise an error if the input *x* is sparse, as determined by `issparse (x)`.

**See also:** [issparse], page 734.

`mustBeNonzero (x)`
: Require that input $x$ is not zero.

  Raise an error if any element of the input $x$ is zero, as determined by `x == 0`.

  **See also:** [mustBeNonnegative], page 221, [mustBePositive], page 222.

`mustBeNumeric (x)`
: Require that input $x$ is numeric.

  Raise an error if the input $x$ is not numeric, as determined by `isnumeric (x)`.

  **See also:** [mustBeNumericOrLogical], page 222, [isnumeric], page 69.

`mustBeNumericOrLogical (x)`
: Require that input $x$ is numeric or logical.

  Raise an error if the input $x$ is not numeric or logical, as determined by `isnumeric (x) || islogical (x)`.

  **See also:** [mustBeNumeric], page 222, [isnumeric], page 69, [islogical], page 69.

`mustBePositive (x)`
: Require that input $x$ is positive.

  Raise an error if any element of the input $x$ is not positive, as determined by `x > 0`.

  **See also:** [mustBeNonnegative], page 221, [mustBeNonzero], page 222.

`mustBeReal (x)`
: Require that input $x$ is real.

  Raise an error if the input $x$ is not real, as determined by `isreal (x)`.

  **See also:** [mustBeFinite], page 220, [mustBeNonNan], page 221, [isreal], page 69.

### 11.9.3 Parsing Arguments

If none of the preceding validation functions is sufficient there is also the class `inputParser` which can perform extremely complex input checking for functions.

`p = inputParser ()`
: Create object $p$ of the inputParser class.

  This class is designed to allow easy parsing of function arguments. The class supports four types of arguments:

  1. mandatory (see `addRequired`);
  2. optional (see `addOptional`);
  3. named (see `addParameter`);
  4. switch (see `addSwitch`).

  After defining the function API with these methods, the supplied arguments can be parsed with the `parse` method and the results accessed with the `Results` accessor.

`inputParser.Parameters`
: Return the list of parameter names already defined. (read-only)

`inputParser.Results`
: Return a structure with argument names as fieldnames and corresponding values. (read-only)

`inputParser.Unmatched`
: Return a structure similar to `Results`, but for unmatched parameters. (read-only) See the `KeepUnmatched` property.

`inputParser.UsingDefaults`
: Return cell array with the names of arguments that are using default values. (read-only)

`inputParser.FunctionName` = *`name`*
: Set function name to be used in error messages; Defaults to empty string.

`inputParser.CaseSensitive` = *`boolean`*
: Set whether matching of argument names should be case sensitive; Defaults to false.

`inputParser.KeepUnmatched` = *`boolean`*
: Set whether string arguments which do not match any Parameter are parsed and stored in the `Unmatched` property; Defaults to false. If false, an error will be emitted at the first unrecognized argument and parsing will stop. Note that since `Switch` and `Parameter` arguments can be mixed, it is not possible to know the type of the unmatched argument. Octave assumes that all unmatched arguments are of the `Parameter` type and therefore must be followed by a value.

`inputParser.PartialMatching` = *`boolean`*
: Set whether argument names for `Parameter` and `Switch` options may be given in shortened form as long as the name uniquely identifies an option; Defaults to true. For example, the argument `'opt'` will match a parameter `'opt_color'`, but will fail if there is also a parameter `'opt_case'`.

`inputParser.StructExpand` = *`boolean`*
: Set whether a structure passed to the function is expanded into parameter/value pairs (parameter = fieldname); Defaults to true.

  The following example shows how to use this class:

```
function check (varargin)
  p = inputParser ();                       # create object
  p.FunctionName = "check";                 # set function name
  p.addRequired ("pack", @ischar);          # mandatory argument
  p.addOptional ("path", pwd(), @ischar);  # optional argument

  ## Create anonymous function handle for validators
  valid_vec = @(x) isvector (x) && all (x >= 0) && all (x <= 1);
  p.addOptional ("vec", [0 0], valid_vec);

  ## Create two arguments of type "Parameter"
  vld_type = @(x) any (strcmp (x, {"linear", "quadratic"}));
  p.addParameter ("type", "linear", vld_type);
```

```
  vld_tol = @(x) any (strcmp (x, {"low", "medium", "high"}));
  p.addParameter ("tolerance", "low", vld_tol);

  ## Create a switch type of argument
  p.addSwitch ("verbose");

  p.parse (varargin{:});  # Run created parser on inputs

  ## The rest of the function can access inputs by using p.Results.
  ## For example, get the tolerance input with p.Results.tolerance
endfunction

check ("mech");           # valid, use defaults for other arguments
check ();                 # error, one argument is mandatory
check (1);                # error, since ! ischar
check ("mech", "~/dev");  # valid, use defaults for other arguments

check ("mech", "~/dev", [0 1 0 0], "type", "linear");  # valid

## following is also valid.  Note how the Switch argument type can
## be mixed in with or before the Parameter argument type (but it
## must still appear after any Optional arguments).
check ("mech", "~/dev", [0 1 0 0], "verbose", "tolerance", "high");

## following returns an error since an Optional argument, 'path',
## was given after the Parameter argument 'type'.
check ("mech", "type", "linear", "~/dev");
```

*Note 1*: A function can have any mixture of the four API types but they must appear in a specific order. `Required` arguments must be first and can be followed by any `Optional` arguments. Only the `Parameter` and `Switch` arguments may be mixed together and they must appear following the first two types.

*Note 2*: If both `Optional` and `Parameter` arguments are mixed in a function API then once a string Optional argument fails to validate it will be considered the end of the `Optional` arguments. The remaining arguments will be compared against any `Parameter` or `Switch` arguments.

**See also:** [nargin], page 207, [validateattributes], page 218, [validatestring], page 217, [varargin], page 213.

## 11.10 Function Files

Except for simple one-shot programs, it is not practical to have to define all the functions you need each time you need them. Instead, you will normally want to save them in a file so that you can easily edit them, and save them for use at a later time.

Octave does not require you to load function definitions from files before using them. You simply need to put the function definitions in a place where Octave can find them.

When Octave encounters an identifier that is undefined, it first looks for variables or functions that are already compiled and currently listed in its symbol table. If it fails to

find a definition there, it searches a list of directories (the *path*) for files ending in `.m` that have the same base name as the undefined identifier.[1] Once Octave finds a file with a name that matches, the contents of the file are read. If it defines a *single* function, it is compiled and executed. See Section 11.11 [Script Files], page 238, for more information about how you can define more than one function in a single file.

When Octave defines a function from a function file, it saves the full name of the file it read and the time stamp on the file. If the time stamp on the file changes, Octave may reload the file. When Octave is running interactively, time stamp checking normally happens at most once each time Octave prints the prompt. Searching for new function definitions also occurs if the current working directory changes.

Checking the time stamp allows you to edit the definition of a function while Octave is running, and automatically use the new function definition without having to restart your Octave session.

To avoid degrading performance unnecessarily by checking the time stamps on functions that are not likely to change, Octave assumes that function files in the directory tree *octave-home*`/share/octave/`*version*`/m` will not change, so it doesn't have to check their time stamps every time the functions defined in those files are used. This is normally a very good assumption and provides a significant improvement in performance for the function files that are distributed with Octave.

If you know that your own function files will not change while you are running Octave, you can improve performance by calling `ignore_function_time_stamp ("all")`, so that Octave will ignore the time stamps for all function files. Passing `"system"` to this function resets the default behavior.

`edit` *name*
`edit` *field value*
*value* = `edit (`"*get*", *field*`)`
*value* = `edit (`"*get*", "*all*"`)`

Edit the named function, or change editor settings.

If `edit` is called with the name of a file or function as its argument it will be opened in the default text editor. The default editor for the Octave GUI is specified in the Editor tab of Preferences. The default editor for the CLI is specified by the `EDITOR` function.

- If the function *name* is available in a file on your path, then it will be opened in the editor. If no file is found, then the m-file variant, ending with `".m"`, will be considered. If still no file is found, then variants with a leading `"@"` and then with both a leading `"@"` and trailing `".m"` will be considered.
- If *name* is the name of a command-line function, then an m-file will be created to contain that function along with its current definition.
- If *name*`.cc` is specified, then it will search for *name*`.cc` in the path and open it in the editor. If the file is not found, then a new `.cc` file will be created. If *name* happens to be an m-file or command-line function, then the text of that function will be inserted into the .cc file as a comment.

[1] The '`.m`' suffix was chosen for compatibility with MATLAB.

- If *name*`.ext` is on your path then it will be edited, otherwise the editor will be started with *name*`.ext` in the current directory as the filename.

  **Warning:** You may need to clear *name* before the new definition is available. If you are editing a .cc file, you will need to execute `mkoctfile` *name*`.cc` before the definition will be available.

If `edit` is called with *field* and *value* variables, the value of the control field *field* will be set to *value*.

If an output argument is requested and the first input argument is `get` then `edit` will return the value of the control field *field*. If the control field does not exist, edit will return a structure containing all fields and values. Thus, `edit ("get", "all")` returns a complete control structure.

The following control fields are used:

'`author`' This is the name to put after the `"## Author:"` field of new functions. By default it guesses from the `gecos` field of the password database.

'`email`' This is the e-mail address to list after the name in the author field. By default it guesses `<$LOGNAME@$HOSTNAME>`, and if `$HOSTNAME` is not defined it uses `uname -n`. You probably want to override this. Be sure to use the format `user@host`.

'`license`'

- '`gpl`' GNU General Public License (default).
- '`bsd`' BSD-style license without advertising clause.
- '`pd`' Public domain.
- '`"text"`' Your own default copyright and license.

Unless you specify '`pd`', edit will prepend the copyright statement with `"Copyright (C) YYYY Author"`.

'`mode`' This value determines whether the editor should be started in async mode (editor is started in the background and Octave continues) or sync mode (Octave waits until the editor exits). Set it to `"sync"` to start the editor in sync mode. The default is `"async"` (see [system], page 1058).

'`editinplace`'
Determines whether files should be edited in place, without regard to whether they are modifiable or not. The default is `true`. Set it to `false` to have read-only function files automatically copied to '`home`', if it exists, when editing them.

'`home`' This value indicates a directory that system m-files should be copied into before opening them in the editor. The intent is that this directory is also in the path, so that the edited copy of a system function file shadows the original. This setting only has an effect when '`editinplace`' is set to `false`. The default is the empty matrix (`[]`), which means it is not used. The default in previous versions of Octave was `~/octave`.

**See also:** [EDITOR], page 33, [path], page 229.

`mfilename ()`
`mfilename ("`*fullpath*`")`
`mfilename ("`*fullpathext*`")`

Return the name of the currently executing file.

The base name of the currently executing script or function is returned without any extension. If called from outside an m-file, such as the command line, return the empty string.

Given the argument `"fullpath"`, include the directory part of the filename, but not the extension.

Given the argument `"fullpathext"`, include the directory part of the filename and the extension.

**See also:** [inputname], page 208, [dbstack], page 279.

*val* `= ignore_function_time_stamp ()`
*old_val* `= ignore_function_time_stamp (`*new_val*`)`

Query or set the internal variable that controls whether Octave checks the time stamp on files each time it looks up functions defined in function files.

If the internal variable is set to `"system"`, Octave will not automatically recompile function files in subdirectories of *octave-home*`/share/`*version*`/m` if they have changed since they were last compiled, but will recompile other function files in the search path if they change.

If set to `"all"`, Octave will not recompile any function files unless their definitions are removed with `clear`.

If set to `"none"`, Octave will always check time stamps on files to determine whether functions defined in function files need to recompiled.

### 11.10.1 Manipulating the Load Path

When a function is called, Octave searches a list of directories for a file that contains the function declaration. This list of directories is known as the load path. By default the load path contains a list of directories distributed with Octave plus the current working directory. To see your current load path call the `path` function without any input or output arguments.

It is possible to add or remove directories to or from the load path using `addpath` and `rmpath`. As an example, the following code adds '`~/Octave`' to the load path.

```
addpath ("~/Octave")
```

After this the directory '`~/Octave`' will be searched for functions.

`addpath (`*dir1*`, ...)`
`addpath (`*dir1*`, ..., `*option*`)`
*oldpath* `= addpath (...)`

Add named directories to the function search path.

If *option* is `"-begin"` or 0 (the default), prepend the directory name(s) to the current path. If *option* is `"-end"` or 1, append the directory name(s) to the current path. Directories added to the path must exist.

In addition to accepting individual directory arguments, lists of directory names separated by `pathsep` are also accepted. For example:

```
addpath ("dir1:/dir2:~/dir3")
```

The newly added paths appear in the load path in the same order that they appear in the arguments of `addpath`. When extending the load path to the front, the last path in the list of arguments is added first. When extending the load path to the end, the first path in the list of arguments is added first.

For each directory that is added, and that was not already in the path, `addpath` checks for the existence of a file named `PKG_ADD` (note lack of .m extension) and runs it if it exists.

**See also:** [path], page 229, [rmpath], page 228, [genpath], page 228, [pathdef], page 229, [savepath], page 228, [pathsep], page 229.

*pathstr* = genpath (*dir*)
*pathstr* = genpath (*dir*, *skipdir1*, ...)

Return a path constructed from *dir* and all its subdirectories.

The path does not include package directories (beginning with '+'), old-style class directories (beginning with '@'), `private` directories, or any subdirectories of these types.

If additional string parameters are given, the resulting path will exclude directories with those names.

**See also:** [path], page 229, [addpath], page 227.

rmpath (*dir1*, ...)
*oldpath* = rmpath (*dir1*, ...)

Remove *dir1*, ... from the current function search path.

In addition to accepting individual directory arguments, lists of directory names separated by `pathsep` are also accepted. For example:

```
rmpath ("dir1:/dir2:~/dir3")
```

For each directory that is removed, `rmpath` checks for the existence of a file named `PKG_DEL` (note lack of .m extension) and runs it if it exists.

**See also:** [path], page 229, [addpath], page 227, [genpath], page 228, [pathdef], page 229, [savepath], page 228, [pathsep], page 229.

savepath
savepath *file*
*status* = savepath (...)

Save the unique portion of the current function search path to *file*.

The list of folders that are saved in *file* does *not* include the folders that are added for Octave's own functions, those that belong to Octave packages (see [pkg load], page 1084), and those added via command line switches.

If *file* is omitted, Octave looks in the current directory for a project-specific `.octaverc` file in which to save the path information. If no such file is present then the user's configuration file `~/.octaverc` is used.

If successful, `savepath` returns 0.

The `savepath` function makes it simple to customize a user's configuration file to restore the working paths necessary for a particular instance of Octave. Assuming no filename is specified, Octave will automatically restore the saved directory paths from the appropriate `.octaverc` file when starting up. If a filename has been specified then the paths may be restored manually by calling `source` *`file`*.

**See also:** [path], page 229, [addpath], page 227, [rmpath], page 228, [genpath], page 228, [pathdef], page 229.

`path ()`
*`str`* `= path ()`
*`str`* `= path (`*`path1`*`, ...)`

Modify or display Octave's load path.

If *nargin* and *nargout* are zero, display the elements of Octave's load path in an easy to read format.

If *nargin* is zero and nargout is greater than zero, return the current load path.

If *nargin* is greater than zero, concatenate the arguments, separating them with `pathsep`. Set the internal search path to the result and return it.

No checks are made for duplicate elements.

**See also:** [addpath], page 227, [rmpath], page 228, [genpath], page 228, [pathdef], page 229, [savepath], page 228, [pathsep], page 229.

*`val`* `= pathdef ()`

Return the default path for Octave.

The path information is extracted from one of four sources. The possible sources, in order of preference, are:

1. `.octaverc`
2. `~/.octaverc`
3. `<OCTAVE_HOME>/.../<version>/m/startup/octaverc`
4. Octave's path prior to changes by any octaverc file.

**See also:** [path], page 229, [addpath], page 227, [rmpath], page 228, [genpath], page 228, [savepath], page 228.

*`val`* `= pathsep ()`

Query the character used to separate directories in a path.

**See also:** [filesep], page 1046.

`rehash ()`

Reinitialize Octave's load path directory cache.

*`fname`* `= file_in_loadpath (`*`file`*`)`
*`fname`* `= file_in_loadpath (`*`file`*`, "`*`all`*`")`

Return the absolute name of *file* if it can be found in the list of directories specified by `path`.

If no file is found, return an empty character string.

When *file* is already an absolute name, the name is checked against the file system instead of Octave's loadpath. In this case, if *file* exists it will be returned in *fname*, otherwise an empty string is returned.

If the first argument is a cell array of strings, search each directory of the loadpath for element of the cell array and return the first that matches.

If the second optional argument `"all"` is supplied, return a cell array containing the list of all files that have the same name in the path. If no files are found, return an empty cell array.

**See also:** [file_in_path], page 1045, [dir_in_loadpath], page 230, [path], page 229.

*pathstr* = `restoredefaultpath` ()

Restore Octave's path to its initial state at startup.

The re-initialized path is returned as an output.

**See also:** [path], page 229, [addpath], page 227, [rmpath], page 228, [genpath], page 228, [pathdef], page 229, [savepath], page 228, [pathsep], page 229.

*pathstr* = `command_line_path` ()

Return the path argument given to Octave at the command line when the interpreter was started (`--path` *arg*).

**See also:** [path], page 229, [addpath], page 227, [rmpath], page 228, [genpath], page 228, [pathdef], page 229, [savepath], page 228, [pathsep], page 229.

*dirname* = `dir_in_loadpath` (*dir*)
*dirname* = `dir_in_loadpath` (*dir*, "*all*")

Return the absolute name of the loadpath element matching *dir* if it can be found in the list of directories specified by `path`.

If no match is found, return an empty character string.

The match is performed at the end of each path element. For example, if *dir* is `"foo/bar"`, it matches the path element `"/some/dir/foo/bar"`, but not `"/some/dir/foo/bar/baz"` `"/some/dir/allfoo/bar"`. When *dir* is an absolute name, rather than just a path fragment, it is matched against the file system instead of Octave's loadpath. In this case, if *dir* exists it will be returned in *dirname*, otherwise an empty string is returned.

If the optional second argument is supplied, return a cell array containing all name matches rather than just the first.

**See also:** [file_in_path], page 1045, [file_in_loadpath], page 229, [path], page 229.

*current_encoding* = `mfile_encoding` ()
`mfile_encoding` (*new_encoding*)
*old_encoding* = `mfile_encoding` (*new_encoding*)

Query or set the encoding that is used for reading m-files.

The input and output are strings naming an encoding, e.g., `"utf-8"`.

This encoding is used by Octave's parser when reading m-files unless a different encoding was set for a specific directory containing m-files using the function `dir_encoding` or in a file `.oct-config` in that directory.

The special value `"system"` selects the encoding that matches the system locale.

If the m-file encoding is changed after the m-files have already been parsed, the files have to be parsed again for that change to take effect. That can be triggered with the command `clear all`.

Additionally, this encoding is used to load and save files with the built-in editor in Octave's GUI.

**See also:** [dir_encoding], page 231.

*current_encoding* = dir_encoding (*dir*)
dir_encoding (*dir*, *new_encoding*)
dir_encoding (*dir*, "*delete*")
*old_encoding* = dir_encoding (*dir*, *new_encoding*)

Query or set the *encoding* that is used for reading m-files in *dir*.

The per-directory encoding overrides the (globally set) m-file encoding, see [mfile_encoding], page 230.

The string *DIR* must match how the directory would appear in the load path.

The *new_encoding* input must be a valid encoding identifier or `"delete"`. In the latter case, any per-directory encoding is removed and the (globally set) m-file encoding will be used for the given *dir*.

The currently or previously used encoding is returned only if an output argument is requested.

The directory encoding is automatically read from the file `.oct-config` when a new path is added to the load path (for example with `addpath`). To set the encoding for all files in the same folder, that file must contain a line starting with `"encoding="` followed by the encoding identifier.

For example to set the file encoding for all files in the same folder to ISO 8859-1 (Latin-1), create a file `.oct-config` with the following content:

```
encoding=iso8859-1
```

If the file encoding is changed after the files have already been parsed, the files have to be parsed again for that change to take effect. That can be done with the command `clear all`.

**See also:** [addpath], page 227, [path], page 229, [mfile_encoding], page 230.

### 11.10.2 Subfunctions

A function file may contain secondary functions called *subfunctions*. These secondary functions are only visible to the other functions in the same function file. For example, a file `f.m` containing

```
function f ()
  printf ("in f, calling g\n");
  g ()
endfunction
function g ()
  printf ("in g, calling h\n");
  h ()
endfunction
function h ()
  printf ("in h\n")
endfunction
```

defines a main function `f` and two subfunctions. The subfunctions `g` and `h` may only be called from the main function `f` or from the other subfunctions, but not from outside the file `f.m`.

*`subfcn_list`* `= localfunctions ()`

Return a list of all local functions, i.e., subfunctions, within the current file.

The return value is a column cell array of function handles to all local functions accessible from the function from which `localfunctions` is called. Nested functions are *not* included in the list.

If the call is from the command line, an anonymous function, or a script, the return value is an empty cell array.

**See also:** [functions], page 248.

### 11.10.3 Private Functions

In many cases one function needs to access one or more helper functions. If the helper function is limited to the scope of a single function, then subfunctions as discussed above might be used. However, if a single helper function is used by more than one function, then this is no longer possible. In this case the helper functions might be placed in a subdirectory, called "private", of the directory in which the functions needing access to this helper function are found.

As a simple example, consider a function `func1`, that calls a helper function `func2` to do much of the work. For example:

```
function y = func1 (x)
  y = func2 (x);
endfunction
```

Then if the path to `func1` is `<directory>/func1.m`, and if `func2` is found in the directory `<directory>/private/func2.m`, then `func2` is only available for use of the functions, like `func1`, that are found in `<directory>`.

### 11.10.4 Nested Functions

Nested functions are similar to subfunctions in that only the main function is visible outside the file. However, they also allow for child functions to access the local variables in their parent function. This shared access mimics using a global variable to share information — but a global variable which is not visible to the rest of Octave. As a programming strategy,

sharing data this way can create code which is difficult to maintain. It is recommended to use subfunctions in place of nested functions when possible.

As a simple example, consider a parent function `foo`, that calls a nested child function `bar`, with a shared variable *x*.

```
function y = foo ()
  x = 10;
  bar ();
  y = x;

  function bar ()
    x = 20;
  endfunction
endfunction

foo ()
 ⇒ 20
```

Notice that there is no special syntax for sharing *x*. This can lead to problems with accidental variable sharing between a parent function and its child. While normally variables are inherited, child function parameters and return values are local to the child function.

Now consider the function `foobar` that uses variables *x* and *y*. `foobar` calls a nested function `foo` which takes *x* as a parameter and returns *y*. `foo` then calls `bat` which does some computation.

```
function z = foobar ()
  x = 0;
  y = 0;
  z = foo (5);
  z += x + y;

  function y = foo (x)
    y = x + bat ();

    function z = bat ()
      z = x;
    endfunction
  endfunction
endfunction

foobar ()
    ⇒ 10
```

It is important to note that the *x* and *y* in `foobar` remain zero, as in `foo` they are a return value and parameter respectively. The *x* in `bat` refers to the *x* in `foo`.

Variable inheritance leads to a problem for `eval` and scripts. If a new variable is created in a parent function, it is not clear what should happen in nested child functions. For example, consider a parent function `foo` with a nested child function `bar`:

```
function y = foo (to_eval)
  bar ();
  eval (to_eval);

  function bar ()
    eval ("x = 100;");
    eval ("y = x;");
  endfunction
endfunction

foo ("x = 5;")
    ⇒ error: can not add variable "x" to a static workspace

foo ("y = 10;")
    ⇒ 10

foo ("")
    ⇒ 100
```

The parent function `foo` is unable to create a new variable $x$, but the child function `bar` was successful. Furthermore, even in an `eval` statement $y$ in `bar` is the same $y$ as in its parent function `foo`. The use of `eval` in conjunction with nested functions is best avoided.

As with subfunctions, only the first nested function in a file may be called from the outside. Inside a function the rules are more complicated. In general a nested function may call:

0. Globally visible functions
1. Any function that the nested function's parent can call
2. Sibling functions (functions that have the same parents)
3. Direct children

As a complex example consider a parent function `ex_top` with two child functions, `ex_a` and `ex_b`. In addition, `ex_a` has two more child functions, `ex_aa` and `ex_ab`. For example:

```
function ex_top ()
  ## Can call: ex_top, ex_a, and ex_b
  ## Can NOT call: ex_aa and ex_ab

  function ex_a ()
    ## Can call everything

    function ex_aa ()
      ## Can call everything
    endfunction

    function ex_ab ()
      ## Can call everything
    endfunction
  endfunction
```

```
  function ex_b ()
    ## Can call: ex_top, ex_a, and ex_b
    ## Can NOT call: ex_aa and ex_ab
  endfunction
endfunction
```

### 11.10.5 Overloading and Autoloading

Functions can be overloaded to work with different input arguments. For example, the operator '+' has been overloaded in Octave to work with single, double, uint8, int32, and many other arguments. The preferred way to overload functions is through classes and object oriented programming (see Section 34.4.1 [Function Overloading], page 983). Occasionally, however, one needs to undo user overloading and call the default function associated with a specific type. The `builtin` function exists for this purpose.

`[...] = builtin (`*f*`, ...)`

Call the base function *f* even if *f* is overloaded to another function for the given type signature.

This is normally useful when doing object-oriented programming and there is a requirement to call one of Octave's base functions rather than the overloaded one of a new class.

A trivial example which redefines the `sin` function to be the `cos` function shows how `builtin` works.

```
sin (0)
  ⇒ 0
function y = sin (x), y = cos (x); endfunction
sin (0)
  ⇒ 1
builtin ("sin", 0)
  ⇒ 0
```

A single dynamically linked file might define several functions. However, as Octave searches for functions based on the functions filename, Octave needs a manner in which to find each of the functions in the dynamically linked file. On operating systems that support symbolic links, it is possible to create a symbolic link to the original file for each of the functions which it contains.

However, there is at least one well known operating system that doesn't support symbolic links. Making copies of the original file for each of the functions is undesirable as it increases the amount of disk space used by Octave. Instead Octave supplies the `autoload` function, that permits the user to define in which file a certain function will be found.

*autoload_map* = `autoload ()`
`autoload (`*function*`, `*file*`)`
`autoload (...,` "*remove*"`)`

Define *function* to autoload from *file*.

The second argument, *file*, should be an absolute filename or a file name in the same directory as the function or script from which the autoload command was run. *file should not* depend on the Octave load path.

Normally, calls to `autoload` appear in PKG_ADD script files that are evaluated when a directory is added to Octave's load path. To avoid having to hardcode directory names in *file*, if *file* is in the same directory as the PKG_ADD script then

```
autoload ("foo", "bar.oct");
```

will load the function `foo` from the file `bar.oct`. The above usage when `bar.oct` is not in the same directory, or usages such as

```
autoload ("foo", file_in_loadpath ("bar.oct"))
```

are strongly discouraged, as their behavior may be unpredictable.

With no arguments, return a structure containing the current autoload map.

If a third argument `"remove"` is given, the function is cleared and not loaded anymore during the current Octave session.

**See also:** [PKG_ADD], page 1091.

### 11.10.6 Function Locking

It is sometime desirable to lock a function into memory with the `mlock` function. This is typically used for dynamically linked functions in oct-files or mex-files that contain some initialization, and it is desirable that calling `clear` does not remove this initialization.

As an example,

```
function my_function ()
  mlock ();
  ...
endfunction
```

prevents `my_function` from being removed from memory after it is called, even if `clear` is called. It is possible to determine if a function is locked into memory with the `mislocked`, and to unlock a function with `munlock`, which the following code illustrates.

```
my_function ();
mislocked ("my_function")
⇒ ans = 1
munlock ("my_function");
mislocked ("my_function")
⇒ ans = 0
```

A common use of `mlock` is to prevent persistent variables from being removed from memory, as the following example shows:

```
function count_calls ()
  mlock ();
  persistent calls = 0;
  printf ("count_calls() has been called %d times\n", ++calls);
endfunction

count_calls ();
⊣ count_calls() has been called 1 times

clear count_calls
count_calls ();
⊣ count_calls() has been called 2 times
```

`mlock` might also be used to prevent changes to an m-file, such as in an external editor, from having any effect in the current Octave session; A similar effect can be had with the `ignore_function_time_stamp` function.

`mlock ()`
: Lock the current function into memory so that it can't be removed with `clear`.

  **See also:** [munlock], page 237, [mislocked], page 237, [persistent], page 1193, [clear], page 154.

`munlock ()`
`munlock (`*`fcn`*`)`
: Unlock the named function *fcn* so that it may be removed from memory with `clear`.

  If no function is named then unlock the current function.

  **See also:** [mlock], page 237, [mislocked], page 237, [persistent], page 1193, [clear], page 154.

*`tf`* `= mislocked ()`
*`tf`* `= mislocked (`*`fcn`*`)`
: Return true if the named function *fcn* is locked in memory.

  If no function is named then return true if the current function is locked.

  **See also:** [mlock], page 237, [munlock], page 237, [persistent], page 1193, [clear], page 154.

### 11.10.7 Function Precedence

Given the numerous different ways that Octave can define a function, it is possible and even likely that multiple versions of a function, might be defined within a particular scope. The precedence of which function will be used within a particular scope is given by

1. Subfunction A subfunction with the required function name in the given scope.
2. Private function A function defined within a private directory of the directory which contains the current function.
3. Class constructor A function that constructs a user class as defined in chapter Chapter 34 [Object Oriented Programming], page 973.

4. Class method An overloaded function of a class as in chapter Chapter 34 [Object Oriented Programming], page 973.
5. Command-line Function A function that has been defined on the command-line.
6. Autoload function A function that is marked as autoloaded with [autoload], page 235.
7. A Function on the Path A function that can be found on the users load-path. There can also be Oct-file, mex-file or m-file versions of this function and the precedence between these versions are in that order.
8. Built-in function A function that is a part of core Octave such as `numel`, `size`, etc.

## 11.11 Script Files

A script file is a file containing (almost) any sequence of Octave commands. It is read and evaluated just as if you had typed each command at the Octave prompt, and provides a convenient way to perform a sequence of commands that do not logically belong inside a function.

Unlike a function file, a script file must *not* begin with the keyword `function`. If it does, Octave will assume that it is a function file, and that it defines a single function that should be evaluated as soon as it is defined.

A script file also differs from a function file in that the variables named in a script file are not local variables, but are in the same scope as the other variables that are visible on the command line.

Even though a script file may not begin with the `function` keyword, it is possible to define more than one function in a single script file and load (but not execute) all of them at once. To do this, the first token in the file (ignoring comments and other white space) must be something other than `function`. If you have no other statements to evaluate, you can use a statement that has no effect, like this:

```
# Prevent Octave from thinking that this
# is a function file:

1;

# Define function one:

function one ()
  ...
```

To have Octave read and compile these functions into an internal form, you need to make sure that the file is in Octave's load path (accessible through the `path` function), then simply type the base name of the file that contains the commands. (Octave uses the same rules to search for script files as it does to search for function files.)

If the first token in a file (ignoring comments) is `function`, Octave will compile the function and try to execute it, printing a message warning about any non-whitespace characters that appear after the function definition.

Note that Octave does not try to look up the definition of any identifier until it needs to evaluate it. This means that Octave will compile the following statements if they appear in a script file, or are typed at the command line,

```
# not a function file:
1;
function foo ()
  do_something ();
endfunction
function do_something ()
  do_something_else ();
endfunction
```

even though the function `do_something` is not defined before it is referenced in the function `foo`. This is not an error because Octave does not need to resolve all symbols that are referenced by a function until the function is actually evaluated.

Since Octave doesn't look for definitions until they are needed, the following code will always print '`bar = 3`' whether it is typed directly on the command line, read from a script file, or is part of a function body, even if there is a function or script file called `bar.m` in Octave's path.

```
eval ("bar = 3");
bar
```

Code like this appearing within a function body could fool Octave if definitions were resolved as the function was being compiled. It would be virtually impossible to make Octave clever enough to evaluate this code in a consistent fashion. The parser would have to be able to perform the call to `eval` at compile time, and that would be impossible unless all the references in the string to be evaluated could also be resolved, and requiring that would be too restrictive (the string might come from user input, or depend on things that are not known until the function is evaluated).

Although Octave normally executes commands from script files that have the name *`file`*`.m`, you can use the function `source` to execute commands from any file.

`source (`*`file`*`)`
`source (`*`file`*`, `*`context`*`)`

> Parse and execute the contents of *file*.
>
> Without specifying *context*, this is equivalent to executing commands from a script file, but without requiring the file to be named *`file`*`.m` or to be on the execution path.
>
> Instead of the current context, the script may be executed in either the context of the function that called the present function (`"caller"`), or the top-level context (`"base"`).
>
> **See also:** [run], page 189.

### 11.11.1 Publish Octave Script Files

The function `publish` provides a dynamic possibility to document your script file. Unlike static documentation, `publish` runs the script file, saves any figures and output while running the script, and presents them alongside static documentation in a desired output format. The static documentation can make use of Section 11.11.2 [Publishing Markup], page 242, to enhance and customize the output.

`publish (`*`file`*`)`
`publish (`*`file`*`, `*`output_format`*`)`

`publish (`*`file`*`, `*`option1`*`, `*`value1`*`, ...)`
`publish (`*`file`*`, `*`options`*`)`
*`output_file`* `= publish (`*`file`*`, ...)`

Generate a report from the Octave script file *file* in one of several output formats.

The generated reports interpret Publishing Markup in section comments, which is explained in detail in the GNU Octave manual. Section comments are comment blocks that start with a line with double comment character.

Assume the following example, using some Publishing Markup, to be the contents of the script file `pub_example.m`:

```
## Headline title
#
# Some *bold*, _italic_, or |monospaced| Text with
# a <https://www.octave.org link to *GNU Octave*>.
##

# "Real" Octave commands to be evaluated
sombrero ()

%% MATLAB comment style ('%') is supported as well
%
% * Bulleted list item 1
% * Bulleted list item 2
%
% # Numbered list item 1
% # Numbered list item 2
```

To publish this script file, type `publish ("pub_example.m")`.

When called with one input argument, a HTML report is generated in a subdirectory `html` relative to the current working directory. Any Octave commands in `pub_example.m` are evaluated in a separate context and any figures created while executing the script file are included in the report.

Using `publish (`*`file`*`, `*`output_format`*`)` is equivalent to the function call using a structure

```
options.format = output_format;
publish (file, options)
```

which is described below. The same holds for using option/value pairs

```
options.option1 = value1;
publish (file, options)
```

The structure *options* can have the following field names. If a field name is not specified, the default value is used:

- '`format`' — Output format of the published script file, one of

  '`html`' (default), '`doc`', '`latex`', '`ppt`', '`pdf`', or '`xml`'.

  The output formats '`doc`', '`ppt`', and '`xml`' are not currently supported. To generate a '`doc`' report, open a generated '`html`' report with your office suite.

In Octave custom formats are supported by implementing all callback subfunctions in a function file named '`__publish_<custom format>_output__.m`'. To obtain a template for the HTML format type:

```
edit (fullfile (fileparts (which ("publish")), ...
      "private", "__publish_html_output__.m"))
```

- '`outputDir`' — Full path of the directory where the generated report will be located. If no directory is given, the report is generated in a subdirectory `html` relative to the current working directory.
- '`stylesheet`' — Not supported, only for MATLAB compatibility.
- '`createThumbnail`' — Not supported, only for MATLAB compatibility.
- '`figureSnapMethod`' — Not supported, only for MATLAB compatibility.
- '`imageFormat`' — Desired format for any images produced while evaluating the code. The allowed image formats depend on the output format:
  - '`html`', '`xml`' — '`png`' (default), any image format supported by Octave
  - '`latex`' — '`epsc2`' (default), any image format supported by Octave
  - '`pdf`' — '`jpg`' (default) or '`bmp`', note MATLAB uses '`bmp`' as default
  - '`doc`' or '`ppt`' — '`png`' (default), '`jpg`', '`bmp`', or '`tiff`'
- '`maxWidth`' and '`maxHeight`' — Maximum width (height) of the produced images in pixels. An empty value means no restriction. Both values must be set in order for the option to work properly.

  '`[]`' (default), integer value $\geq 0$
- '`useNewFigure`' — Use a new figure window for figures created by the evaluated code. This avoids side effects with already opened figure windows.

  '`true`' (default) or '`false`'
- '`evalCode`' — Evaluate code of the Octave source file

  '`true`' (default) or '`false`'
- '`catchError`' — Catch errors while evaluating code and continue

  '`true`' (default) or '`false`'
- '`codeToEvaluate`' — Octave commands that should be evaluated prior to publishing the script file. These Octave commands do not appear in the generated report.
- '`maxOutputLines`' — Maximum number of output lines from code evaluation which are included in output.

  '`Inf`' (default) or integer value > 0
- '`showCode`' — Show the evaluated Octave commands in the generated report

  '`true`' (default) or '`false`'

The option output *output_file* is a string with path and file name of the generated report.

**See also:** [grabcode], page 241.

The counterpart to `publish` is `grabcode`:

`grabcode` *`filename`*
`grabcode` *`url`*
*`code_str`* `= grabcode (...)`

Grab the code from a report created by the `publish` function.

The grabbed code inside the published report must be enclosed by the strings '`##### SOURCE BEGIN #####`' and '`##### SOURCE END #####`'. The `publish` function creates this format automatically.

If no return value is requested the code is saved to a temporary file and opened in the default editor. NOTE: The temporary file must be saved to a new filename or the code will be lost.

If an output is requested the grabbed code will be returned as string *code_str*.

Example:

```
publish ("my_script.m");
grabcode ("html/my_script.html");
```

The example above publishes `my_script.m` to the default location `html/my_script.html`. Next, the published Octave script is grabbed to edit its content in a new temporary file.

**See also:** [publish], page 239.

## 11.11.2 Publishing Markup

### 11.11.2.1 Using Publishing Markup in Script Files

To use Publishing Markup, start by typing '`##`' or '`%%`' at the beginning of a new line. For MATLAB compatibility '`%%%`' is treated the same way as '`%%`'.

The lines following '`##`' or '`%%`' start with one of either '`#`' or '`%`' followed by at least one space. These lines are interpreted as section. A section ends at the first line not starting with '`#`' or '`%`', or when the end of the document is reached.

A section starting in the first line of the document, followed by another start of a section that might be empty, is interpreted as a document title and introduction text.

See the example below for clarity:

```
%% Headline title
%
% Some *bold*, _italic_, or |monospaced| Text with
% a <https://www.octave.org link to GNU Octave>.
%%

# "Real" Octave commands to be evaluated
sombrero ()

## Octave comment style supported as well
#
# * Bulleted list item 1
# * Bulleted list item 2
#
# # Numbered list item 1
# # Numbered list item 2
```

#### 11.11.2.2 Text Formatting

Basic text formatting is supported inside sections, see the example given below:

```
##
# *bold*, _italic_, or |monospaced| Text
```

Additionally two trademark symbols are supported, just embrace the letters 'TM' or 'R'.

```
##
# (TM) or (R)
```

#### 11.11.2.3 Sections

A section is started by typing '##' or '%%' at the beginning of a new line. A section title can be provided by writing it, separated by a space, in the first line after '##' or '%%'. Without a section title, the section is interpreted as a continuation of the previous section. For MATLAB compatibility '%%%' is treated the same way as '%%'.

```
some_code ();

## Section 1
#
## Section 2

some_code ();

##
# Still in section 2

some_code ();

%%% Section 3
%
%
```

#### 11.11.2.4 Preformatted Code

To write preformatted code inside a section, indent the code by three spaces after '#' at the beginning of each line and leave the lines above and below the code blank, except for '#' at the beginning of those lines.

```
##
# This is a syntax highlighted for-loop:
#
#   for i = 1:5
#     disp (i);
#   endfor
#
# And more usual text.
```

#### 11.11.2.5 Preformatted Text

To write preformatted text inside a section, indent the code by two spaces after '#' at the beginning of each line and leave the lines above and below the preformatted text blank, except for '#' at the beginning of those lines.

```
##
# This following text is preformatted:
#
#  "To be, or not to be: that is the question:
#  Whether 'tis nobler in the mind to suffer
#  The slings and arrows of outrageous fortune,
#  Or to take arms against a sea of troubles,
#  And by opposing end them?  To die: to sleep;"
#
#  --"Hamlet" by W. Shakespeare
```

#### 11.11.2.6 Bulleted Lists

To create a bulleted list, type

```
##
#
# * Bulleted list item 1
# * Bulleted list item 2
#
```

to get output like

- Bulleted list item 1
- Bulleted list item 2

Notice the blank lines, except for the '#' or '%' before and after the bulleted list!

#### 11.11.2.7 Numbered Lists

To create a numbered list, type

```
##
#
# # Numbered list item 1
# # Numbered list item 2
#
```

to get output like

1. Numbered list item 1
2. Numbered list item 2

Notice the blank lines, except for the '`#`' or '`%`' before and after the numbered list!

### 11.11.2.8 Including File Content

To include the content of an external file, e.g., a file called '`my_function.m`' at the same location as the published Octave script, use the following syntax to include it with Octave syntax highlighting.

Alternatively, you can write the full or relative path to the file.

```
##
#
# <include>my_function.m</include>
#
# <include>/full/path/to/my_function.m</include>
#
# <include>../relative/path/to/my_function.m</include>
#
```

### 11.11.2.9 Including Graphics

To include external graphics, e.g., a graphic called '`my_graphic.png`' at the same location as the published Octave script, use the following syntax.

Alternatively, you can write the full path to the graphic.

```
##
#
# <<my_graphic.png>>
#
# <</full/path/to/my_graphic.png>>
#
# <<../relative/path/to/my_graphic.png>>
#
```

### 11.11.2.10 Including URLs

Basically, a URL is written between an opening '`<`' and a closing '`>`' angle.

```
##
# <https://www.octave.org>
```

Text that is within these angles and separated by at least one space from the URL is a displayed text for the link.

```
##
# <https://www.octave.org GNU Octave>
```

A link starting with '<octave:' followed by the name of a GNU Octave function, optionally with a displayed text, results in a link to the online GNU Octave documentations function index.

```
##
# <octave:DISP The display function>
```

#### 11.11.2.11 Mathematical Equations

One can insert LaTeX inline math, surrounded by single '$' signs, or displayed math, surrounded by double '$$' signs, directly inside sections.

```
##
# Some shorter inline equation $e^{ix} = \cos x + i\sin x$.
#
# Or more complicated formulas as displayed math:
# $$e^x = \lim_{n\rightarrow\infty}\left(1+\dfrac{x}{n}\right)^{n}.$$
```

#### 11.11.2.12 HTML Markup

If the published output is a HTML report, you can insert HTML markup, that is only visible in this kind of output.

```
##
# <html>
# <table style="border:1px solid black;">
# <tr><td>1</td><td>2</td></tr>
# <tr><td>3</td><td>3</td></tr>
# </html>
```

#### 11.11.2.13 LaTeX Markup

If the published output is a LaTeX or PDF report, you can insert LaTeX markup, that is only visible in this kind of output.

```
##
# <latex>
# Some output only visible in LaTeX or PDF reports.
# \begin{equation}
# e^x = \lim\limits_{n\rightarrow\infty}\left(1+\dfrac{x}{n}\right)^{n}
# \end{equation}
# </latex>
```

### 11.11.3 Jupyter Notebooks

Jupyter notebooks are one popular technique for displaying code, text, and graphical output together in a comprehensive manner. Octave can publish results to a Jupyter notebook with the function `jupyter_notebook`.

*notebook* = jupyter_notebook (*notebook_filename*)
*notebook* = jupyter_notebook (*notebook_filename*, *options*)

Run and fill the Jupyter Notebook in file *notebook_filename* from within GNU Octave.

Both text and graphical Octave outputs are supported.

This class has a public property `notebook` which is a structure representing the JSON-decoded Jupyter Notebook. This property is intentionally public to enable advanced notebook manipulations.

Note: Jupyter Notebook versions (`nbformat`) lower than 4.0 are not supported.

The optional second argument *options* is a struct with fields:

- `tmpdir` to set the temporary working directory.

`%plot` magic is supported with the following settings:

- `"%plot -f <format>"` or `"%plot --format <format>"`: specifies the image storage format. Supported formats are:
  - PNG (default)
  - SVG (Note: If SVG images do not appear in the notebook, it is most likely related to Jupyter Notebook security mechanisms and explicitly "trusting" them will be necessary).
  - JPG
- `"%plot -r <number>"` or `"%plot --resolution <number>"`: specifies the image resolution.
- `"%plot -w <number>"` or `"%plot --width <number>"`: specifies the image width.
- `"%plot -h <number>"` or `"%plot --height <number>"`: specifies the image height.

Examples:

```
## Run all cells and generate the filled notebook

## Instantiate an object from a notebook file
notebook = jupyter_notebook ("myNotebook.ipynb");
## Run the code and embed the results in the notebook property
notebook.run_all ();
## Generate a new notebook by overwriting the original notebook
notebook.generate_notebook ("myNotebook.ipynb");

## Run just the second cell and generate the filled notebook

## Instantiate an object from a notebook file
notebook = jupyter_notebook ("myNotebook.ipynb");
## Run the code and embed the results in the notebook property
notebook.run (2)
## Generate a new notebook in a new file
notebook.generate_notebook ("myNewNotebook.ipynb");

## Generate an Octave script from a notebook

## Instantiate an object from a notebook file
notebook = jupyter_notebook ("myNotebook.ipynb");
## Generate the Octave script
notebook.generate_octave_script ("jup_script.m");
```

**See also:** [jsondecode], page 111, [jsonencode], page 109.

## 11.12 Function Handles and Anonymous Functions

It can be very convenient store a function in a variable so that it can be passed to a different function. For example, a function that performs numerical minimization needs access to the function that should be minimized.

### 11.12.1 Function Handles

A function handle is a pointer to another function and is defined with the syntax

```
@function-name
```

For example,

```
f = @sin;
```

creates a function handle called `f` that refers to the function `sin`.

Function handles are used to call other functions indirectly, or to pass a function as an argument to another function like `quad` or `fsolve`. For example:

```
f = @sin;
quad (f, 0, pi)
    ⇒ 2
```

You may use `feval` to call a function using function handle, or simply write the name of the function handle followed by an argument list. If there are no arguments, you must use an empty argument list '`()`'. For example:

```
f = @sin;
feval (f, pi/4)
    ⇒ 0.70711
f (pi/4)
    ⇒ 0.70711
```

*tf* = `is_function_handle` (*x*)

Return true if *x* is a function handle.

**See also:** [isa], page 41, [typeinfo], page 41, [class], page 41, [functions], page 248.

*s* = `functions` (*fcn_handle*)

Return a structure containing information about the function handle *fcn_handle*.

The structure *s* always contains these three fields:

function
: The function name. For an anonymous function (no name) this will be the actual function definition.

type
: Type of the function.

  anonymous
  : The function is anonymous.

  private
  : The function is private.

  overloaded
  : The function overloads an existing function.

simple The function is a built-in or m-file function.

subfunction
The function is a subfunction within an m-file.

nested The function is nested.

file The m-file that will be called to perform the function. This field is empty for anonymous and built-in functions.

In addition, some function types may return more information in additional fields.

**Warning:** `functions` is provided for debugging purposes only. Its behavior may change in the future and programs should not depend on any particular output format.

**See also:** [func2str], page 249, [str2func], page 249.

*str* = `func2str` (*fcn_handle*)

Return a string containing the name of the function referenced by the function handle *fcn_handle*.

**See also:** [str2func], page 249, [functions], page 248.

*hfcn* = `str2func` (*str*)

Return a function handle constructed from the string *str*.

The input may be the name of a function such as `"sin"` or a string defining a function such as `"@(x) sin (x + pi)"`.

Programming Note: In most cases it will be better to use anonymous function syntax and let the Octave parser create the function handle rather than use `str2func`. For example:

```
hfcn = @sin ;
hfcn = @(x) sin (x + pi) ;
```

**See also:** [func2str], page 249, [functions], page 248.

*vars* = `symvar` (*str*)

Identify the symbolic variable names in the string *str*.

Common constant names such as `i`, `j`, `pi`, `Inf` and Octave functions such as `sin` or `plot` are ignored.

Any names identified are returned in a cell array of strings. The array is empty if no variables were found.

Example:

```
symvar ("x^2 + y^2 == 4")
⇒ {
    [1,1] = x
    [2,1] = y
  }
```

### 11.12.2 Anonymous Functions

Anonymous functions are defined using the syntax

```
@(argument-list) expression
```

Any variables that are not found in the argument list are inherited from the enclosing scope. Anonymous functions are useful for creating simple unnamed functions from expressions or for wrapping calls to other functions to adapt them for use by functions like `quad`. For example,

```
f = @(x) x.^2;
quad (f, 0, 10)
    ⇒ 333.33
```

creates a simple unnamed function from the expression `x.^2` and passes it to `quad`,

```
quad (@(x) sin (x), 0, pi)
    ⇒ 2
```

wraps another function, and

```
a = 1;
b = 2;
quad (@(x) betainc (x, a, b), 0, 0.4)
    ⇒ 0.13867
```

adapts a function with several parameters to the form required by `quad`. In this example, the values of $a$ and $b$ that are passed to `betainc` are inherited from the current environment.

Note that for performance reasons it is better to use handles to existing Octave functions, rather than to define anonymous functions which wrap an existing function. The integration of `sin (x)` is 5X faster if the code is written as

```
quad (@sin, 0, pi)
```

rather than using the anonymous function `@(x) sin (x)`. There are many operators which have functional equivalents that may be better choices than an anonymous function. Instead of writing

```
f = @(x, y) x + y
```

this should be coded as

```
f = @plus
```

See Section 34.4.2 [Operator Overloading], page 984, for a list of operators which also have a functional form.

## 11.13 Command Syntax and Function Syntax

In addition to the function syntax described above (i.e., calling a function like `fun (arg1, arg2, ...)`), a function can be called using command syntax (for example, calling a function like `fun arg1 arg2 ...`). In that case, all arguments are passed to the function as strings. For example,

```
my_command hello world
```

is equivalent to

```
my_command ("hello", "world")
```

The general form of a command call is

```
cmdname arg1 arg2 ...
```

which translates directly to

```
cmdname ("arg1", "arg2", ...)
```

If an argument including spaces should be passed to a function in command syntax, (double-)quotes can be used. For example,

```
my_command "first argument" "second argument"
```

is equivalent to

```
my_command ("first argument", "second argument")
```

Any function can be used as a command if it accepts string input arguments. For example:

```
upper lower_case_arg
   ⇒ ans = LOWER_CASE_ARG
```

Since the arguments are passed as strings to the corresponding function, it is not possible to pass input arguments that are stored in variables. In that case, a command must be called using the function syntax. For example:

```
strvar = "hello world";
upper strvar
   ⇒ ans = STRVAR
upper (strvar)
   ⇒ ans = HELLO WORLD
```

Additionally, the return values of functions cannot be assigned to variables using the command syntax. Only the first return argument is assigned to the built-in variable `ans`. If the output argument of a command should be assigned to a variable, or multiple output arguments of a function should be returned, the function syntax must be used.

It should be noted that mixing command syntax and binary operators can create apparent ambiguities with mathematical and logical expressions that use function syntax. For example, all three of the statements

```
arg1 - arg2
arg1 -arg2
arg1-arg2
```

could be intended by a user to be subtraction operations between `arg1` and `arg2`. The first two, however, could also have been meant as a command syntax call to function `arg1`, in the first case with options `-` and `arg2`, and in the second case with option `-arg2`.

Octave uses whitespace to interpret such expressions according to the following rules:

- Statements consisting of plain symbols without any operators that are separated only by whitespace are always treated as command syntax:

  ```
  arg1 arg2 arg3 ... argn
  ```

- Statements without any whitespace are always treated as function syntax:

  ```
  arg1+arg2
  arg1&&arg2||arg3
  arg1+=arg2*arg3
  ```

- If the first symbol is a constant (or special-valued named constant pi, i, I, j, J, e, NaN, or Inf) followed by a binary operator, the statement is treated as function syntax regardless of any whitespace or what follows the second symbol:

```
7 -arg2
pi+ arg2
j * arg2 -arg3
```

- If the first symbol is a function or variable and there is no whitespace separating the operator and the second symbol, the statement is treated as command syntax:

```
arg1 -arg2
arg1 &&arg2 ||arg3
arg1 +=arg2*arg3
```

- Any other whitespace combination will result in the statement being treated as function syntax.

Note 1: If a special-valued named constant has been redefined as a variable, the interpreter will still process the statement with function syntax.

Note 2: Attempting to use a variable as `arg1` in a command being processed as command syntax will result in an error.

## 11.14 Organization of Functions Distributed with Octave

Many of Octave's standard functions are distributed as function files. They are loosely organized by topic, in subdirectories of *octave-home*`/share/octave/`*version*`/m`, to make it easier to find them.

The following is a list of all the function file subdirectories, and the types of functions you will find there.

`@ftp` Class functions for the FTP object.

`+containers`
: Package for the containers classes.

`audio` Functions for playing and recording sounds.

`deprecated`
: Out-of-date functions which will eventually be removed from Octave.

`elfun` Elementary functions, principally trigonometric.

`general` Miscellaneous matrix manipulations, like `flipud`, `rot90`, and `triu`, as well as other basic functions, like `ismatrix`, `narginchk`, etc.

`geometry` Functions related to Delaunay triangulation.

`gui` Functions for GUI elements like dialog, message box, etc.

`help` Functions for Octave's built-in help system.

`image` Image processing tools. These functions require the X Window System.

`io` Input-output functions.

`java` Functions related to the Java integration.

`linear-algebra`
: Functions for linear algebra.

`miscellaneous`
: Functions that don't really belong anywhere else.

`ode`
: Functions to solve ordinary differential equations (ODEs).

`optimization`
: Functions related to minimization, optimization, and root finding.

`path`
: Functions to manage the directory path Octave uses to find functions.

`pkg`
: Package manager for installing external packages of functions in Octave.

`plot`
: Functions for displaying and printing two- and three-dimensional graphs.

`polynomial`
: Functions for manipulating polynomials.

`prefs`
: Functions implementing user-defined preferences.

`set`
: Functions for creating and manipulating sets of unique values.

`signal`
: Functions for signal processing applications.

`sparse`
: Functions for handling sparse matrices.

`specfun`
: Special functions such as `bessel` or `factor`.

`special-matrix`
: Functions that create special matrix forms such as Hilbert or Vandermonde matrices.

`startup`
: Octave's system-wide startup file.

`statistics`
: Statistical functions.

`strings`
: Miscellaneous string-handling functions.

`testfun`
: Functions for performing unit tests on other functions.

`time`
: Functions related to time and date processing.

# 12 Errors and Warnings

Octave includes several functions for printing error and warning messages. When you write functions that need to take special action when they encounter abnormal conditions, you should print the error messages using the functions described in this chapter.

Since many of Octave's functions use these functions, it is also useful to understand them, so that errors and warnings can be handled.

## 12.1 Handling Errors

An error is something that occurs when a program is in a state where it doesn't make sense to continue. An example is when a function is called with too few input arguments. In this situation the function should abort with an error message informing the user of the lacking input arguments.

Since an error can occur during the evaluation of a program, it is very convenient to be able to detect that an error occurred, so that the error can be fixed. This is possible with the `try` statement described in Section 10.9 [The try Statement], page 202.

### 12.1.1 Raising Errors

The most common use of errors is for checking input arguments to functions. The following example calls the `error` function if the function `f` is called without any input arguments.

```
function f (arg1)
  if (nargin == 0)
    error ("not enough input arguments");
  endif
endfunction
```

When the `error` function is called, it prints the given message and returns to the Octave prompt. This means that no code following a call to `error` will be executed.

It is also possible to assign an identification string to an error. If an error has such an ID the user can catch this error as will be described in the next section. To assign an ID to an error, simply call `error` with two string arguments, where the first is the identification string, and the second is the actual error. Note that error IDs are in the format `"NAMESPACE:ERROR-NAME"`. The namespace `"Octave"` is used for Octave's own errors. Any other string is available as a namespace for user's own errors.

`error` (*msg*)
`error` (*template*, ...)
`error` (*id*, *template*, ...)
`error` (*errstruct*)

Display an error message and stop m-file execution.

The input *msg* is a simple string to which the text '`error:` ' is prepended. The resulting message is printed on the `stderr` stream. Alternatively, the first input may be a template string *template* which uses the same rules as the `printf` family of functions (see Section 14.2.4 [Formatted Output], page 315). Formatting is only done for single-quoted character vectors if there are additional arguments following the template string. If there are no additional arguments, the template string is used

literally (i.e., without interpreting any escape sequences in single-quoted character vectors).

The optional *id* argument allows programmers to tag an error with a specific identifier so that users can later retrieve it (using `lasterr` or `lasterror`) and know the origin of the error. The identifier must contain at least one colon character (`':'`) and must not contain any whitespace characters. It should be a string of the form `"NAMESPACE:ERROR-NAME"`. Octave's own errors use the `"Octave"` namespace (see [error_ids], page 260). For example:

```
error ("MyNameSpace:wrong-type-argument",
       "fcn_name: argument should be numeric");
```

Calling `error` also sets Octave's internal error state such that control will return to the top level without evaluating any further commands. This is useful for aborting from functions or scripts.

If the error message does not end with a newline character, Octave will print a traceback of all the function calls leading to the error. For example, given the following function definitions:

```
function f () g (); end
function g () h (); end
function h () nargin == 1 || error ("nargin != 1"); end
```

calling the function `f` will result in a list of messages that can help you to quickly find the exact location of the error:

```
f ()
error: nargin != 1
error: called from:
error:   h at line 1, column 27
error:   g at line 1, column 15
error:   f at line 1, column 15
```

If the error message ends in a newline character, Octave will print the message but will not display any traceback messages as it returns control to the top level. For example, modifying the error message in the previous example to end in a newline causes Octave to only print a single message:

```
function h () nargin == 1 || error ("nargin != 1\n"); end
f ()
error: nargin != 1
```

A null string (`""`) input to `error` will be ignored and the code will continue running as if the statement were a NOP. This is for compatibility with MATLAB. It also makes it possible to write code such as

```
err_msg = "";
if (CONDITION 1)
  err_msg = "CONDITION 1 found";
elseif (CONDITION2)
  err_msg = "CONDITION 2 found";
...
endif
error (err_msg);
```

which will only stop execution if an error has been found.

The function may also be called with an error structure such as that returned from `lasterror`. The *errstruct* argument must contain fields `message`, `identifier`, and `stack`. The first two fields are strings with the meanings discussed above. The `stack` field must be a structure or structure array with fields `file`, `name`, and `line`.

Implementation Note: For compatibility with MATLAB, escape sequences in *template* (e.g., `"\n"` => newline) are processed regardless of whether *template* has been defined with single quotes, as long as there are two or more input arguments. To disable escape sequence expansion use a second backslash before the sequence (e.g., `"\\n"`) or use the `regexptranslate` function.

**See also:** [warning], page 263, [lasterror], page 259.

Since it is common to use errors when there is something wrong with the input to a function, Octave supports functions to simplify such code. When the `print_usage` function is called, it reads the help text of the function calling `print_usage`, and presents a useful error. If the help text is written in Texinfo it is possible to present an error message that only contains the function prototypes as described by the `@deftypefn` parts of the help text. When the help text isn't written in Texinfo, the error message contains the entire help message.

Consider the following function.

```
## -*- texinfo -*-
## @deftypefn {} f (@var{arg1})
## Function help text goes here...
## @end deftypefn
function f (arg1)
  if (nargin == 0)
    print_usage ();
  endif
endfunction
```

When it is called with no input arguments it produces the following error.

```
f ()

 -|  error: Invalid call to f.  Correct usage is:
 -|
 -|   -- f (ARG1)
 -|
 -|
 -|  Additional help for built-in functions and operators is
 -|  available in the online version of the manual.  Use the command
 -|  'doc <topic>' to search the manual index.
 -|
 -|  Help and information about Octave is also available on the WWW
 -|  at https://www.octave.org and via the help@octave.org
 -|  mailing list.
```

`print_usage ()`
`print_usage (`*`name`*`)`
Print the usage message for the function *name*.

When called with no input arguments the `print_usage` function displays the usage message of the currently executing function.

**See also:** [help], page 21.

`beep ()`
Produce a beep from the speaker (or visual bell).

This function sends the alarm character `"\a"` to the terminal. Depending on the user's configuration this may produce an audible beep, a visual bell, or nothing at all.

**See also:** [puts], page 314, [fputs], page 313, [printf], page 315, [fprintf], page 315.

*`val`* `= beep_on_error ()`
*`old_val`* `= beep_on_error (`*`new_val`*`)`
*`old_val`* `= beep_on_error (`*`new_val`*`, "`*`local`*`")`
Query or set the internal variable that controls whether Octave will try to ring the terminal bell before printing an error message.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

### 12.1.2 Catching Errors

When an error occurs, it can be detected and handled using the `try` statement as described in Section 10.9 [The try Statement], page 202. As an example, the following piece of code counts the number of errors that occurs during a `for` loop.

```
number_of_errors = 0;
for n = 1:100
  try
    ...
  catch
    number_of_errors++;
  end_try_catch
endfor
```

The above example treats all errors the same. In many situations it can however be necessary to discriminate between errors, and take different actions depending on the error. The `lasterror` function returns a structure containing information about the last error that occurred. As an example, the code above could be changed to count the number of errors related to the '`*`' operator.

```
number_of_errors = 0;
for n = 1:100
  try
    ...
  catch
    msg = lasterror.message;
    if (strfind (msg, "operator *"))
      number_of_errors++;
    endif
  end_try_catch
endfor
```

Alternatively, the output of the `lasterror` function can be found in a variable indicated immediately after the `catch` keyword, as in the example below showing how to redirect an error as a warning:

```
try
  ...
catch err
  warning(err.identifier, err.message);
  ...
end_try_catch
```

*lasterr* = `lasterror` ()
`lasterror` (*err*)
`lasterror` ("*reset*")

Query or set the last error message structure.

When called without arguments, return a structure containing the last error message and other information related to this error. The elements of the structure are:

`message` The text of the last error message

`identifier`
: The message identifier of this error message

`stack` A structure containing information on where the message occurred. This may be an empty structure if the information cannot be obtained. The fields of the structure are:

  `file` The name of the file where the error occurred

  `name` The name of function in which the error occurred

  `line` The line number at which the error occurred

  `column` An optional field with the column number at which the error occurred

The last error structure may be set by passing a scalar structure, *err*, as input. Any fields of *err* that match those above are set while any unspecified fields are initialized with default values.

If `lasterror` is called with the argument `"reset"`, all fields are set to their default values.

**See also:** [lasterr], page 260, [error], page 255, [lastwarn], page 265.

`[`*msg*`, `*msgid*`] = lasterr ()`
`lasterr (`*msg*`)`
`lasterr (`*msg*`, `*msgid*`)`

Query or set the last error message.

When called without input arguments, return the last error message and message identifier.

With one argument, set the last error message to *msg*.

With two arguments, also set the last message identifier.

**See also:** [lasterror], page 259, [error], page 255, [lastwarn], page 265.

The next example counts indexing errors. The errors are caught using the field identifier of the structure returned by the function `lasterror`.

```
number_of_errors = 0;
for n = 1:100
  try
    ...
  catch
    id = lasterror.identifier;
    if (strcmp (id, "Octave:invalid-indexing"))
      number_of_errors++;
    endif
  end_try_catch
endfor
```

The functions distributed with Octave can issue one of the following errors.

`Octave:bad-alloc`
: Indicates that memory couldn't be allocated.

`Octave:invalid-context`
: Indicates the error was generated by an operation that cannot be executed in the scope from which it was called. For example, the function `print_usage ()` when called from the Octave prompt raises this error.

`Octave:invalid-fun-call`
: Indicates that a function was called in an incorrect way, e.g., wrong number of input arguments.

`Octave:invalid-indexing`
: Indicates that a data-type was indexed incorrectly, e.g., real-value index for arrays, nonexistent field of a structure.

`Octave:invalid-input-arg`
: Indicates that a function was called with invalid input arguments.

`Octave:parse-error`
: The interpreter failed to parse (read) specified code.

`Octave:undefined-function`
: Indicates a call to a function that is not defined. The function may exist but Octave is unable to find it in the search path.

When an error has been handled it is possible to raise it again. This can be useful when an error needs to be detected, but the program should still abort. This is possible using the `rethrow` function. The previous example can now be changed to count the number of errors related to the '`*`' operator, but still abort if another kind of error occurs.

```
number_of_errors = 0;
for n = 1:100
  try
    ...
  catch
    msg = lasterror.message;
    if (strfind (msg, "operator *"))
      number_of_errors++;
    else
      rethrow (lasterror);
    endif
  end_try_catch
endfor
```

`rethrow (`*err*`)`

Reissue a previous error as defined by *err*.

*err* is a structure that must contain at least the `"message"` and `"identifier"` fields. *err* can also contain a field `"stack"` that gives information on the assumed location of the error. Typically *err* is returned from `lasterror`.

**See also:** [lasterror], page 259, [lasterr], page 260, [error], page 255.

*err* `= errno ()`
*err* `= errno (`*val*`)`
*err* `= errno (`*name*`)`

Query or set the system-dependent variable errno.

When called with no inputs, return the current value of errno.

When called with a numeric input *val*, set the current value of errno to the specified value. The previous value of errno is returned as *err*.

When called with a character string *name*, return the numeric value of errno which corresponds to the specified error code. If *name* is not a recognized error code then -1 is returned.

**See also:** [errno_list], page 261.

*S* `= errno_list ()`

Return a structure containing the system-dependent errno values.

**See also:** [errno], page 261.

### 12.1.3 Recovering From Errors

Octave provides several ways of recovering from errors. There are `try/catch` blocks, `unwind_protect/unwind_protect_cleanup` blocks, and finally the `onCleanup` command.

The `onCleanup` command associates an ordinary Octave variable (the trigger) with an arbitrary function (the action). Whenever the Octave variable ceases to exist—whether

due to a function return, an error, or simply because the variable has been removed with `clear`—then the assigned function is executed.

The function can do anything necessary for cleanup such as closing open file handles, printing an error message, or restoring global variables to their initial values. The last example is a very convenient idiom for Octave code. For example:

```
function rand42
  old_state = rand ("state");
  restore_state = onCleanup (@() rand ("state", old_state));
  rand ("state", 42);
  ...
endfunction  # rand generator state restored by onCleanup
```

*obj* = onCleanup (*function*)

Create a special object that executes a given *function* upon destruction.

If the object is copied to multiple variables (or cell or struct array elements) or returned from a function, then *function* will be executed only after the last copy of the object is cleared.

The input *function* is a handle to a function. The handle may point to an anonymous function in order to directly place commands in the `onCleanup` call.

Programming Note: If multiple local `onCleanup` variables are created, the order in which they are called is unspecified. For similar functionality See Section 10.8 [The unwind_protect Statement], page 202.

Example

```
octave:1> trigger = onCleanup (@() disp ('onCleanup was executed'));
octave:2> clear trigger
onCleanup was executed
octave:3
```

## 12.2 Handling Warnings

Like an error, a warning is issued when something unexpected happens. Unlike an error, a warning doesn't abort the currently running program. A simple example of a warning is when a number is divided by zero. In this case Octave will issue a warning and assign the value `Inf` to the result.

```
a = 1/0
     ⊣ warning: division by zero
     ⇒ a = Inf
```

### 12.2.1 Issuing Warnings

It is possible to issue warnings from any code using the `warning` function. In its most simple form, the `warning` function takes a string describing the warning as its input argument. As an example, the following code controls if the variable 'a' is non-negative, and if not issues a warning and sets 'a' to zero.

```
a = -1;
if (a < 0)
  warning ("'a' must be non-negative.  Setting 'a' to zero.");
  a = 0;
endif
     ⊣ 'a' must be non-negative.  Setting 'a' to zero.
```

Since warnings aren't fatal to a running program, it is not possible to catch a warning using the `try` statement or something similar. It is however possible to access the last warning as a string using the `lastwarn` function.

It is also possible to assign an identification string to a warning. If a warning has such an ID the user can enable and disable this warning as will be described in the next section. To assign an ID to a warning, simply call `warning` with two string arguments, where the first is the identification string, and the second is the actual warning. Note that warning IDs are in the format `"NAMESPACE:WARNING-NAME"`. The namespace `"Octave"` is used for Octave's own warnings. Any other string is available as a namespace for user's own warnings.

`warning (`*`template`*`, ...)`
`warning (`*`id`*`, `*`template`*`, ...)`
`warning ("`*`on`*`", `*`id`*`)`
`warning ("`*`off`*`", `*`id`*`)`
`warning ("`*`error`*`", `*`id`*`)`
`warning ("`*`query`*`", `*`id`*`)`
`warning (`*`state`*`, `*`id`*`, "`*`local`*`")`
`warning (`*`warning_struct`*`)`
*`warning_struct`* `= warning (...)`
`warning (`*`state`*`, `*`mode`*`)`

Display a warning message or control the behavior of Octave's warning system.

The first call form uses a template *template* and optional additional arguments to display a message on the `stderr` stream. The message is formatted using the same rules as the `printf` family of functions (see Section 14.2.4 [Formatted Output], page 315) and prefixed by the character string '`warning: `'. You should use this function when you want to notify the user of an unusual condition, but only when it makes sense for your program to go on. For example:

```
warning ("foo: maybe something wrong here");
```

If the warning message does not end with a newline character, Octave will print a traceback of all the function calls leading to the warning. If the warning message does end in a newline character, Octave will suppress the traceback messages as it returns control to the top level. For more details and examples, see [`error`], page 255.

The optional warning identifier *id* allows users to enable or disable warnings tagged by this identifier. A message identifier is a string of the form `"NAMESPACE:WARNING-NAME"`. Octave's own warnings use the `"Octave"` namespace (see [`warning_ids`], page 265). For example:

```
warning ("MyNameSpace:check-something",
         "foo: maybe something wrong here");
```

The second call form is meant to change and/or query the state of warnings. The first input argument must be a string *state* (`"on"`, `"off"`, `"error"`, or `"query"`) followed by an optional warning identifier *id* or `"all"` (default).

The optional output argument *warning_struct* is a structure or structure array with fields `"state"` and `"identifier"`. The *state* argument may have the following values:

`"on"|"off"`:
: Enable or disable the display of warnings identified by *id* and optionally return their previous state *stout*.

`"error"`:
: Turn warnings identified by *id* into errors and optionally return their previous state *stout*.

`"query"`:
: Return the current state of warnings identified by *id*.

A structure or structure array *warning_struct*, with fields `"state"` and `"identifier"`, may be given as an input to achieve equivalent results. The following example shows how to temporarily disable a warning and then restore its original state:

```
loglog (-1:10);
## Disable the previous warning and save its original state
[~, id] = lastwarn ();
warnstate = warning ("off", id);
loglog (-1:10);
## Restore its original state
warning (warnstate);
```

If a final argument `"local"` is provided then the warning state will be set temporarily until the end of the current function. Changes to warning states that are set locally affect the current function and all functions called from the current scope. The previous warning state is restored on return from the current function. The `"local"` option is ignored if used in the top-level workspace.

With no input argument `warning ()` is equivalent to `warning ("query", "all")` except that in the absence of an output argument, the state of warnings is displayed on `stderr`.

The level of verbosity of the warning system may also be controlled by two modes *mode*:

`"backtrace"`:
: enable/disable the display of the stack trace after the warning message

`"verbose"`:
: enable/disable the display of additional information after the warning message

In this case the *state* argument may only be `"on"` or `"off"`.

Implementation Note: For compatibility with MATLAB, escape sequences in *template* (e.g., `"\n"` => newline) are processed regardless of whether *template* has been defined with single quotes, as long as there are two or more input arguments. To disable escape sequence expansion use a second backslash before the sequence (e.g., `"\\n"`) or use the `regexptranslate` function.

**See also:** [warning_ids], page 265, [lastwarn], page 265, [error], page 255.

`[`*msg*`, `*msgid*`] = lastwarn ()`
`lastwarn (`*msg*`)`
`lastwarn (`*msg*`, `*msgid*`)`

Query or set the last warning message.

When called without input arguments, return the last warning message and message identifier.

With one argument, set the last warning message to *msg*.

With two arguments, also set the last message identifier to *msgid*.

**See also:** [warning], page 263, [lasterror], page 259, [lasterr], page 260.

The functions distributed with Octave can issue one of the following warnings.

`Octave:abbreviated-property-match`

If the `Octave:abbreviated-property-match` warning is enabled, a warning is printed if a non-exact or ambiguous match is being used for a operation specifying an object property. For example, for a graphics object, *fig*, with the property `'displayname'`, `get (`*fig*`, 'dis')` elicits a warning if the `Octave:abbreviated-property-match` warning is enabled. By default, the `Octave:abbreviated-property-match` warning is enabled.

`Octave:addpath-pkg`

If the `Octave:addpath-pkg` warning is enabled, Octave will warn when a package directory (i.e., +package_name) is added to the `path`. Typically, only the parent directory which contains the package directory should be added to the load path. By default, the `Octave:addpath-pkg` warning is enabled.

`Octave:array-as-logical`

If the `Octave:array-as-logical` warning is enabled, Octave will warn when an array of size greater than 1x1 is used as a truth value in an if, while, or until statement. By default, the `Octave:array-as-logical` warning is disabled.

`Octave:array-to-scalar`

If the `Octave:array-to-scalar` warning is enabled, Octave will warn when an implicit conversion from an array to a scalar value is attempted. By default, the `Octave:array-to-scalar` warning is disabled.

`Octave:array-to-vector`

If the `Octave:array-to-vector` warning is enabled, Octave will warn when an implicit conversion from an array to a vector value is attempted. By default, the `Octave:array-to-vector` warning is disabled.

`Octave:assign-as-truth-value`

If the `Octave:assign-as-truth-value` warning is enabled, a warning is issued for statements like

```
if (s = t)
  ...
```

since such statements are not common, and it is likely that the intent was to write

```
if (s == t)
  ...
```

instead.

There are times when it is useful to write code that contains assignments within the condition of a `while` or `if` statement. For example, statements like

```
while (c = getc ())
  ...
```

are common in C programming.

It is possible to avoid all warnings about such statements by disabling the `Octave:assign-as-truth-value` warning, but that may also let real errors like

```
if (x = 1)  # intended to test (x == 1)!
  ...
```

slip by.

In such cases, it is possible suppress errors for specific statements by writing them with an extra set of parentheses. For example, writing the previous example as

```
while ((c = getc ()))
  ...
```

will prevent the warning from being printed for this statement, while allowing Octave to warn about other assignments used in conditional contexts.

By default, the `Octave:assign-as-truth-value` warning is enabled.

`Octave:autoload-relative-file-name`

If the `Octave:autoload-relative-file-name` is enabled, Octave will warn when parsing autoload() function calls with relative paths to function files. This usually happens when using autoload() calls in PKG_ADD files, when the PKG_ADD file is not in the same directory as the .oct file referred to by the autoload() command. By default, the `Octave:autoload-relative-file-name` warning is enabled.

`Octave:charmat-truncated`

If the `Octave:charmat-truncated` warning is enabled, a warning is printed when a character matrix with multiple rows is converted to a string. In this case, the Octave interpreter keeps only the first row and discards the others. By default, the `Octave:charmat-truncated` warning is enabled.

`Octave:classdef-to-struct`

If the `Octave:classdef-to-struct` warning is enabled, a warning is issued when a classdef object is forcibly converted into a struct with `struct (`*`CLASSDEF_OBJ`*`)`. Conversion removes the access restrictions from the object and makes private and protected properties visible. By default, the `Octave:classdef-to-struct` warning is enabled.

`Octave:colon-complex-argument`

If the `Octave:colon-complex-argument` warning is enabled, a warning is issued when one of the three arguments to the colon operator (base, increment, limit) is a complex value. For example, `1:3*i` will cause a warning to be emitted. By default, the `Octave:colon-complex-argument` warning is enabled.

`Octave:colon-nonscalar-argument`
: If the `Octave:colon-nonscalar-argument` warning is enabled, a warning is issued when one of the three arguments to the colon operator (base, increment, limit) is not a scalar. For example, `1:[3, 5]` will cause a warning to be emitted. By default, the `Octave:colon-nonscalar-argument` warning is enabled.

`Octave:data-file-in-path`
: If the `Octave:data-file-in-path` warning is enabled, a warning is issued when Octave does not find the target of a file operation such as `load` or `fopen` directly, but is able to locate the file in Octave's search `path` for files. The warning could indicate that a different file target than the programmer intended is being used. By default, the `Octave:data-file-in-path` warning is enabled.

`Octave:datevec:date-format-spec`
: If the `Octave:datevec:date-format-spec` warning is enabled, a warning is printed if the date format specification contains questionable date or time specifiers. Typical bad patterns are using uppercase date specifiers or lowercase time specifiers. By default, the `Octave:datevec:date-format-spec` warning is enabled.

`Octave:deprecated-function`
: If the `Octave:deprecated-function` warning is enabled, a warning is issued when Octave encounters a function that is obsolete and scheduled for removal from Octave. By default, the `Octave:deprecated-function` warning is enabled.

`Octave:deprecated-keyword`
: If the `Octave:deprecated-keyword` warning is enabled, a warning is issued when Octave encounters a keyword that is obsolete and scheduled for removal from Octave. By default, the `Octave:deprecated-keyword` warning is enabled.

`Octave:deprecated-option`
: If the `Octave:deprecated-option` warning is enabled, a warning is issued when an obsolete option or input to a function is used. By default, the `Octave:deprecated-option` warning is enabled.

`Octave:deprecated-property`
: If the `Octave:deprecated-property` warning is enabled, a warning is issued when Octave encounters a graphics property that is obsolete and scheduled for removal from Octave. By default, the `Octave:deprecated-property` warning is enabled.

`Octave:eigs:UnconvergedEigenvalues`
: If the `Octave:eigs:UnconvergedEigenvalues` warning is enabled then the eigs function will issue a warning if the number of calculated eigenvalues is less than the number of requested eigenvalues. By default, the `Octave:eigs:UnconvergedEigenvalues` warning is enabled.

`Octave:empty-index`
: If the `Octave:empty-index` warning is enabled then Octave will emit a warning whenever indexing operators are used without an index, for example `x()`. By default, the `Octave:empty-index` warning is enabled.

`Octave:erase:chararray`
: If the `Octave:erase:chararray` warning is enabled then the erase function will issue a warning if the input pattern is a character array rather than a string or cell array of strings. By default, the `Octave:erase:chararray` warning is enabled.

`Octave:function-name-clash`
: If the `Octave:function-name-clash` warning is enabled, a warning is issued when Octave finds that the name of a function defined in a function file differs from the name of the file. (If the names disagree, the name declared inside the file is ignored.) By default, the `Octave:function-name-clash` warning is enabled.

`Octave:future-time-stamp`
: If the `Octave:future-time-stamp` warning is enabled, Octave will print a warning if it finds a function file with a time stamp that is in the future. By default, the `Octave:future-time-stamp` warning is enabled.

`Octave:glyph-render`
: If the `Octave:glyph-render` warning is enabled, Octave will print a warning if the glyph for a character couldn't be rendered with the current font. By default, the `Octave:glyph-render` warning is enabled.

`Octave:imag-to-real`
: If the `Octave:imag-to-real` warning is enabled, a warning is printed for implicit conversions of complex numbers to real numbers. By default, the `Octave:imag-to-real` warning is disabled.

`Octave:infinite-loop`
: If the `Octave:infinite-loop` warning is enabled, a warning is printed when an infinite loop is detected such as `for i = 1:Inf` or `while (1)`. By default, the `Octave:infinite-loop` warning is enabled.

`Octave:language-extension`
: Print warnings when using features that are unique to the Octave language and that may still be missing in MATLAB. By default, the `Octave:language-extension` warning is disabled. The `--traditional` or `--braindead` startup options for Octave may also be of use, see Section 2.1.1 [Command Line Options], page 15.

`Octave:legacy-function`
: If the `Octave:legacy-function` warning is enabled, a warning is issued when Octave encounters a function that MATLAB has suggested should be avoided. The function may become obsolete at some point in the future and removed, in which case the warning will change to `Octave:deprecated-function`, and the function will continue to exist for two further versions of Octave before being removed. By default, the `Octave:legacy-function` warning is enabled.

`Octave:logical-conversion`
: If the `Octave:logical-conversion` warning is enabled, a warning is printed if an implicit conversion of an array from numerical to boolean occurs and any

of the elements in the array are not equal to zero or one. By default, the `Octave:logical-conversion` warning is enabled.

`Octave:lu:sparse_input`
: If the `Octave:lu:sparse_input` warning is enabled, Octave will warn when the lu function is called with a sparse input and less than four output arguments. In this case, sparsity-preserving column permutations are not performed and the result may be inaccurate. By default, the `Octave:lu:sparse_input` warning is enabled.

`Octave:missing-glyph`
: If the `Octave:glyph-render` warning is enabled, Octave will print a warning if the current font doesn't provide a glyph for a used character. By default, the `Octave:missing-glyph` warning is enabled.

`Octave:missing-semicolon`
: If the `Octave:missing-semicolon` warning is enabled, Octave will warn when statements in function definitions don't end in semicolons. By default the `Octave:missing-semicolon` warning is disabled.

`Octave:mixed-string-concat`
: If the `Octave:mixed-string-concat` warning is enabled, print a warning when concatenating a mixture of single and double quoted strings. By default, the `Octave:mixed-string-concat` warning is disabled.

`Octave:nearly-singular-matrix`
`Octave:singular-matrix`
: These warnings are emitted if a (nearly) singular matrix is inverted. By default, the `Octave:nearly-singular-matrix` and `Octave:singular-matrix` warnings are enabled.

`Octave:neg-dim-as-zero`
: If the `Octave:neg-dim-as-zero` warning is enabled, print a warning for expressions like

```
eye (-1)
```

By default, the `Octave:neg-dim-as-zero` warning is disabled.

`Octave:noninteger-range-as-index`
: If the `Octave:noninteger-range-as-index` warning is enabled, a warning is printed if an array is indexed with a range that contains non-integer values. For example,

```
a = [1 2 3 4 5];
b = 2.2:4.2
⇒ 1.2  2.2  3.2
a(b)
⇒ 2 3 4
```

elicits a warning if the `Octave:noninteger-range-as-index` warning is enabled. By default, the `Octave:noninteger-range-as-index` warning is enabled.

`Octave:num-to-str`
: If the `Octave:num-to-str` warning is enabled, a warning is printed for implicit conversions of numbers to their UTF-8 encoded character equivalents when strings are constructed using a mixture of strings and numbers in matrix notation. For example,

```
[ "f", 111, 111 ]
⇒ "foo"
```

elicits a warning if the `Octave:num-to-str` warning is enabled. By default, the `Octave:num-to-str` warning is enabled.

`Octave:possible-matlab-short-circuit-operator`
: If the `Octave:possible-matlab-short-circuit-operator` warning is enabled, Octave will warn about using the not short circuiting operators `&` and `|` inside `if` or `while` conditions. They normally never short circuit, but they do short circuit when used in a condition. By default, the `Octave:possible-matlab-short-circuit-operator` warning is enabled.

`Octave:pow2:imaginary-ignored`
: If the `Octave:pow2:imaginary-ignored` warning is enabled, a warning is printed if either input to `pow2` is complex. By default, the `Octave:pow2:imaginary-ignored` warning is enabled.

`Octave:recursive-path-search`
: If the `Octave:recursive-path-search` warning is enabled, Octave will issue a warning if `addpath` is used with double trailing slashes. By default, the `Octave:recursive-path-search` warning is enabled.

`Octave:remove-init-dir`
: The `path` function changes the search path that Octave uses to find functions. It is possible to set the path to a value which excludes Octave's own built-in functions. If the `Octave:remove-init-dir` warning is enabled then Octave will warn when the `path` function has been used in a way that may render Octave unworkable. By default, the `Octave:remove-init-dir` warning is enabled.

`Octave:reload-forces-clear`
: If several functions have been loaded from the same file, Octave must clear all the functions before any one of them can be reloaded. If the `Octave:reload-forces-clear` warning is enabled, Octave will warn you when this happens, and print a list of the additional functions that it is forced to clear. By default, the `Octave:reload-forces-clear` warning is enabled.

`Octave:separator-insert`
: Print warning if commas or semicolons might be inserted automatically in literal matrices. By default, the `Octave:separator-insert` warning is disabled.

`Octave:shadowed-function`
: If the `Octave:shadowed-function` warning is enabled, Octave will warn if a path is added to the search path that contains functions that shadow core functions. By default, the `Octave:shadowed-function` warning is enabled.

`Octave:single-quote-string`
: Print warning if a single quote character is used to introduce a string constant. By default, the `Octave:single-quote-string` warning is disabled.

`Octave:sparse:double-conversion`
: If the `Octave:sparse:double-conversion` warning is enabled, a warning is printed when an implicit conversion from a full, single array occurs during the creation of a sparse array. By default, the `Octave:sparse:double-conversion` warning is enabled.

`Octave:sqrtm:SingularMatrix`
: If the `Octave:sqrtm:SingularMatrix` warning is enabled, a warning is printed if the matrix square root function `sqrtm` is called with an input matrix that is singular. By default, the `Octave:sqrtm:SingularMatrix` warning is enabled.

`Octave:str-to-num`
: If the `Octave:str-to-num` warning is enabled, a warning is printed for implicit conversions of strings to their numeric UTF-8 encoded byte sequences. For example,

```
"abc" + 0
⇒ 97 98 99
```

  elicits a warning if the `Octave:str-to-num` warning is enabled. By default, the `Octave:str-to-num` warning is disabled.

`Octave:LaTeX:internal-error`
: If the `Octave:LaTeX:internal-error` warning is enabled, a warning is printed whenever the LaTeX renderer for text in plots encounters an issue. By default, the `Octave:LaTeX:internal-error` warning is enabled.

`Octave:unimplemented-matlab-functionality`
: If the `Octave:unimplemented-matlab-functionality` warning is enabled, a warning is printed when a MATLAB code construct is used which the Octave interpreter parses as valid, but for which Octave does not yet implement the functionality. By default, the `Octave:unimplemented-matlab-functionality` warning is enabled.

`Octave:variable-switch-label`
: If the `Octave:variable-switch-label` warning is enabled, Octave will print a warning if a switch label is not a constant or constant expression. By default, the `Octave:variable-switch-label` warning is disabled.

### 12.2.2 Enabling and Disabling Warnings

The `warning` function also allows you to control which warnings are actually printed to the screen. If the `warning` function is called with a string argument that is either `"on"` or `"off"` all warnings will be enabled or disabled.

It is also possible to enable and disable individual warnings through their string identifications. The following code will issue a warning

```
warning ("example:non-negative-variable",
         "'a' must be non-negative.  Setting 'a' to zero.");
```

while the following won't issue a warning

```
warning ("off", "example:non-negative-variable");
warning ("example:non-negative-variable",
         "'a' must be non-negative.  Setting 'a' to zero.");
```

# 13 Debugging

Octave includes a built-in debugger to aid in the development of scripts. This can be used to interrupt the execution of an Octave script at a certain point, or when certain conditions are met. Once execution has stopped, and debug mode is entered, the symbol table at the point where execution has stopped can be examined and modified to check for errors.

The normal command-line editing and history functions are available in debug mode.

## 13.1 Entering Debug Mode

There are two basic means of interrupting the execution of an Octave script. These are breakpoints (see Section 13.3 [Breakpoints], page 274), discussed in the next section, and interruption based on some condition.

Octave supports three means to stop execution based on the values set in the functions `debug_on_interrupt`, `debug_on_warning`, and `debug_on_error`.

*val* = debug_on_interrupt ()
*old_val* = debug_on_interrupt (*new_val*)
*old_val* = debug_on_interrupt (*new_val*, "*local*")

Query or set the internal variable that controls whether Octave will try to enter debugging mode when it receives an interrupt signal (typically generated with `C-c`).

If a second interrupt signal is received before reaching the debugging mode, a normal interrupt will occur.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [debug_on_error], page 273, [debug_on_warning], page 273.

*val* = debug_on_warning ()
*old_val* = debug_on_warning (*new_val*)
*old_val* = debug_on_warning (*new_val*, "*local*")

Query or set the internal variable that controls whether Octave will try to enter the debugger when a warning is encountered.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [debug_on_error], page 273, [debug_on_interrupt], page 273.

*val* = debug_on_error ()
*old_val* = debug_on_error (*new_val*)
*old_val* = debug_on_error (*new_val*, "*local*")

Query or set the internal variable that controls whether Octave will try to enter the debugger when an error is encountered.

This will also inhibit printing of the normal traceback message (you will only see the top-level error message).

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [debug_on_warning], page 273, [debug_on_interrupt], page 273.

## 13.2 Leaving Debug Mode

Use either `dbcont` or `return` to leave the debug mode and continue the normal execution of the script.

`dbcont`
: Leave command-line debugging mode and continue code execution normally.

  **See also:** [dbstep], page 279, [dbquit], page 274.

To quit debug mode and return directly to the prompt without executing any additional code use `dbquit`.

`dbquit`
`dbquit` *all*
: Quit debugging mode immediately without further code execution.

  With no arguments, exit the current debugging level. With argument `all`, exit all debugging levels and return to the Octave prompt.

  **See also:** [dbcont], page 274, [dbstep], page 279.

Finally, typing `exit` or `quit` at the debug prompt will result in Octave terminating normally.

## 13.3 Breakpoints

Breakpoints can be set in any m-file function by using the `dbstop` function.

`dbstop` *fcn*
`dbstop` *fcn line*
`dbstop` *fcn line1 line2 ...*
`dbstop` *line1 ...*
`dbstop` *in fcn*
`dbstop` *in fcn at line*
`dbstop` *in fcn at line if "condition"*
`dbstop` *in class at method*
`dbstop` *if event*
`dbstop` *if event ID*
`dbstop (`*bp_struct*`)`
*rline* `= dbstop (...)`
: Set breakpoints for the built-in debugger.

  *fcn* is the name of a function on the current `path`. When already in debug mode the *fcn* argument can be omitted and the current function will be used. Breakpoints at subfunctions are set with the scope operator '`>`'. For example, If `file.m` has a subfunction `fcn2`, then a breakpoint in `fcn2` can be specified by `file>fcn2`.

*line* is the line number at which to break. If *line* is not specified, it defaults to the first executable line in the file `fcn.m`. Multiple lines can be specified in a single command; when function syntax is used, the lines may also be passed as a single vector argument (`[line1, line2, ...]`).

*condition* is any Octave expression that can be evaluated in the code context that exists at the breakpoint. When the breakpoint is encountered, *condition* will be evaluated, and execution will stop if *condition* is true. If *condition* cannot be evaluated, for example because it refers to an undefined variable, an error will be thrown. Expressions with side effects (such as `y++ > 1`) will alter variables, and should generally be avoided. Conditions containing quotes (‘`"`’, ‘`'`’) or comment characters (‘`#`’, ‘`%`’) must be enclosed in quotes. (This does not apply to conditions entered from the editor's context menu.) For example:

```
dbstop in axis at 246 if 'any (opt == "x")'
```

The form specifying *event* does not cause a specific breakpoint at a given function and line number. Instead it causes debug mode to be entered when certain unexpected events are encountered. Possible values are

`error` Stop when an error is reported. This is equivalent to specifying both `debug_on_error (true)` and `debug_on_interrupt (true)`.

`caught error`
Stop when an error is caught by a try-catch block (not yet implemented).

`interrupt`
Stop when an interrupt (*Ctrl-C*) occurs.

`naninf` Stop when code returns a non-finite value (not yet implemented).

`warning` Stop when a warning is reported. This is equivalent to specifying `debug_on_warning (true)`.

The events `error`, `caught error`, and `warning` can all be followed by a string specifying an error ID or warning ID. If that is done, only errors with the specified ID will cause execution to stop. To stop on one of a set of IDs, multiple `dbstop` commands must be issued.

Breakpoints and events can be removed using the `dbclear` command with the same syntax.

It is possible to save all breakpoints and restore them at once by issuing the commands `bp_state = dbstatus; ...; dbstop (bp_state)`.

The optional output *rline* is the real line number where the breakpoint was set. This can differ from the specified line if the line is not executable. For example, if a breakpoint attempted on a blank line then Octave will set the real breakpoint at the next executable line.

When a file is re-parsed, such as when it is modified outside the GUI, all breakpoints within the file are cleared.

**See also:** [dbclear], page 276, [dbstatus], page 276, [dbstep], page 279, [debug_on_error], page 273, [debug_on_warning], page 273, [debug_on_interrupt], page 273.

Breakpoints in class methods are also supported (e.g., `dbstop ("@class/method")`). However, breakpoints cannot be set in built-in functions (e.g., `sin`, etc.) or dynamically loaded functions (i.e., oct-files).

To set a breakpoint immediately upon entering a function use line number 1, or omit the line number entirely and just give the function name. When setting the breakpoint Octave will ignore the leading comment block, and the breakpoint will be set on the first executable statement in the function. For example:

```
dbstop ("asind", 1)
⇒ 29
```

Note that the return value of `29` means that the breakpoint was effectively set to line 29. The status of breakpoints in a function can be queried with `dbstatus`.

`dbstatus`
`dbstatus` *`fcn`*
*`bp_list`* `= dbstatus ()`
*`bp_list`* `= dbstatus (`*`fcn`*`)`

Report the location of active breakpoints.

When called with no input or output arguments, print the list of all functions with breakpoints and the line numbers where those breakpoints are set.

If a function name *fcn* is specified then only report breakpoints for the named function and its subfunctions.

The optional return argument *bp_list* is a struct array with the following fields:

| | |
|---|---|
| name | The name of the function with a breakpoint. A subfunction, say `fcn2` within an m-file, say `file.m`, is specified as `file>fcn2`. |
| file | The name of the m-file where the function code is located. |
| line | The line number with the breakpoint. |
| cond | The condition that must be satisfied for the breakpoint to be active, or the empty string for unconditional breakpoints. |

If `dbstop if error` is true but no explicit IDs are specified, the return value will have an empty field called `"errs"`. If IDs are specified, the `errs` field will have one row per ID. If `dbstop if error` is false, there is no `"errs"` field. The `"warn"` field is set similarly by `dbstop if warning`.

**See also:** [dbstop], page 274, [dbclear], page 276, [dbwhere], page 278, [dblist], page 278, [dbstack], page 279.

Reusing the previous example, `dbstatus ("asind")` will return 29. The breakpoints listed can then be cleared with the `dbclear` function.

`dbclear` *`fcn`*
`dbclear` *`fcn line`*
`dbclear` *`fcn line1 line2 ...`*
`dbclear` *`line ...`*
`dbclear` *all*
`dbclear` *in `fcn`*

`dbclear` *in* `fcn` *at* `line`
`dbclear` *if* `event`
`dbclear` ("`fcn`")
`dbclear` ("`fcn`", `line`)
`dbclear` ("`fcn`", `line1`, `line2`, ...)
`dbclear` ("`fcn`", `line1`, ...)
`dbclear` (`line`, ...)
`dbclear` ("*all*")

Delete a breakpoint at line number *line* in the function *fcn*.

Arguments are

*fcn* Function name as a string variable. When already in debug mode this argument can be omitted and the current function will be used.

*line* Line number from which to remove a breakpoint. Multiple lines may be given as separate arguments or as a vector.

*event* An event such as `error`, `interrupt`, or `warning` (see [dbstop], page 274, for details).

When called without a line number specification all breakpoints in the named function are cleared.

If the requested line is not a breakpoint no action is performed.

The special keyword `"all"` will clear all breakpoints from all files.

**See also:** [dbstop], page 274, [dbstatus], page 276, [dbwhere], page 278.

A breakpoint may also be set in a subfunction. For example, if a file contains the functions

```
function y = func1 (x)
  y = func2 (x);
endfunction
function y = func2 (x)
  y = x + 1;
endfunction
```

then a breakpoint can be set at the start of the subfunction directly with

```
dbstop func1>func2
⇒ 5
```

Note that '`>`' is the character that distinguishes subfunctions from the m-file containing them.

Another simple way of setting a breakpoint in an Octave script is the use of the `keyboard` function.

`keyboard` ()
`keyboard` ("`prompt`")

Stop m-file execution and enter debug mode.

When the `keyboard` function is executed, Octave prints a prompt and waits for user input. The input strings are then evaluated and the results are printed. This makes it

possible to examine the values of variables within a function, and to assign new values if necessary. To leave the prompt and return to normal execution type '`return`' or '`dbcont`'. The `keyboard` function does not return an exit status.

If `keyboard` is invoked without arguments, a default prompt of '`debug>` ' is used.

**See also:** [dbstop], page 274, [dbcont], page 274, [dbquit], page 274.

The `keyboard` function is placed in a script at the point where the user desires that the execution be stopped. It automatically sets the running script into the debug mode.

## 13.4 Debug Mode

There are three additional support functions that allow the user to find out where in the execution of a script Octave entered the debug mode, and to print the code in the script surrounding the point where Octave entered debug mode.

`dbwhere`

In debugging mode, report the current file and line number where execution is stopped.

**See also:** [dbstack], page 279, [dblist], page 278, [dbstatus], page 276, [dbcont], page 274, [dbstep], page 279, [dbup], page 280, [dbdown], page 280.

`dbtype`
`dbtype` *`lineno`*
`dbtype` *`startl:endl`*
`dbtype` *`startl`*`:`*`end`*
`dbtype` *`fcn`*
`dbtype` *`fcn lineno`*
`dbtype` *`fcn startl:endl`*
`dbtype` *`fcn`* *`startl`*`:`*`end`*

Display a script file with line numbers.

When called with no arguments in debugging mode, display the script file currently being debugged.

An optional range specification can be used to list only a portion of the file. The special keyword `"end"` is a valid line number specification for the last line of the file.

When called with the name of a function, list that script file with line numbers.

**See also:** [dblist], page 278, [dbwhere], page 278, [dbstatus], page 276, [dbstop], page 274.

`dblist`
`dblist` *`n`*

In debugging mode, list *n* lines of the function being debugged centered around the current line to be executed.

If unspecified *n* defaults to 10 (+/- 5 lines)

**See also:** [dbwhere], page 278, [dbtype], page 278, [dbstack], page 279.

You may also use `isdebugmode` to determine whether the debugger is currently active.

*tf* = isdebugmode ()
: Return true if in debugging mode, otherwise false.

  **See also:** [dbwhere], page 278, [dbstack], page 279, [dbstatus], page 276.

Debug mode also allows single line stepping through a function using the command `dbstep`.

dbstep
dbstep *n*
dbstep *in*
dbstep *out*
dbnext ...
: In debugging mode, execute the next *n* lines of code.

  If *n* is omitted, execute the next single line of code. If the next line of code is itself defined in terms of an m-file remain in the existing function.

  Using `dbstep in` will cause execution of the next line to step into any m-files defined on the next line.

  Using `dbstep out` will cause execution to continue until the current function returns.

  Programming Note: `dbnext` is an alias for `dbstep` and can be used interchangeably.

  **See also:** [dbcont], page 274, [dbquit], page 274.

When in debug mode the `RETURN` key will execute the last entered command. This is useful, for example, after hitting a breakpoint and entering `dbstep` once. After that, one can advance line by line through the code with only a single key stroke. This feature may be disabled using the `auto_repeat_debug_command` function.

*val* = auto_repeat_debug_command ()
*old_val* = auto_repeat_debug_command (*new_val*)
*old_val* = auto_repeat_debug_command (*new_val*, "*local*")
: Query or set the internal variable that controls whether debugging commands are automatically repeated when the input line is empty (typing just `RET`).

  When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

## 13.5 Call Stack

The function being debugged may be the leaf node of a series of function calls. After examining values in the current subroutine it may turn out that the problem occurred in earlier pieces of code. Use `dbup` and `dbdown` to move up and down through the series of function calls to locate where variables first took on the wrong values. `dbstack` shows the entire series of function calls and at what level debugging is currently taking place.

dbstack
dbstack *n*
dbstack *-completenames*

`[stack, idx] = dbstack (...)`

Display or return current debugging function stack information.

With optional argument *n*, omit the *n* innermost stack frames.

Although accepted, the argument `-completenames` is silently ignored. Octave always returns absolute filenames.

The arguments *n* and `-completenames` can both be specified and may appear in any order.

The optional return argument *stack* is a struct array with the following fields:

| | |
|---|---|
| file | The name of the m-file where the function code is located. |
| name | The name of the function with a breakpoint. |
| line | The line number of an active breakpoint. |
| column | The column number of the line where the breakpoint begins. |

The return argument *idx* specifies which element of the *stack* struct array is currently active.

**See also:** [dbup], page 280, [dbdown], page 280, [dbwhere], page 278, [dblist], page 278, [dbstatus], page 276.

`dbup`
`dbup n`

In debugging mode, move up the execution stack *n* frames.

If *n* is omitted, move up one frame.

**See also:** [dbstack], page 279, [dbdown], page 280.

`dbdown`
`dbdown n`

In debugging mode, move down the execution stack *n* frames.

If *n* is omitted, move down one frame.

**See also:** [dbstack], page 279, [dbup], page 280.

## 13.6 Profiling

Octave supports profiling of code execution on a per-function level. If profiling is enabled, each call to a function (supporting built-ins, operators, functions in oct- and mex-files, user-defined functions in Octave code and anonymous functions) is recorded while running Octave code. After that, this data can aid in analyzing the code behavior, and is in particular helpful for finding "hot spots" in the code which use up a lot of computation time and are the best targets to spend optimization efforts on.

The main command for profiling is `profile`, which can be used to start or stop the profiler and also to query collected data afterwards. The data is returned in an Octave data structure which can then be examined or further processed by other routines or tools.

`profile` *on*
`profile` *off*
`profile` *resume*
`profile` *clear*
*S* = `profile` ("*status*")
*T* = `profile` ("*info*")
Control the built-in profiler.

`profile on`
Start the profiler. Any previously collected data is cleared.

`profile off`
Stop profiling. The collected data can later be retrieved and examined with `T = profile ("info")`.

`profile clear`
Clear all collected profiler data and stop profiling.

`profile resume`
Restart profiling without clearing the old data. All newly collected statistics are added to the existing ones.

*S* = `profile ("status")`
Return a structure with information about the current status of the profiler. At the moment, the only field is `ProfilerStatus` which is either `"on"` or `"off"`.

*T* = `profile ("info")`
Return the collected profiling statistics in the structure *T*. The flat profile is returned in the field `FunctionTable` which is an array of structures, each entry corresponding to a function which was called and for which profiling statistics are present. In addition, the field `Hierarchical` contains the hierarchical call tree. Each node has an index into the `FunctionTable` identifying the function it corresponds to as well as data fields for number of calls and time spent at this level in the call tree.

**See also:** [profshow], page 281, [profexplore], page 282.

An easy way to get an overview over the collected data is `profshow`. This function takes the profiler data returned by `profile` as input and prints a flat profile, for instance:

```
  Function Attr     Time (s)        Calls
 ----------------------------------------
    >myfib    R        2.195        13529
binary <=             0.061        13529
 binary -             0.050        13528
 binary +             0.026         6764
```

This shows that most of the run time was spent executing the function '`myfib`', and some minor proportion evaluating the listed binary operators. Furthermore, it is shown how often the function was called and the profiler also records that it is recursive.

`profshow (`*`data`*`)`
`profshow (`*`data`*`, `*`n`*`)`
`profshow ()`
`profshow (`*`n`*`)`

Display flat per-function profiler results.

Print out profiler data (execution time, number of calls) for the most critical *n* functions. The results are sorted in descending order by the total time spent in each function. If *n* is unspecified it defaults to 20.

The input *data* is the structure returned by `profile ("info")`. If unspecified, `profshow` will use the current profile dataset.

The attribute column displays 'R' for recursive functions, and is blank for all other function types.

**See also:** [profexplore], page 282, [profile], page 280.

`profexport (`*`dir`*`)`
`profexport (`*`dir`*`, `*`data`*`)`
`profexport (`*`dir`*`, `*`name`*`)`
`profexport (`*`dir`*`, `*`name`*`, `*`data`*`)`

Export profiler data as HTML.

Export the profiling data in *data* into a series of HTML files in the folder *dir*. The initial file will be `data/index.html`.

If *name* is specified, it must be a string that contains a "name" for the profile being exported. This name is included in the HTML.

The input *data* is the structure returned by `profile ("info")`. If unspecified, `profexport` will use the current profile dataset.

**See also:** [profshow], page 281, [profexplore], page 282, [profile], page 280.

`profexplore ()`
`profexplore (`*`data`*`)`

Interactively explore hierarchical profiler output.

Assuming *data* is the structure with profile data returned by `profile ("info")`, this command opens an interactive prompt that can be used to explore the call-tree. Type *`help`* to get a list of possible commands. If *data* is omitted, `profile ("info")` is called and used in its place.

**See also:** [profile], page 280, [profshow], page 281.

## 13.7 Profiler Example

Below, we will give a short example of a profiler session. See Section 13.6 [Profiling], page 280, for the documentation of the profiler functions in detail. Consider the code:

```
global N A;

N = 300;
A = rand (N, N);
```

```
function xt = timesteps (steps, x0, expM)
  global N;

  if (steps == 0)
    xt = NA (N, 0);
  else
    xt = NA (N, steps);
    x1 = expM * x0;
    xt(:, 1) = x1;
    xt(:, 2 : end) = timesteps (steps - 1, x1, expM);
  endif
endfunction

function foo ()
  global N A;

  initial = @(x) sin (x);
  x0 = (initial (linspace (0, 2 * pi, N)))';

  expA = expm (A);
  xt = timesteps (100, x0, expA);
endfunction

function fib = bar (N)
  if (N <= 2)
    fib = 1;
  else
    fib = bar (N - 1) + bar (N - 2);
  endif
endfunction
```

If we execute the two main functions, we get:

```
tic; foo; toc;
⇒ Elapsed time is 2.37338 seconds.

tic; bar (20); toc;
⇒ Elapsed time is 2.04952 seconds.
```

But this does not give much information about where this time is spent; for instance, whether the single call to `expm` is more expensive or the recursive time-stepping itself. To get a more detailed picture, we can use the profiler.

```
profile on;
foo;
profile off;

data = profile ("info");
profshow (data, 10);
```

This prints a table like:

```
   #  Function Attr     Time (s)        Calls
-----------------------------------------------
   7      expm             1.034            1
   3  binary *             0.823          117
  41  binary \             0.188            1
  38  binary ^             0.126            2
  43 timesteps    R        0.111          101
  44        NA             0.029          101
  39  binary +             0.024            8
  34      norm             0.011            1
  40  binary -             0.004          101
  33   balance             0.003            1
```

The entries are the individual functions which have been executed (only the 10 most important ones), together with some information for each of them. The entries like '`binary *`' denote operators, while other entries are ordinary functions. They include both built-ins like `expm` and our own routines (for instance `timesteps`). From this profile, we can immediately deduce that `expm` uses up the largest proportion of the processing time, even though it is only called once. The second expensive operation is the matrix-vector product in the routine `timesteps`.[1]

Timing, however, is not the only information available from the profile. The attribute column shows us that `timesteps` calls itself recursively. This may not be that remarkable in this example (since it's clear anyway), but could be helpful in a more complex setting. As to the question of why is there a '`binary \`' in the output, we can easily shed some light on that too. Note that `data` is a structure array (Section 6.1.2 [Structure Arrays], page 121) which contains the field `FunctionTable`. This stores the raw data for the profile shown. The number in the first column of the table gives the index under which the shown function can be found there. Looking up `data.FunctionTable(41)` gives:

```
scalar structure containing the fields:

  FunctionName = binary \
  TotalTime =  0.18765
  NumCalls =  1
  IsRecursive = 0
  Parents =  7
  Children = [](1x0)
```

Here we see the information from the table again, but have additional fields `Parents` and `Children`. Those are both arrays, which contain the indices of functions which have directly called the function in question (which is entry 7, `expm`, in this case) or been called by it (no functions). Hence, the backslash operator has been used internally by `expm`.

Now let's take a look at `bar`. For this, we start a fresh profiling session (`profile on` does this; the old data is removed before the profiler is restarted):

[1] We only know it is the binary multiplication operator, but fortunately this operator appears only at one place in the code and thus we know which occurrence takes so much time. If there were multiple places, we would have to use the hierarchical profile to find out the exact place which uses up the time which is not covered in this example.

```
profile on;
bar (20);
profile off;

profshow (profile ("info"));
```

This gives:

```
   #            Function Attr     Time (s)        Calls
-----------------------------------------------------------
   1                 bar    R        2.091        13529
   2           binary <=             0.062        13529
   3            binary -             0.042        13528
   4            binary +             0.023         6764
   5             profile             0.000            1
   8               false             0.000            1
   6              nargin             0.000            1
   7           binary !=             0.000            1
   9 __profiler_enable__             0.000            1
```

Unsurprisingly, `bar` is also recursive. It has been called 13,529 times in the course of recursively calculating the Fibonacci number in a suboptimal way, and most of the time was spent in `bar` itself.

Finally, let's say we want to profile the execution of both `foo` and `bar` together. Since we already have the run-time data collected for `bar`, we can restart the profiler without clearing the existing data and collect the missing statistics about `foo`. This is done by:

```
profile resume;
foo;
profile off;

profshow (profile ("info"), 10);
```

As you can see in the table below, now we have both profiles mixed together.

```
   #  Function Attr     Time (s)        Calls
-----------------------------------------------
   1       bar    R        2.091        13529
  16      expm             1.122            1
  12  binary *             0.798          117
  46  binary \             0.185            1
  45  binary ^             0.124            2
  48 timesteps    R        0.115          101
   2 binary <=             0.062        13529
   3  binary -             0.045        13629
   4  binary +             0.041         6772
  49        NA             0.036          101
```

# 14 Input and Output

Octave supports several ways of reading and writing data to or from the prompt or a file. The simplest functions for data Input and Output (I/O) are easy to use, but only provide limited control of how data is processed. For more control, a set of functions modeled after the C standard library are also provided by Octave.

## 14.1 Basic Input and Output

### 14.1.1 Terminal Output

Since Octave normally prints the value of an expression as soon as it has been evaluated, the simplest of all I/O functions is a simple expression. For example, the following expression will display the value of 'pi'

```
pi
     ⊣ ans =  3.1416
```

This works well as long as it is acceptable to have the name of the variable (or 'ans') printed along with the value. To print the value of a variable without printing its name, use the function disp.

The format command offers some control over the way Octave prints values with disp and through the normal echoing mechanism.

disp (*x*)
*str* = disp (*x*)

Display the value of *x*.

For example:

```
disp ("The value of pi is:"), disp (pi)

     ⊣ the value of pi is:
     ⊣ 3.1416
```

Note that the output from disp always ends with a newline.

If an output value is requested, disp prints nothing and returns the formatted output in a string.

**See also:** [fdisp], page 300.

*str* = list_in_columns (*arg*, *width*, *prefix*)

Return a string containing the elements of *arg* listed in columns with an overall maximum width of *width* and optional prefix *prefix*.

The argument *arg* must be a cell array of character strings or a character array.

If *width* is not specified or is an empty matrix, or less than or equal to zero, the width of the terminal screen is used. Newline characters are used to break the lines in the output string. For example:

```
list_in_columns ({"abc", "def", "ghijkl", "mnop", "qrs", "tuv"}, 20)
     ⇒ abc     mnop
       def     qrs
       ghijkl  tuv

whos ans
     ⇒
     Variables in the current scope:

       Attr Name        Size                     Bytes  Class
       ==== ====        ====                     =====  =====
            ans         1x37                        37  char

     Total is 37 elements using 37 bytes
```

**See also:** [terminal_size], page 288.

`[`*rows*`, `*cols*`] = terminal_size ()`
`terminal_size (`[*rows*, *cols*]`)`

Query or set the size of the terminal window. If called with no arguments, return a two-element row vector containing the current size of the terminal window in characters (rows and columns). If called with a two-element vector of integer values, set the terminal size and return the previous setting. Setting the size manually should not be needed when using readline for command-line editing.

**See also:** [list_in_columns], page 287.

`format`
`format` *options*
`format (`*options*`)`
`[`*format*`, `*formatspacing*`, `*uppercase*`] = format`

Reset or specify the format of the output produced by `disp` and Octave's normal echoing mechanism.

This command only affects the display of numbers, not how they are stored or computed. To change the internal representation from the default double use one of the conversion functions such as `single`, `uint8`, `int64`, etc. Any `format` options that change the number of displayed significant digits will also be reflected by the `output_precision` function.

By default, Octave displays 5 significant digits in a human readable form (option '`short`', option '`lowercase`', and option '`loose`' format for matrices). If `format` is invoked without any options, or the option '`default`' is specified, then this default format is restored.

Valid format options for floating point numbers are listed in the following table.

`default` Restore the default format state described above.

`short` Fixed point format with 5 significant figures (default).

`long` Fixed point format with 16 significant figures.

As with the '`short`' format, Octave will switch to an exponential '`e`' format if it is unable to format a matrix properly using the current format.

`shorte`
`longe` Exponential format. The number to be represented is split between a mantissa and an exponent (power of 10). The mantissa has 5 significant digits in the short format. In the long format, double values are displayed with 16 significant digits and single values are displayed with 8. For example, with the '`shorte`' format, `pi` is displayed as `3.1416e+00`. Optionally, the trailing '`e`' can be split into a second argument.

`shortg`
`longg` Optimally choose between fixed point and exponential format based on the magnitude of the number. For example, with the '`shortg`' format, `pi .^ [2; 4; 8; 16; 32]` is displayed as

```
ans =

     9.8696
     97.409
     9488.5
  9.0032e+07
  8.1058e+15
```

Optionally, the trailing '`g`' can be split into a second argument.

`shorteng`
`longeng` Identical to '`shorte`' or '`longe`' but displays the value using an engineering format, where the exponent is divisible by 3. For example, with the '`shorteng`' format, `10 * pi` is displayed as `31.416e+00`. Optionally, the trailing '`eng`' can be split into a second argument.

`free`
`none` Print output in free format, without trying to line up columns of matrices on the decimal point. This is a raw format equivalent to the C++ code `std::cout << variable`. In general, the result is a presentation with 6 significant digits where unnecessary precision (such as trailing zeros for integers) is suppressed. Complex numbers are formatted as numeric pairs like this `(0.60419, 0.60709)` instead of like this `0.60419 + 0.60709i`.

The following formats affect all numeric output (floating point and integer types).

`"+"`
`"+" "chars"`
`plus`
`plus chars`
Print a '`+`' symbol for matrix elements greater than zero, a '`-`' symbol for elements less than zero, and a space for zero matrix elements. This format can be useful for examining the sparsity structure of a large matrix. For very large matrices the function `spy` which plots the sparsity pattern will be clearer.

The optional argument *chars* specifies a list of 3 characters to use for printing values greater than zero, less than zero, and equal to zero. For

example, with the format `"+" "+-."`, the matrix `[1, 0, -1; -1, 0, 1]` is displayed as

```
ans =

+.-
-.+
```

`bank` Print variable in a format appropriate for a currency (fixed format with two digits to the right of the decimal point). Only the real part of a variable is displayed, as the imaginary part makes no sense for a currency.

`bit` Print the bit representation of numbers in memory, always with the most significant bit first. For example, `pi` is printed like this:

```
0 10000000000 1001001000011111101101010100010001000010110100011000
```

where spaces have been added for clarity to show the sign bit, the 11-bit exponent, and the 52-bit mantissa, in that order. Together they represent `pi` as an IEEE 754 double precision floating point number in the normal form. Single precision floating point numbers are analogous.

`native-bit`
Print the bit representation of numbers as stored in memory. For big-endian machines, this is identical to the `format bit` layout seen above. For little-endian machines, it will print the bytes in the opposite order, though bits within a byte will still be presented with the most significant bit on the left.

For example, the value of `pi` in this format on x86-64 is:

```
00011000 00101101 01000100 01010100 11111011 00100001 00001001 01000000
```

shown here with spaces added for clarity. Compare with the previous bit string from `format bit` to see the grouping into bytes and their ordering.

`hex` The same as `format bit` above, except that bits are grouped four at a time into hexadecimal digits for brevity. Thus `pi` is represented as:

```
400921fb54442d18
```

`native-hex`
The same as `format native-bit` above, except that bits are grouped four at a time into hexadecimal digits for brevity. Thus `pi` is represented on an x86-64 as:

```
182d4454fb210940
```

`rat` Print a rational approximation, i.e., values are approximated as the ratio of small integers. For example, with the '`rat`' format, `pi` is displayed as `355/113`.

The following two options affect the display of scientific and hex notations.

`lowercase (default)`
Use a lowercase '`e`' for the exponent character in scientific notation and lowercase '`a-f`' for the hex digits representing 10-15.

`uppercase`
: Use an uppercase '`E`' for the exponent character in scientific notation and uppercase '`A-F`' for the hex digits representing 10-15.

The following two options affect the display of all matrices.

`compact`
: Remove blank lines around column number labels and between matrices producing more compact output with more data per page.

`loose (default)`
: Insert blank lines above and below column number labels and between matrices to produce a more readable output with less data per page.

If `format` is called with multiple competing options, the rightmost one is used, except for '`default`' which will override all other options. In case of an error the format remains unchanged.

If called with one to three output arguments, and no inputs, return the current format, format spacing, and uppercase preference. Specifying both outputs and inputs will produce an error.

**See also:** [fixed_point_format], page 55, [output_precision], page 54, [split_long_rows], page 54, [print_empty_dimensions], page 56, [rats], page 640.

#### 14.1.1.1 Paging Screen Output

When running interactively, Octave normally sends all output directly to the Command Window. However, when using the CLI version of Octave this can create a problem because large volumes of data will stream by before you can read them. In such cases, it is better to use a paging program such as `less` or `more` which displays just one screenful at a time. With `less` (and some versions of `more`) you can also scan forward and backward, and search for specific items. The pager is enabled by the command `more on`.

Normally, no output is displayed by the pager until just before Octave is ready to print the top level prompt, or read from the standard input (for example, by using the `fscanf` or `scanf` functions). This means that there may be some delay before any output appears on your screen if you have asked Octave to perform a significant amount of work with a single command statement. The function `fflush` may be used to force output to be sent to the pager (or any other stream) immediately.

You can select the program to run as the pager with the `PAGER` function, and configure the pager itself with the `PAGER_FLAGS` function.

`more`
`more` *on*
`more` *off*
: Turn output pagination on or off.

Without an argument, `more` toggles the current state.

The current state can be determined via `page_screen_output`.

Using the pager is not supported on Windows in the GUI.

**See also:** [page_screen_output], page 292, [page_output_immediately], page 292, [PAGER], page 292, [PAGER_FLAGS], page 292.

*val* = PAGER ()
*old_val* = PAGER (*new_val*)
*old_val* = PAGER (*new_val*, "*local*")

Query or set the internal variable that specifies the program to use to display terminal output on your system.

The default value is normally `"less"`, `"more"`, or `"pg"`, depending on what programs are installed on your system. See Appendix E [Installing Octave], page 1179.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [PAGER_FLAGS], page 292, [page_output_immediately], page 292, [more], page 291, [page_screen_output], page 292.

*val* = PAGER_FLAGS ()
*old_val* = PAGER_FLAGS (*new_val*)
*old_val* = PAGER_FLAGS (*new_val*, "*local*")

Query or set the internal variable that specifies the options to pass to the pager.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [PAGER], page 292, [more], page 291, [page_screen_output], page 292, [page_output_immediately], page 292.

*val* = page_screen_output ()
*old_val* = page_screen_output (*new_val*)
*old_val* = page_screen_output (*new_val*, "*local*")

Query or set the internal variable that controls whether output intended for the terminal window that is longer than one page is sent through a pager.

This allows you to view one screenful at a time. Some pagers (such as `less`—see Appendix E [Installing Octave], page 1179) are also capable of moving backward on the output.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [more], page 291, [page_output_immediately], page 292, [PAGER], page 292, [PAGER_FLAGS], page 292.

*val* = page_output_immediately ()
*old_val* = page_output_immediately (*new_val*)
*old_val* = page_output_immediately (*new_val*, "*local*")

Query or set the internal variable that controls whether Octave sends output to the pager as soon as it is available.

When the value is `false`, Octave buffers its output and waits until just before the prompt is printed to flush it to the pager. This is the default.

When `page_screen_output` is `false`, this variable has no effect.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [page_screen_output], page 292, [more], page 291, [PAGER], page 292, [PAGER_FLAGS], page 292.

*status* = `fflush` (*fid*)

Flush output to file descriptor *fid*.

`fflush` returns 0 on success and an OS dependent error value (−1 on Unix) on error.

Programming Note: Flushing is useful for ensuring that all pending output makes it to the screen before some other event occurs. For example, it is always a good idea to flush the standard output stream before calling `input`.

**See also:** [fopen], page 311, [fclose], page 313.

### 14.1.2 Terminal Input

Octave has three functions that make it easy to prompt users for input. The `input` and `menu` functions are normally used for managing an interactive dialog with a user, and the `keyboard` function is normally used for doing simple debugging.

*ans* = `input` (*prompt*)
*ans* = `input` (*prompt*, "*s*")

Print *prompt* and wait for user input.

For example,

```
input ("Pick a number, any number! ")
```

prints the prompt

```
Pick a number, any number!
```

and waits for the user to enter a value. The string entered by the user is evaluated as an expression, so it may be a literal constant, a variable name, or any other valid Octave code.

The number of return arguments, their size, and their class depend on the expression entered.

If you are only interested in getting a literal string value, you can call `input` with the character string `"s"` as the second argument. This tells Octave to return the string entered by the user directly, without evaluating it first.

Because there may be output waiting to be displayed by the pager, it is a good idea to always call `fflush (stdout)` before calling `input`. This will ensure that all pending output is written to the screen before your prompt.

**See also:** [yes_or_no], page 294, [kbhit], page 294, [pause], page 1032, [menu], page 293, [listdlg], page 1002.

*choice* = `menu` (*title*, *opt1*, ...)
*choice* = `menu` (*title*, {*opt1*, ...})

Display a menu with heading *title* and options *opt1*, ..., and wait for user input.

If the GUI is running, the menu is displayed graphically using `listdlg`. Otherwise, the title and menu options are printed on the console.

*title* is a string and the options may be input as individual strings or as a cell array of strings.

The return value *choice* is the number of the option selected by the user counting from 1. If the user aborts the dialog or makes an invalid selection then 0 is returned.

This function is useful for interactive programs. There is no limit to the number of options that may be passed in, but it may be confusing to present more than will fit easily on one screen.

**See also:** [input], page 293, [listdlg], page 1002.

*ans* = yes_or_no ("*prompt*")

Ask the user a yes-or-no question.

Return logical true if the answer is yes or false if the answer is no.

Takes one argument, *prompt*, which is the string to display when asking the question. *prompt* should end in a space; `yes-or-no` adds the string '`(yes or no)` ' to it. The user must confirm the answer with `RET` and can edit it until it has been confirmed.

**See also:** [input], page 293.

For `input`, the normal command line history and editing functions are available at the prompt.

Octave also has a function that makes it possible to get a single character from the keyboard without requiring the user to type a carriage return.

*c* = kbhit ()
*c* = kbhit (*1*)

Read a single keystroke from the keyboard.

If called with an argument (typically 1), don't wait for a keypress and immediately return the next keystroke in the keyboard input buffer or an empty string (`""`) if no keystroke is available.

For example,

```
c = kbhit ();
```

will set *c* to the next character typed at the keyboard as soon as it is typed.

```
c = kbhit (1);
```

is identical to the above example, but doesn't wait for a keypress, returning the empty string if no key is available.

**See also:** [input], page 293, [pause], page 1032.

### 14.1.3 Simple File I/O

The `save` and `load` commands allow data to be written to and read from disk files in various formats. The default format of files written by the `save` command can be controlled using the functions `save_default_options` and `save_precision`.

As an example the following code creates a 3-by-3 matrix and saves it to the file '`myfile.mat`'.

```
A = [ 1:3; 4:6; 7:9 ];
save myfile.mat A
```

Once one or more variables have been saved to a file, they can be read into memory using the `load` command.

```
load myfile.mat
A
     ⊣ A =
     ⊣
     ⊣    1   2   3
     ⊣    4   5   6
     ⊣    7   8   9
```

`save` *file*
`save` *options file*
`save` *options file v1 v2 ...*
`save` *options file -struct STRUCT*
`save` *options file -struct STRUCT f1 f2 ...*
`save - v1 v2 ...`
*str* = `save ("-", "v1", "v2", ...)`

Save the named variables *v1*, *v2*, ..., in the file *file*.

If no variable names are listed, Octave saves all the variables in the current scope. Otherwise, full variable names or pattern syntax can be used to specify the variables to save. If the `-struct` modifier is used then the fields of the **scalar** struct are saved as if they were variables with the corresponding field names. The `-struct` option can be combined with specific field names *f1*, *f2*, ... to write only certain fields to the file.

The `save` command may also be invoked using the functional form

```
save ("-option1", ..., "file", "v1", ...)
```

where the *options*, *file*, and variable name arguments (*v1*, ...) must be specified as character strings.

Valid options for the `save` command are listed in the following table. Options that modify the output format override the format specified by `save_default_options`.

`-append` Append to the file instead of overwriting.

`-ascii` Save a matrix in a text file without a header or any other information. The matrix must be 2-D and only the real part of any complex value is written to the file. Numbers are stored in single-precision format and separated by spaces. Additional options for the `-ascii` format are

`-double` Store numbers in double-precision format.

`-tabs` Separate numbers with tabs.

`-binary` Save the data in Octave's binary data format.

`-float-binary`
Save the data in Octave's binary data format using just single precision. Use this format **only** if you know that all the values to be saved can be represented in single precision.

`-hdf5` Save the data in HDF5 format. (HDF5 is a free, portable, binary format developed by the National Center for Supercomputing Applications at the University of Illinois.) This format is only available if Octave was built with a link to the HDF5 libraries.

`-float-hdf5`
Save the data in HDF5 format using just single precision. Use this format **only** if you know that all the values to be saved can be represented in single precision.

`-text (default)`
Save the data in Octave's text data format. The [save_precision], page 297, function specifies the number of significant figures to use when saving data (default: 17). The header of the text data file can be configure with [save_header_format_string], page 298.

`-v7.3`
`-V7.3`
`-7.3` Octave does **not** yet implement saving in MATLAB's v7.3 binary data format.

`-v7`
`-V7`
`-7`
`-mat7-binary`
Save the data in MATLAB's v7 binary data format.

`-v6`
`-V6`
`-6`
`-mat`
`-mat-binary`
Save the data in MATLAB's v6 binary data format.

`-v4`
`-V4`
`-4`
`-mat4-binary`
Save the data in MATLAB's v4 binary data format.

`-zip`
`-z` Use the gzip algorithm to compress the file. This works on files that are compressed with gzip outside of Octave, and gzip can also be used to convert the files for backward compatibility. This option is only available if Octave was built with a link to the zlib libraries.

The list of variables to save may use wildcard patterns (glob patterns) containing the following special characters:

`?` Match any single character.

`*` Match zero or more characters.

`[ list ]` Match the list of characters specified by *list*. If the first character is `!` or `^`, match all characters except those specified by *list*. For example, the pattern `[a-zA-Z]` will match all lower and uppercase alphabetic characters.

Wildcards may also be used in the field name specifications when using the `-struct` modifier (but not in the struct name itself).

Programming Notes: If called with the special filename `"-"` the data to be saved is returned as a string rather than writing it to an actual file.

When saving global variables the global status of the variable is also stored. If the variable is restored at a later time using `load`, it will be restored as a global variable. Global status is *not* preserved if using a MATLAB binary data file format or the `-ascii` format.

Note that `classdef` objects will be saved as `struct` in file formats that support saving `struct`. That means they will not be loaded as `classdef` objects when loading the variables from the save file at a later point.

Example:

The command

```
save -binary data a b*
```

saves the variable 'a' and all variables beginning with 'b' to the file `data` in Octave's binary format.

**See also:** [load], page 298, [save_default_options], page 297, [save_header_format_string], page 298, [save_precision], page 297, [csvwrite], page 302, [dlmwrite], page 301, [fwrite], page 327.

There are three functions that modify the behavior of `save`.

*val* = save_default_options ()
*old_val* = save_default_options (*new_val*)
*old_val* = save_default_options (*new_val*, "*local*")

Query or set the internal variable that specifies the default options for the `save` command, and defines the default format.

The default value is `"-text"` (Octave's own text-based file format). See the documentation of the `save` command for other choices.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [save], page 295, [save_header_format_string], page 298, [save_precision], page 297.

*val* = save_precision ()
*old_val* = save_precision (*new_val*)
*old_val* = save_precision (*new_val*, "*local*")

Query or set the internal variable that specifies the number of digits to keep when saving data in text format.

The default value is 17 which is the minimum necessary for the lossless saving and restoring of IEEE 754 double values; For IEEE 754 single values the minimum value is 9. If file size is a concern, it is probably better to choose a binary format for saving data rather than to reduce the precision of the saved values.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [save_default_options], page 297.

*val* = save_header_format_string ()
*old_val* = save_header_format_string (*new_val*)
*old_val* = save_header_format_string (*new_val*, "*local*")

Query or set the internal variable that specifies the format string used for the comment line written at the beginning of text-format data files saved by Octave.

The format string is passed to `strftime` and must begin with the character '`#`' and contain no newline characters. If the value of `save_header_format_string` is the empty string, the header comment is omitted from text-format data files. The default value is

```
"# Created by Octave VERSION, %a %b %d %H:%M:%S %Y %Z"
```

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [strftime], page 1027, [save_default_options], page 297.

load *file*
load *options file*
load *options file v1 v2 ...*
S = load ("*options*", "*file*", "*v1*", "*v2*", ...)
load *file options*
load *file options v1 v2 ...*
S = load ("*file*", "*options*", "*v1*", "*v2*", ...)

Load the named variables *v1*, *v2*, ..., from the file *file*.

If no variables are specified then all variables found in the file will be loaded. Otherwise, full variable names or pattern syntax can be used to specify the variables to save. The format of the file is automatically detected but may be overridden by supplying the appropriate option.

The `load` command may also be invoked using the functional form

```
load ("-option1", ..., "file", "v1", ...)
```

where the *options*, *file*, and variable name arguments (*v1*, ...) must be specified as character strings.

Valid options for `load` are listed in the following table.

`-ascii` Force Octave to assume the file contains columns of numbers in text format without any header or other information. Data in the file will be loaded as a single numeric matrix with the name of the variable derived from the name of the file.

`-binary` Force Octave to assume the file is in Octave's binary format.

`-hdf5` Force Octave to assume the file is in HDF5 format. (HDF5 is a free, portable binary format developed by the National Center for Supercomputing Applications at the University of Illinois.) Note that `load` is only designed to read HDF5 files that were created by Octave with `save`, and attempts to read other HDF5 files may fail or produce unpredictable results. The `-hdf5` option provides a limited ability to read files created using MATLAB's `-v7.3` option (which saves in HDF5 format) although many data types are not yet supported. This format is only available if Octave was built with a link to the HDF5 libraries.

`-text` Force Octave to assume the file is in Octave's text format.

`-v7.3`
`-V7.3`
`-7.3` Octave does **not** yet implement MATLAB's v7.3 binary data format. As the v7.3 format is an HDF5 based format, the `"-hdf5"` option may be used to attempt to open a v7.3 format file, although most non-numeric data types are not yet supported. Note that Octave **can not** currently save in this format.

`-v7`
`-V7`
`-7`
`-mat7-binary`
Force Octave to assume the file is in MATLAB's version 7 binary format.

`-v6`
`-V6`
`-6`
`-mat`
`-mat-binary`
Force Octave to assume the file is in MATLAB's version 6 binary format.

`-v4`
`-V4`
`-4`
`-mat4-binary`
Force Octave to assume the file is in MATLAB's version 4 binary format.

The list of variables to load may use wildcard patterns (glob patterns) containing the following special characters:

`?` Match any single character.

`*` Match zero or more characters.

`[ list ]` Match the list of characters specified by *list*. If the first character is `!` or `^`, match all characters except those specified by *list*. For example, the pattern `[a-zA-Z]` will match all lower and uppercase alphabetic characters.

If invoked with a single output argument, Octave assigns loaded data to the output instead of inserting variables in the symbol table. If the data file contains only numbers (TAB- or space-delimited columns), a matrix of values is returned. Otherwise, `load` returns a structure with members corresponding to the names of the variables in the file.

The `load` command can read data stored in Octave's text and binary formats, MATLAB's binary format, and many simple formats such as comma-separated-values (CSV). If compiled with zlib support, it can also load gzip-compressed files. It will automatically detect the type of file and do conversion from different floating point formats (currently only IEEE big and little endian, though other formats may be added in the future).

Programming Note: If a variable that is not marked as global is loaded from a file when a global symbol with the same name already exists, it is loaded in the global symbol table. Also, if a variable is marked as global in a file and a local symbol exists, the local symbol is moved to the global symbol table and given the value from the file.

**See also:** [save], page 295, [csvread], page 302, [dlmread], page 302, [fread], page 325, [textscan], page 304.

*str* = fileread (*filename*)
*str* = fileread (*filename*, *param*, *value*, ...)

Read the contents of *filename* and return it as a string.

*param*, *value* are optional pairs of parameters and values. Valid options are:

`"Encoding"`

Specify encoding used when reading from the file. This is a character string of a valid encoding identifier. The default is `"utf-8"`.

**See also:** [fopen], page 311, [fread], page 325, [fscanf], page 321, [importdata], page 308, [textscan], page 304, [type], page 156.

*fmtstr* = native_float_format ()

Return the native floating point format as a string.

It is possible to write data to a file in a similar way to the `disp` function for writing data to the screen. The `fdisp` works just like `disp` except its first argument is a file pointer as created by `fopen`. As an example, the following code writes to data '`myfile.txt`'.

```
fid = fopen ("myfile.txt", "w");
fdisp (fid, "3/8 is ");
fdisp (fid, 3/8);
fclose (fid);
```

See Section 14.2.1 [Opening and Closing Files], page 311, for details on how to use `fopen` and `fclose`.

fdisp (*fid*, *x*)

Display the value of *x* on the stream *fid*.

For example:

```
fdisp (stdout, "The value of pi is:"), fdisp (stdout, pi)
```

```
⊣ the value of pi is:
⊣ 3.1416
```

Note that the output from `fdisp` always ends with a newline.

**See also:** [disp], page 287.

Octave can also read and write matrices text files such as comma separated lists.

`dlmwrite (`*`file`*`, `*`M`*`)`
`dlmwrite (`*`file`*`, `*`M`*`, `*`delim`*`, `*`r`*`, `*`c`*`)`
`dlmwrite (`*`file`*`, `*`M`*`, `*`key`*`, `*`val ...`*`)`
`dlmwrite (`*`file`*`, `*`M`*`, "-`*`append`*`", ...)`
`dlmwrite (`*`fid`*`, ...)`

Write the numeric matrix *M* to the text file *file* using a delimiter.

*file* should be a filename or a writable file ID given by `fopen`.

The parameter *delim* specifies the delimiter to use to separate values on a row. If no delimiter is specified the comma character '`,`' is used.

The value of *r* specifies the number of delimiter-only lines to add to the start of the file.

The value of *c* specifies the number of delimiters to prepend to each line of data.

If the argument `"-append"` is given, append to the end of *file*.

In addition, the following keyword value pairs may appear at the end of the argument list:

`"append"` Either `"on"` or `"off"`. See `"-append"` above.

`"delimiter"`
: See *delim* above.

`"newline"`
: The character(s) to separate each row. Three special cases exist for this option. `"unix"` is changed into `"\n"`, `"pc"` is changed into `"\r\n"`, and `"mac"` is changed into `"\r"`. Any other value is used directly as the newline separator.

`"roffset"`
: See *r* above.

`"coffset"`
: See *c* above.

`"precision"`
: The precision to use when writing the file. It can either be a format string (as used by fprintf) or a number of significant digits.

```
dlmwrite ("file.csv", reshape (1:16, 4, 4));
dlmwrite ("file.tex", a, "delimiter", "&", "newline", "\n")
```

**See also:** [dlmread], page 302, [csvread], page 302, [csvwrite], page 302.

*data* = dlmread (*file*)
*data* = dlmread (*file*, *sep*)
*data* = dlmread (*file*, *sep*, *r0*, *c0*)
*data* = dlmread (*file*, *sep*, *range*)
*data* = dlmread (..., "*emptyvalue*", *EMPTYVAL*)

Read numeric data from the text file *file* which uses the delimiter *sep* between data values.

If *sep* is not defined the separator between fields is determined from the file itself.

The optional scalar arguments *r0* and *c0* define the starting row and column of the data to be read. These values are indexed from zero, i.e., the first data row corresponds to an index of zero.

The *range* parameter specifies exactly which data elements are read. The first form of the parameter is a 4-element vector containing the upper left and lower right corners `[R0,C0,R1,C1]` where the indices are zero-based. To specify the last column—the equivalent of `end` when indexing—use the specifier `Inf`. Alternatively, a spreadsheet style form such as `"A2..Q15"` or `"T1:AA5"` can be used. The lowest alphabetical index `'A'` refers to the first column. The lowest row index is 1.

*file* should be a filename or a file id given by `fopen`. In the latter case, the file is read until end of file is reached.

The `"emptyvalue"` option may be used to specify the value used to fill empty fields. The default is zero. Note that any non-numeric values, such as text, are also replaced by the `"emptyvalue"`.

**See also:** [csvread], page 302, [textscan], page 304, [dlmwrite], page 301.

csvwrite (*filename*, *x*)
csvwrite (*filename*, *x*, *dlm_opt1*, ...)

Write the numeric matrix *x* to the file *filename* in comma-separated-value (CSV) format.

This function is equivalent to

```
dlmwrite (filename, x, ",", dlm_opt1, ...)
```

Any optional arguments are passed directly to `dlmwrite` (see [dlmwrite], page 301).

**See also:** [csvread], page 302, [dlmwrite], page 301, [dlmread], page 302.

*x* = csvread (*filename*)
*x* = csvread (*filename*, *dlm_opt1*, ...)

Read the comma-separated-value (CSV) file *filename* into the matrix *x*.

Note: only CSV files containing numeric data can be read.

This function is equivalent to

```
x = dlmread (filename, "," , dlm_opt1, ...)
```

Any optional arguments are passed directly to `dlmread` (see [dlmread], page 302).

**See also:** [dlmread], page 302, [textscan], page 304, [csvwrite], page 302, [dlmwrite], page 301.

Formatted data from can be read from, or written to, text files as well.

`[a, ...] = textread (`*`filename`*`)`
`[a, ...] = textread (`*`filename`*`, `*`format`*`)`
`[a, ...] = textread (`*`filename`*`, `*`format`*`, `*`n`*`)`
`[a, ...] = textread (`*`filename`*`, `*`format`*`, `*`prop1`*`, `*`value1`*`, ...)`
`[a, ...] = textread (`*`filename`*`, `*`format`*`, `*`n`*`, `*`prop1`*`, `*`value1`*`, ...)`
This function is obsolete. Use `textscan` instead.

Read data from a text file.

The file *filename* is read and parsed according to *format*. The function behaves like `strread` except it works by parsing a file instead of a string. See the documentation of `strread` for details.

In addition to the options supported by `strread`, this function supports two more:

- `"headerlines"`: The first *value* number of lines of *filename* are skipped.
- `"endofline"`: Specify a single character or `"\r\n"`. If no value is given, it will be inferred from the file. If set to `""` (empty string) EOLs are ignored as delimiters.

The optional input *n* (format repeat count) specifies the number of times the format string is to be used or the number of lines to be read, whichever happens first while reading. The former is equivalent to requesting that the data output vectors should be of length *N*. Note that when reading files with format strings referring to multiple lines, *n* should rather be the number of lines to be read than the number of format string uses.

If the format string is empty (not just omitted) and the file contains only numeric data (excluding headerlines), textread will return a rectangular matrix with the number of columns matching the number of numeric fields on the first data line of the file. Empty fields are returned as zero values.

Examples:

```
Assume a data file like:
1 a 2 b
3 c 4 d
5 e

[a, b] = textread (f, "%f %s")
returns two columns of data, one with doubles, the other a
cellstr array:
a = [1; 2; 3; 4; 5]
b = {"a"; "b"; "c"; "d"; "e"}

[a, b] = textread (f, "%f %s", 3)
(read data into two culumns, try to use the format string
three times)
returns
a = [1; 2; 3]
b = {"a"; "b"; "c"}
```

```
With a data file like:
1
a
2
b

[a, b] = textread (f, "%f %s", 2)
returns a = 1 and b = {"a"}; i.e., the format string is used
only once because the format string refers to 2 lines of the
data file.  To obtain 2x1 data output columns, specify N = 4
(number of data lines containing all requested data) rather
than 2.
```

**See also:** [textscan], page 304, [load], page 298, [dlmread], page 302, [fscanf], page 321, [strread], page 107.

*C* = textscan (*fid*, *format*)
*C* = textscan (*fid*, *format*, *repeat*)
*C* = textscan (*fid*, *format*, *param*, *value*, ...)
*C* = textscan (*fid*, *format*, *repeat*, *param*, *value*, ...)
*C* = textscan (*str*, ...)
[*C*, *position*, *errmsg*] = textscan (...)

Read data from a text file or string.

The string *str* or file associated with *fid* is read from and parsed according to *format*. The function is an extension of `strread` and `textread`. Differences include: the ability to read from either a file or a string, additional options, and additional format specifiers.

The input is interpreted as a sequence of words, delimiters (such as whitespace), and literals. The characters that form delimiters and whitespace are determined by the options. The format consists of format specifiers interspersed between literals. In the format, whitespace forms a delimiter between consecutive literals, but is otherwise ignored.

The output *C* is a cell array where the number of columns is determined by the number of format specifiers.

The first word of the input is matched to the first specifier of the format and placed in the first column of the output; the second is matched to the second specifier and placed in the second column and so forth. If there are more words than specifiers then the process is repeated until all words have been processed or the limit imposed by *repeat* has been met (see below).

The string *format* describes how the words in *str* should be parsed. As in *fscanf*, any (non-whitespace) text in the format that is not one of these specifiers is considered a literal. If there is a literal between two format specifiers then that same literal must appear in the input stream between the matching words.

The following specifiers are valid:

`%f`
`%f64`
`%n` The word is parsed as a number and converted to double.

`%f32` The word is parsed as a number and converted to single (float).

`%d`
`%d8`
`%d16`
`%d32`
`%d64` The word is parsed as a number and converted to int8, int16, int32, or int64. If no size is specified then int32 is used.

`%u`
`%u8`
`%u16`
`%u32`
`%u64` The word is parsed as a number and converted to uint8, uint16, uint32, or uint64. If no size is specified then uint32 is used.

`%s` The word is parsed as a string ending at the last character before whitespace, an end-of-line, or a delimiter specified in the options.

`%q` The word is parsed as a "quoted string". If the first character of the string is a double quote (") then the string includes everything until a matching double quote—including whitespace, delimiters, and end-of-line characters. If a pair of consecutive double quotes appears in the input, it is replaced in the output by a single double quote. For examples, the input "He said ""Hello""" would return the value 'He said "Hello"'.

`%c` The next character of the input is read. This includes delimiters, whitespace, and end-of-line characters.

`%[...]`
`%[^...]` In the first form, the word consists of the longest run consisting of only characters between the brackets. Ranges of characters can be specified by a hyphen; for example, %[0-9a-zA-Z] matches all alphanumeric characters (if the underlying character set is ASCII). Since MATLAB treats hyphens literally, this expansion only applies to alphanumeric characters. To include '-' in the set, it should appear first or last in the brackets; to include ']', it should be the first character. If the first character is '^' then the word consists of characters **not** listed.

`%N...` For %s, %c %d, %f, %n, %u, an optional width can be specified as %Ns, etc. where N is an integer > 1. For %c, this causes exactly N characters to be read instead of a single character. For the other specifiers, it is an upper bound on the number of characters read; normal delimiters can cause fewer characters to be read. For complex numbers, this limit applies to the real and imaginary components individually. For %f and %n, format specifiers like %N.Mf are allowed, where M is an upper bound on number of characters after the decimal point to be considered; subsequent digits are skipped. For example, the specifier %8.2f would read 12.345e6 as 1.234e7.

`%*...` The word specified by the remainder of the conversion specifier is skipped.

`literals` In addition the format may contain literal character strings; these will be skipped during reading. If the input string does not match this literal, the processing terminates.

Parsed words corresponding to the first specifier are returned in the first output argument and likewise for the rest of the specifiers.

By default, if there is only one input argument, *format* is `"%f"`. This means that numbers are read from the input into a single column vector. If *format* is explicitly empty (`""`) then textscan will return data in a number of columns matching the number of fields on the first data line of the input. Either of these is suitable only when the input is exclusively numeric.

For example, the string

```
str = "\
Bunny Bugs   5.5\n\
Duck Daffy  -7.5e-5\n\
Penguin Tux   6"
```

can be read using

```
a = textscan (str, "%s %s %f");
```

The optional numeric argument *repeat* can be used for limiting the number of items read:

-1 Read all of the string or file until the end (default).

N Read until the first of two conditions occurs: 1) the format has been processed N times, or 2) N lines of the input have been processed. Zero (0) is an acceptable value for *repeat*. Currently, end-of-line characters inside %q, %c, and %[...]$ conversions do not contribute to the line count. This is incompatible with MATLAB and may change in future.

The behavior of `textscan` can be changed via property/value pairs. The following properties are recognized:

`"BufSize"`
This specifies the number of bytes to use for the internal buffer. A modest speed improvement may be obtained by setting this to a large value when reading a large file, especially if the input contains long strings. The default is 4096, or a value dependent on *n* if that is specified.

`"CollectOutput"`
A value of 1 or true instructs `textscan` to concatenate consecutive columns of the same class in the output cell array. A value of 0 or false (default) leaves output in distinct columns.

`"CommentStyle"`
Specify parts of the input which are considered comments and will be skipped. *value* is the comment style and can be either (1) A string or 1x1 cell string, to skip everything to the right of it; (2) A cell array of two strings, to skip everything between the first and second strings. Comments are only parsed where whitespace is accepted and do not act as delimiters.

`"Delimiter"`
: If *value* is a string, any character in *value* will be used to split the input into words. If *value* is a cell array of strings, any string in the array will be used to split the input into words. (default value = any whitespace.)

`"EmptyValue"`
: Value to return for empty numeric values in non-whitespace delimited data. The default is NaN. When the data type does not support NaN (int32 for example), then the default is zero.

`"EndOfLine"`
: *value* can be either an empty or one character specifying the end-of-line character, or the pair `"\r\n"` (CRLF). In the latter case, any of `"\r"`, `"\n"` or `"\r\n"` is counted as a (single) newline. If no value is given, `"\r\n"` is used.

`"HeaderLines"`
: The first *value* number of lines of *fid* are skipped. Note that this does not refer to the first non-comment lines, but the first lines of any type.

`"MultipleDelimsAsOne"`
: If *value* is nonzero, treat a series of consecutive delimiters, without whitespace in between, as a single delimiter. Consecutive delimiter series need not be vertically aligned. Without this option, a single delimiter before the end of the line does not cause the line to be considered to end with an empty value, but a single delimiter at the start of a line causes the line to be considered to start with an empty value.

`"TreatAsEmpty"`
: Treat single occurrences (surrounded by delimiters or whitespace) of the string(s) in *value* as missing values.

`"ReturnOnError"`
: If set to numerical 1 or true, return normally as soon as an error is encountered, such as trying to read a string using `%f`. If set to 0 or false, return an error and no data.

`"Whitespace"`
: Any character in *value* will be interpreted as whitespace and trimmed; The default value for whitespace is `" \b\r\n\t"` (note the space). Unless whitespace is set to `""` (empty) AND at least one `"%s"` format conversion specifier is supplied, a space is always part of whitespace.

When the number of words in *str* or *fid* doesn't match an exact multiple of the number of format conversion specifiers, `textscan`'s behavior depends on whether the last character of the string or file is an end-of-line as specified by the `EndOfLine` option:

last character = end-of-line
: Data columns are padded with empty fields, NaN or 0 (for integer fields) so that all columns have equal length

last character is not end-of-line
: Data columns are not padded; `textscan` returns columns of unequal length

The second output *position* provides the location, in characters from the beginning of the file or string, where processing stopped.

**See also:** [dlmread], page 302, [fscanf], page 321, [load], page 298, [strread], page 107, [textread], page 302.

The `importdata` function has the ability to work with a wide variety of data.

```
A = importdata (fname)
A = importdata (fname, delimiter)
A = importdata (fname, delimiter, header_rows)
[A, delimiter] = importdata (...)
[A, delimiter, header_rows] = importdata (...)
```

Import data from the file *fname*.

Input parameters:

- *fname* The name of the file containing data.
- *delimiter* The character separating columns of data. Use `\t` for tab. (Only valid for ASCII files)
- *header_rows* The number of header rows before the data begins. (Only valid for ASCII files)

Different file types are supported:

- ASCII table

  Import ASCII table using the specified number of header rows and the specified delimiter.
- Image file
- MATLAB file
- Spreadsheet files (depending on external software)
- WAV file

**See also:** [textscan], page 304, [dlmread], page 302, [csvread], page 302, [load], page 298.

After importing, the data may need to be transformed before further analysis. The `rescale` function can shift and normalize a data set to a specified range.

```
B = rescale (A)
B = rescale (A, l, u)
B = rescale (..., "inputmin", inmin)
B = rescale (..., "inputmax", inmax)
```

Scale matrix elements to a specified range of values.

When called with a single matrix argument $A$, rescale elements to occupy the interval [0, 1].

The optional inputs `[l, u]` will scale $A$ to the interval with lower bound $l$ and upper bound $u$.

The optional input `"inputmin"` replaces all elements less than the specified value *inmin* with *inmin*. Similarly, the optional input `"inputmax"` replaces all elements greater than the specified value *inmax* with *inmax*. If unspecified the minimum and maximum are taken from the data itself (*`inmin`* `= min (A(:))` and *`inmax`* `= max (A(:))`).

Programming Notes: The applied formula is

$$B = l + \frac{A - inmin}{inmax - inmin} \cdot (u - l)$$

The class of the output matrix $B$ is single if the input $A$ is single, but otherwise is of class double for inputs which are of double, integer, or logical type.

**See also:** [bounds], page 832, [min], page 618, [max], page 617.

### 14.1.3.1 Saving Data on Unexpected Exits

If Octave for some reason exits unexpectedly it will by default save the variables available in the workspace to a file in the current directory. By default this file is named '`octave-workspace`' and can be loaded into memory with the `load` command. While the default behavior most often is reasonable it can be changed through the following functions.

*`val`* `= crash_dumps_octave_core ()`
*`old_val`* `= crash_dumps_octave_core (`*`new_val`*`)`
*`old_val`* `= crash_dumps_octave_core (`*`new_val`*`, "`*local*`")`

Query or set the internal variable that controls whether Octave tries to save all current variables to the file `octave-workspace` if it crashes or receives a hangup, terminate or similar signal.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [octave_core_file_limit], page 310, [octave_core_file_name], page 310, [octave_core_file_options], page 310.

*`val`* `= sighup_dumps_octave_core ()`
*`old_val`* `= sighup_dumps_octave_core (`*`new_val`*`)`
*`old_val`* `= sighup_dumps_octave_core (`*`new_val`*`, "`*local*`")`

Query or set the internal variable that controls whether Octave tries to save all current variables to the file `octave-workspace` if it receives a hangup signal.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

*`val`* `= sigquit_dumps_octave_core ()`
*`old_val`* `= sigquit_dumps_octave_core (`*`new_val`*`)`
*`old_val`* `= sigquit_dumps_octave_core (`*`new_val`*`, "`*local*`")`

Query or set the internal variable that controls whether Octave tries to save all current variables to the file `octave-workspace` if it receives a quit signal.

When called from inside a function with the "local" option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

*val* = sigterm_dumps_octave_core ()
*old_val* = sigterm_dumps_octave_core (*new_val*)
*old_val* = sigterm_dumps_octave_core (*new_val*, "*local*")

Query or set the internal variable that controls whether Octave tries to save all current variables to the file octave-workspace if it receives a terminate signal.

When called from inside a function with the "local" option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

*val* = octave_core_file_options ()
*old_val* = octave_core_file_options (*new_val*)
*old_val* = octave_core_file_options (*new_val*, "*local*")

Query or set the internal variable that specifies the options used for saving the workspace data if Octave aborts.

The value of octave_core_file_options should follow the same format as the options for the save function. The default value is Octave's binary format.

When called from inside a function with the "local" option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [crash_dumps_octave_core], page 309, [octave_core_file_name], page 310, [octave_core_file_limit], page 310.

*val* = octave_core_file_limit ()
*old_val* = octave_core_file_limit (*new_val*)
*old_val* = octave_core_file_limit (*new_val*, "*local*")

Query or set the internal variable that specifies the maximum amount of memory that Octave will save when writing a crash dump file.

The limit is measured in kilobytes and is applied to the top-level workspace. The name of the crash dump file is specified by *octave_core_file_name*.

If *octave_core_file_options* flags specify a binary format, then *octave_core_file_limit* will be approximately the maximum size of the file. If a text file format is used, then the file could be much larger than the limit. The default value is -1 (unlimited).

When called from inside a function with the "local" option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [crash_dumps_octave_core], page 309, [octave_core_file_name], page 310, [octave_core_file_options], page 310.

*val* = octave_core_file_name ()
*old_val* = octave_core_file_name (*new_val*)
*old_val* = octave_core_file_name (*new_val*, "*local*")

Query or set the internal variable that specifies the name of the file used for saving data from the top-level workspace if Octave aborts.

The default value is `"octave-workspace"`

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [crash_dumps_octave_core], page 309, [octave_core_file_name], page 310, [octave_core_file_options], page 310.

## 14.2 C-Style I/O Functions

Octave's C-style input and output functions provide most of the functionality of the C programming language's standard I/O library. The argument lists for some of the input functions are slightly different, however, because Octave has no way of passing arguments by reference.

In the following, *file* refers to a filename and `fid` refers to an integer file number, as returned by `fopen`.

There are three files that are always available. Although these files can be accessed using their corresponding numeric file ids, you should always use the symbolic names given in the table below, since it will make your programs easier to understand.

*fid* = `stdin` ()

Return the numeric value corresponding to the standard input stream.

When Octave is used interactively, stdin is filtered through the command line editing functions.

**See also:** [stdout], page 311, [stderr], page 311.

*fid* = `stdout` ()

Return the numeric value corresponding to the standard output stream.

Data written to the standard output may be filtered through the pager.

**See also:** [stdin], page 311, [stderr], page 311, [page_screen_output], page 292.

*fid* = `stderr` ()

Return the numeric value corresponding to the standard error stream.

Even if paging is turned on, the standard error is not sent to the pager. It is useful for error messages and prompts.

**See also:** [stdin], page 311, [stdout], page 311.

### 14.2.1 Opening and Closing Files

When reading data from a file it must be opened for reading first, and likewise when writing to a file. The `fopen` function returns a pointer to an open file that is ready to be read or written. Once all data has been read from or written to the opened file it should be closed. The `fclose` function does this. The following code illustrates the basic pattern for writing to a file, but a very similar pattern is used when reading a file.

```
filename = "myfile.txt";
fid = fopen (filename, "w");
# Do the actual I/O here...
fclose (fid);
```

*fid* = fopen (*name*)
*fid* = fopen (*name*, *mode*)
*fid* = fopen (*name*, *mode*, *arch*)
*fid* = fopen (*name*, *mode*, *arch*, *encoding*)
[*fid*, *msg*] = fopen (...)
*fid_list* = fopen ("*all*")
[*file*, *mode*, *arch*, *encoding*] = fopen (*fid*)

Open a file for low-level I/O or query open files and file descriptors.

The first form of the `fopen` function opens the named file with the specified mode (read-write, read-only, etc.), architecture interpretation (IEEE big endian, IEEE little endian, etc.) and file encoding, and returns an integer value that may be used to refer to the file later. If an error occurs, *fid* is set to −1 and *msg* contains the corresponding system error message. The *mode* is a one or two character string that specifies whether the file is to be opened for reading, writing, or both. The *encoding* is a character string with a valid encoding identifier. This encoding is used when strings are read from or written to the file. By default, that is UTF-8.

The second form of the `fopen` function returns a vector of file ids corresponding to all the currently open files, excluding the `stdin`, `stdout`, and `stderr` streams.

The third form of the `fopen` function returns information about the open file given its file id.

For example,

```
myfile = fopen ("splat.dat", "r", "ieee-le");
```

opens the file `splat.dat` for reading. If necessary, binary numeric values will be read assuming they are stored in IEEE 754 format with the least significant bit first, and then converted to the native representation.

Opening a file that is already open simply opens it again and returns a separate file id. It is not an error to open a file several times, though writing to the same file through several different file ids may produce unexpected results.

The possible values of *mode* are

‘`r`’ (default)
: Open a file for reading.

‘`w`’
: Open a file for writing. The previous contents are discarded.

‘`a`’
: Open or create a file for writing at the end of the file.

‘`r+`’
: Open an existing file for reading and writing.

‘`w+`’
: Open a file for reading or writing. The previous contents are discarded.

‘`a+`’
: Open or create a file for reading or writing at the end of the file.

Append a `"t"` to the mode string to open the file in text mode or a `"b"` to open in binary mode. On Windows systems, text mode reading and writing automatically converts linefeeds to the appropriate line end character for the system (carriage-return linefeed on Windows). The default when no mode is specified is binary.

Additionally, you may append a `"z"` to the mode string to open a gzipped file for reading or writing. For this to be successful, you must also open the file in binary mode.

The parameter *arch* is a string specifying the default data format for the file. Valid values for *arch* are:

`"native"` or `"n"` (default)
: The format of the current machine.

`"ieee-be"` or `"b"`
: IEEE big endian format.

`"ieee-le"` or `"l"`
: IEEE little endian format.

When opening a new file that does not yet exist, permissions will be set to `0666 - umask`.

Compatibility Note: Octave opens files using buffered I/O. Small writes are accumulated until an internal buffer is filled, and then everything is written in a single operation. This is very efficient and improves performance. MATLAB, however, opens files using flushed I/O where every write operation is immediately performed. If the write operation must be performed immediately after data has been written then the write should be followed by a call to `fflush` to flush the internal buffer.

**See also:** [fclose], page 313, [fgets], page 314, [fgetl], page 314, [fscanf], page 321, [fread], page 325, [fputs], page 313, [fdisp], page 300, [fprintf], page 315, [fwrite], page 327, [fskipl], page 315, [fseek], page 330, [frewind], page 331, [ftell], page 330, [feof], page 329, [ferror], page 329, [fclear], page 330, [fflush], page 293, [freport], page 330, [umask], page 1041.

*status* = fclose (*fid*)
*status* = fclose ("*all*")
: Close the file specified by the file descriptor *fid*.

If successful, `fclose` returns 0, otherwise, it returns -1. The second form of the `fclose` call closes all open files except `stdin`, `stdout`, `stderr`, and any FIDs associated with gnuplot.

**See also:** [fopen], page 311, [fflush], page 293, [freport], page 330.

*tf* = is_valid_file_id (*fid*)
: Return true if *fid* refers to an open file.

**See also:** [freport], page 330, [fopen], page 311.

### 14.2.2 Simple Output

Once a file has been opened for writing a string can be written to the file using the `fputs` function. The following example shows how to write the string '`Free Software is needed for Free Science`' to the file '`free.txt`'.

```
filename = "free.txt";
fid = fopen (filename, "w");
fputs (fid, "Free Software is needed for Free Science");
fclose (fid);
```

*status* = fputs (*fid*, *string*)

Write the string *string* to the file with file descriptor *fid*.

The string is written to the file with no additional formatting. Use `fdisp` instead to automatically append a newline character appropriate for the local machine.

The optional output *status* is 0 for success, or -1 if an error was encountered.

**See also:** [fdisp], page 300, [fprintf], page 315, [fwrite], page 327, [fopen], page 311.

A function much similar to `fputs` is available for writing data to the screen. The `puts` function works just like `fputs` except it doesn't take a file pointer as its input.

*status* = puts (*string*)

Write a string to the standard output with no formatting.

The string is written verbatim to the standard output. Use `disp` to automatically append a newline character appropriate for the local machine.

The optional output *status* is 0 for success, or -1 if an error was encountered.

**See also:** [fputs], page 313, [disp], page 287.

### 14.2.3 Line-Oriented Input

To read from a file it must be opened for reading using `fopen`. Then a line can be read from the file using `fgetl` as the following code illustrates

```
fid = fopen ("free.txt");
txt = fgetl (fid)
      ⊣ Free Software is needed for Free Science
fclose (fid);
```

This of course assumes that the file '`free.txt`' exists and contains the line '`Free Software is needed for Free Science`'.

*str* = fgetl (*fid*)
*str* = fgetl (*fid*, *len*)

Read characters from a file, stopping after a newline, or EOF, or *len* characters have been read.

The characters read, excluding the possible trailing newline, are returned as a string.

If *len* is omitted, `fgetl` reads until the next newline character.

If there are no more characters to read, `fgetl` returns $-1$.

To read a line and return the terminating newline, see [`fgets`], page 314.

**See also:** [fgets], page 314, [fscanf], page 321, [fread], page 325, [fopen], page 311.

*str* = fgets (*fid*)
*str* = fgets (*fid*, *len*)

Read characters from a file, stopping after a newline, or EOF, or *len* characters have been read.

The characters read, including the possible trailing newline, are returned as a string.

If *len* is omitted, `fgets` reads until the next newline character.

If there are no more characters to read, `fgets` returns $-1$.

To read a line and discard the terminating newline, see [fgetl], page 314.

**See also:** [fputs], page 313, [fgetl], page 314, [fscanf], page 321, [fread], page 325, [fopen], page 311.

*nlines* = fskipl (*fid*)
*nlines* = fskipl (*fid*, *count*)
*nlines* = fskipl (*fid*, *Inf*)

Read and skip *count* lines from the file specified by the file descriptor *fid*.

`fskipl` discards characters until an end-of-line is encountered exactly *count*-times, or until the end-of-file marker is found.

If *count* is omitted, it defaults to 1. *count* may also be `Inf`, in which case lines are skipped until the end of the file. This form is suitable for counting the number of lines in a file.

Returns the number of lines skipped (end-of-line sequences encountered).

**See also:** [fgetl], page 314, [fgets], page 314, [fscanf], page 321, [fopen], page 311.

### 14.2.4 Formatted Output

This section describes how to call `printf` and related functions.

The following functions are available for formatted output. They are modeled after the C language functions of the same name, but they interpret the format template differently in order to improve the performance of printing vector and matrix values.

Implementation Note: For compatibility with MATLAB, escape sequences in the template string (e.g., `"\n"` => newline) are expanded even when the template string is defined with single quotes.

printf (*template*, ...)
*numbytes* = printf (...)

Print optional arguments under the control of the template string *template* to the stream `stdout` and return the number of characters printed.

See the Formatted Output section of the GNU Octave manual for a complete description of the syntax of the template string.

The optional output *numbytes* returns the number of bytes printed.

Implementation Note: For compatibility with MATLAB, escape sequences in the template string (e.g., `"\n"` => newline) are expanded even when the template string is defined with single quotes.

**See also:** [fprintf], page 315, [sprintf], page 316, [scanf], page 321.

fprintf (*fid*, *template*, ...)
fprintf (*template*, ...)
*numbytes* = fprintf (...)

This function is equivalent to `printf`, except that the output is written to the file descriptor *fid* instead of `stdout`.

If *fid* is omitted, the output is written to `stdout` making the function exactly equivalent to `printf`.

The optional output *numbytes* returns the number of bytes written to the file.

Implementation Note: For compatibility with MATLAB, escape sequences in the template string (e.g., `"\n"` => newline) are expanded even when the template string is defined with single quotes.

**See also:** [fputs], page 313, [fdisp], page 300, [fwrite], page 327, [fscanf], page 321, [printf], page 315, [sprintf], page 316, [fopen], page 311.

*str* = `sprintf` (*template*, ...)

This is like `printf`, except that the output is returned as a string.

Unlike the C library function, which requires you to provide a suitably sized string as an argument, Octave's `sprintf` function returns the string, automatically sized to hold all of the items converted.

Implementation Note: For compatibility with MATLAB, escape sequences in the template string (e.g., `"\n"` => newline) are expanded even when the template string is defined with single quotes.

**See also:** [printf], page 315, [fprintf], page 315, [sscanf], page 322.

The `printf` function can be used to print any number of arguments. The template string argument you supply in a call provides information not only about the number of additional arguments, but also about their types and what style should be used for printing them.

Ordinary characters in the template string are simply written to the output stream as-is, while *conversion specifications* introduced by a '`%`' character in the template cause subsequent arguments to be formatted and written to the output stream. For example,

```
pct = 37;
filename = "foo.txt";
printf ("Processed %d%% of '%s'.\nPlease be patient.\n",
        pct, filename);
```

produces output like

```
Processed 37% of 'foo.txt'.
Please be patient.
```

This example shows the use of the '`%d`' conversion to specify that a scalar argument should be printed in decimal notation, the '`%s`' conversion to specify printing of a string argument, and the '`%%`' conversion to print a literal '`%`' character.

There are also conversions for printing an integer argument as an unsigned value in octal, decimal, or hexadecimal radix ('`%o`', '`%u`', or '`%x`', respectively); or as a character value ('`%c`').

Floating-point numbers can be printed in normal, fixed-point notation using the '`%f`' conversion or in exponential notation using the '`%e`' conversion. The '`%g`' conversion uses either '`%e`' or '`%f`' format, depending on what is more appropriate for the magnitude of the particular number.

You can control formatting more precisely by writing *modifiers* between the '`%`' and the character that indicates which conversion to apply. These slightly alter the ordinary behavior of the conversion. For example, most conversion specifications permit you to

specify a minimum field width and a flag indicating whether you want the result left- or right-justified within the field.

The specific flags and modifiers that are permitted and their interpretation vary depending on the particular conversion. They're all described in more detail in the following sections.

### 14.2.5 Output Conversion for Matrices

When given a matrix value, Octave's formatted output functions cycle through the format template until all the values in the matrix have been printed. For example:

```
printf ("%4.2f %10.2e %8.4g\n", hilb (3));

     ⊣ 1.00   5.00e-01   0.3333
     ⊣ 0.50   3.33e-01     0.25
     ⊣ 0.33   2.50e-01      0.2
```

If more than one value is to be printed in a single call, the output functions do not return to the beginning of the format template when moving on from one value to the next. This can lead to confusing output if the number of elements in the matrices are not exact multiples of the number of conversions in the format template. For example:

```
printf ("%4.2f %10.2e %8.4g\n", [1, 2], [3, 4]);

     ⊣ 1.00   2.00e+00        3
     ⊣ 4.00
```

If this is not what you want, use a series of calls instead of just one.

### 14.2.6 Output Conversion Syntax

This section provides details about the precise syntax of conversion specifications that can appear in a `printf` template string.

Characters in the template string that are not part of a conversion specification are printed as-is to the output stream.

The conversion specifications in a `printf` template string have the general form:

`%` *`flags width`* [ `.` *`precision`* ] *`type conversion`*

For example, in the conversion specifier '`%-10.8ld`', the '`-`' is a flag, '`10`' specifies the field width, the precision is '`8`', the letter '`l`' is a type modifier, and '`d`' specifies the conversion style. (This particular type specifier says to print a numeric argument in decimal notation, with a minimum of 8 digits left-justified in a field at least 10 characters wide.)

In more detail, output conversion specifications consist of an initial '`%`' character followed in sequence by:

- Zero or more *flag characters* that modify the normal behavior of the conversion specification.
- An optional decimal integer specifying the *minimum field width*. If the normal conversion produces fewer characters than this, the field is padded with spaces to the specified width. This is a *minimum* value; if the normal conversion produces more characters than this, the field is *not* truncated. Normally, the output is right-justified within the field.

You can also specify a field width of '`*`'. This means that the next argument in the argument list (before the actual value to be printed) is used as the field width. The value is rounded to the nearest integer. If the value is negative, this means to set the '`-`' flag (see below) and to use the absolute value as the field width.

- An optional *precision* to specify the number of digits to be written for the numeric conversions. If the precision is specified, it consists of a period ('`.`') followed optionally by a decimal integer (which defaults to zero if omitted).

  You can also specify a precision of '`*`'. This means that the next argument in the argument list (before the actual value to be printed) is used as the precision. The value must be an integer, and is ignored if it is negative.
- An optional *type modifier character*. This character is ignored by Octave's `printf` function, but is recognized to provide compatibility with the C language `printf`.
- A character that specifies the conversion to be applied.

The exact options that are permitted and how they are interpreted vary between the different conversion specifiers. See the descriptions of the individual conversions for information about the particular options that they use.

### 14.2.7 Table of Output Conversions

Here is a table summarizing what all the different conversions do:

'`%d`', '`%i`' Print an integer as a signed decimal number. See Section 14.2.8 [Integer Conversions], page 319, for details. '`%d`' and '`%i`' are synonymous for output, but are different when used with `scanf` for input (see Section 14.2.13 [Table of Input Conversions], page 323).

'`%o`' Print an integer as an unsigned octal number. See Section 14.2.8 [Integer Conversions], page 319, for details.

'`%u`' Print an integer as an unsigned decimal number. See Section 14.2.8 [Integer Conversions], page 319, for details.

'`%x`', '`%X`' Print an integer as an unsigned hexadecimal number. '`%x`' uses lowercase letters and '`%X`' uses uppercase. See Section 14.2.8 [Integer Conversions], page 319, for details.

'`%f`' Print a floating-point number in normal (fixed-point) notation. See Section 14.2.9 [Floating-Point Conversions], page 320, for details.

'`%e`', '`%E`' Print a floating-point number in exponential notation. '`%e`' uses lowercase letters and '`%E`' uses uppercase. See Section 14.2.9 [Floating-Point Conversions], page 320, for details.

'`%g`', '`%G`' Print a floating-point number in either normal (fixed-point) or exponential notation, whichever is more appropriate for its magnitude. '`%g`' uses lowercase letters and '`%G`' uses uppercase. See Section 14.2.9 [Floating-Point Conversions], page 320, for details.

'`%c`' Print a single character. See Section 14.2.10 [Other Output Conversions], page 320.

'`%s`' Print a string. See Section 14.2.10 [Other Output Conversions], page 320.

'%%' Print a literal '%' character. See Section 14.2.10 [Other Output Conversions], page 320.

If the syntax of a conversion specification is invalid, unpredictable things will happen, so don't do this. In particular, MATLAB allows a bare percentage sign '%' with no subsequent conversion character. Octave will emit an error and stop if it sees such code. When the string variable to be processed cannot be guaranteed to be free of potential format codes it is better to use the two argument form of any of the `printf` functions and set the format string to `%s`. Alternatively, for code which is not required to be backwards-compatible with MATLAB the Octave function `puts` or `disp` can be used.

```
printf (strvar);        # Unsafe if strvar contains format codes
printf ("%s", strvar);  # Safe
puts (strvar);          # Safe
```

If there aren't enough function arguments provided to supply values for all the conversion specifications in the template string, or if the arguments are not of the correct types, the results are unpredictable. If you supply more arguments than conversion specifications, the extra argument values are simply ignored; this is sometimes useful.

### 14.2.8 Integer Conversions

This section describes the options for the '%d', '%i', '%o', '%u', '%x', and '%X' conversion specifications. These conversions print integers in various formats.

The '%d' and '%i' conversion specifications both print an numeric argument as a signed decimal number; while '%o', '%u', and '%x' print the argument as an unsigned octal, decimal, or hexadecimal number (respectively). The '%X' conversion specification is just like '%x' except that it uses the characters 'ABCDEF' as digits instead of 'abcdef'.

The following flags are meaningful:

'-' Left-justify the result in the field (instead of the normal right-justification).

'+' For the signed '%d' and '%i' conversions, print a plus sign if the value is positive.

' ' For the signed '%d' and '%i' conversions, if the result doesn't start with a plus or minus sign, prefix it with a space character instead. Since the '+' flag ensures that the result includes a sign, this flag is ignored if you supply both of them.

'#' For the '%o' conversion, this forces the leading digit to be '0', as if by increasing the precision. For '%x' or '%X', this prefixes a leading '0x' or '0X' (respectively) to the result. This doesn't do anything useful for the '%d', '%i', or '%u' conversions.

'0' Pad the field with zeros instead of spaces. The zeros are placed after any indication of sign or base. This flag is ignored if the '-' flag is also specified, or if a precision is specified.

If a precision is supplied, it specifies the minimum number of digits to appear; leading zeros are produced if necessary. If you don't specify a precision, the number is printed with as many digits as it needs. If you convert a value of zero with an explicit precision of zero, then no characters at all are produced.

### 14.2.9 Floating-Point Conversions

This section discusses the conversion specifications for floating-point numbers: the '`%f`', '`%e`', '`%E`', '`%g`', and '`%G`' conversions.

The '`%f`' conversion prints its argument in fixed-point notation, producing output of the form [`-`]*ddd*`.`*ddd*, where the number of digits following the decimal point is controlled by the precision you specify.

The '`%e`' conversion prints its argument in exponential notation, producing output of the form [`-`]*d*`.`*ddd*`e`[`+|-`]*dd*. Again, the number of digits following the decimal point is controlled by the precision. The exponent always contains at least two digits. The '`%E`' conversion is similar but the exponent is marked with the letter '`E`' instead of '`e`'.

The '`%g`' and '`%G`' conversions print the argument in the style of '`%e`' or '`%E`' (respectively) if the exponent would be less than -4 or greater than or equal to the precision; otherwise they use the '`%f`' style. Trailing zeros are removed from the fractional portion of the result and a decimal-point character appears only if it is followed by a digit.

The following flags can be used to modify the behavior:

'`-`' Left-justify the result in the field. Normally the result is right-justified.

'`+`' Always include a plus or minus sign in the result.

' ' If the result doesn't start with a plus or minus sign, prefix it with a space instead. Since the '`+`' flag ensures that the result includes a sign, this flag is ignored if you supply both of them.

'`#`' Specifies that the result should always include a decimal point, even if no digits follow it. For the '`%g`' and '`%G`' conversions, this also forces trailing zeros after the decimal point to be left in place where they would otherwise be removed.

'`0`' Pad the field with zeros instead of spaces; the zeros are placed after any sign. This flag is ignored if the '`-`' flag is also specified.

The precision specifies how many digits follow the decimal-point character for the '`%f`', '`%e`', and '`%E`' conversions. For these conversions, the default precision is `6`. If the precision is explicitly `0`, this suppresses the decimal point character entirely. For the '`%g`' and '`%G`' conversions, the precision specifies how many significant digits to print. Significant digits are the first digit before the decimal point, and all the digits after it. If the precision is `0` or not specified for '`%g`' or '`%G`', it is treated like a value of `1`. If the value being printed cannot be expressed precisely in the specified number of digits, the value is rounded to the nearest number that fits.

### 14.2.10 Other Output Conversions

This section describes miscellaneous conversions for `printf`.

The '`%c`' conversion prints a single character. The '`-`' flag can be used to specify left-justification in the field, but no other flags are defined, and no precision or type modifier can be given. For example:

```
printf ("%c%c%c%c%c", "h", "e", "l", "l", "o");
```

prints '`hello`'.

The '`%s`' conversion prints a string. The corresponding argument must be a string. A precision can be specified to indicate the maximum number of characters to write; otherwise characters in the string up to but not including the terminating null character are written to the output stream. The '`-`' flag can be used to specify left-justification in the field, but no other flags or type modifiers are defined for this conversion. For example:

```
printf ("%3s%-6s", "no", "where");
```

prints '`nowhere`' (note the leading and trailing spaces).

### 14.2.11 Formatted Input

Octave provides the `scanf`, `fscanf`, and `sscanf` functions to read formatted input. There are two forms of each of these functions. One can be used to extract vectors of data from a file, and the other is more 'C-like'.

`[`*val*`,` *count*`,` *errmsg*`] = fscanf (`*fid*`,` *template*`,` *size*`)`
`[`*v1*`,` *v2*`, ...,` *count*`,` *errmsg*`] = fscanf (`*fid*`,` *template*`, "`*C*`")`

In the first form, read from *fid* according to *template*, returning the result in the matrix *val*.

The optional argument *size* specifies the amount of data to read and may be one of

`Inf` Read as much as possible, returning a column vector.

*nr* Read up to *nr* elements, returning a column vector.

`[`*nr*`, Inf]` Read as much as possible, returning a matrix with *nr* rows. If the number of elements read is not an exact multiple of *nr*, the last column is padded with zeros.

`[`*nr*`,` *nc*`]` Read up to *nr* `*` *nc* elements, returning a matrix with *nr* rows. If the number of elements read is not an exact multiple of *nr*, the last column is padded with zeros.

If *size* is omitted, a value of `Inf` is assumed.

A string is returned if *template* specifies only character conversions.

The number of items successfully read is returned in *count*.

If an error occurs, *errmsg* contains a system-dependent error message.

In the second form, read from *fid* according to *template*, with each conversion specifier in *template* corresponding to a single scalar return value. This form is more "C-like", and also compatible with previous versions of Octave. The number of successful conversions is returned in *count*

See the Formatted Input section of the GNU Octave manual for a complete description of the syntax of the template string.

**See also:** [fgets], page 314, [fgetl], page 314, [fread], page 325, [scanf], page 321, [sscanf], page 322, [fopen], page 311.

`[`*val*`,` *count*`,` *errmsg*`] = scanf (`*template*`,` *size*`)`
`[`*v1*`,` *v2*`, ...,` *count*`,` *errmsg*`] = scanf (`*template*`, "`*C*`")`

This is equivalent to calling `fscanf` with *fid* = `stdin`.

It is currently not useful to call `scanf` in interactive programs.

**See also:** [fscanf], page 321, [sscanf], page 322, [printf], page 315.

`[`*val*`,` *count*`,` *errmsg*`,` *pos*`] = sscanf (`*string*`,` *template*`,` *size*`)`
`[`*v1*`,` *v2*`, ..., ` *count*`,` *errmsg*`] = sscanf (`*string*`,` *template*`, "`*C*`")`

This is like `fscanf`, except that the characters are taken from the string *string* instead of from a stream.

Reaching the end of the string is treated as an end-of-file condition. In addition to the values returned by `fscanf`, the index of the next character to be read is returned in *pos*.

**See also:** [fscanf], page 321, [scanf], page 321, [sprintf], page 316.

Calls to `scanf` are superficially similar to calls to `printf` in that arbitrary arguments are read under the control of a template string. While the syntax of the conversion specifications in the template is very similar to that for `printf`, the interpretation of the template is oriented more towards free-format input and simple pattern matching, rather than fixed-field formatting. For example, most `scanf` conversions skip over any amount of "white space" (including spaces, tabs, and newlines) in the input file, and there is no concept of precision for the numeric input conversions as there is for the corresponding output conversions. Ordinarily, non-whitespace characters in the template are expected to match characters in the input stream exactly. For example, note that `sscanf` parses the string and whitespace differently when using mixed numeric and string output types:

```
teststr = "1 is a lonely number";
sscanf (teststr, "%s is a %s")
⇒ 1lonelynumber

sscanf (teststr, "%g is a %s")
⇒
     1
   108
   111
   110
   101
   108
   121

[a, b, c] = sscanf ("1 is a lonely number", "%g is a %s %s", "C")
⇒ a =  1
⇒ b = lonely
⇒ c = number
```

When a *matching failure* occurs, `scanf` returns immediately, leaving the first non-matching character as the next character to be read from the stream, and `scanf` returns all the items that were successfully converted.

The formatted input functions are not used as frequently as the formatted output functions. Partly, this is because it takes some care to use them properly. Another reason is that it is difficult to recover from a matching error.

The specific flags and modifiers that are permitted in the template string and their interpretation are all described in more detail in the following sections.

### 14.2.12 Input Conversion Syntax

A `scanf` template string is a string that contains ordinary multibyte characters interspersed with conversion specifications that start with '`%`'.

Any whitespace character in the template causes any number of whitespace characters in the input stream to be read and discarded. The whitespace characters that are matched need not be exactly the same whitespace characters that appear in the template string. For example, write ' `,` ' in the template to recognize a comma with optional whitespace before and after.

Other characters in the template string that are not part of conversion specifications must match characters in the input stream exactly; if this is not the case, a matching failure occurs.

The conversion specifications in a `scanf` template string have the general form:

```
% flags width type conversion
```

In more detail, an input conversion specification consists of an initial '`%`' character followed in sequence by:

- An optional *flag character* '`*`', which says to ignore the text read for this specification. When `scanf` finds a conversion specification that uses this flag, it reads input as directed by the rest of the conversion specification, but it discards this input, does not return any value, and does not increment the count of successful assignments.
- An optional decimal integer that specifies the *maximum field width*. Reading of characters from the input stream stops either when this maximum is reached or when a non-matching character is found, whichever happens first. Most conversions discard initial whitespace characters, and these discarded characters don't count towards the maximum field width. Conversions that do not discard initial whitespace are explicitly documented.
- An optional type modifier character. This character is ignored by Octave's `scanf` function, but is recognized to provide compatibility with the C language `scanf`.
- A character that specifies the conversion to be applied.

The exact options that are permitted and how they are interpreted vary between the different conversion specifiers. See the descriptions of the individual conversions in Section 14.2.13 [Table of Input Conversions], page 323, for information about the particular options that they allow.

### 14.2.13 Table of Input Conversions

Here is a table that summarizes the various conversion specifications:

'`%d`' Matches an optionally signed integer written in decimal. See Section 14.2.14 [Numeric Input Conversions], page 324.

'`%i`' Matches an optionally signed integer in any of the formats that the C language defines for specifying an integer constant. See Section 14.2.14 [Numeric Input Conversions], page 324.

'`%o`' Matches an unsigned integer written in octal radix. See Section 14.2.14 [Numeric Input Conversions], page 324.

'%u' Matches an unsigned integer written in decimal radix. See Section 14.2.14 [Numeric Input Conversions], page 324.

'%x', '%X' Matches an unsigned integer written in hexadecimal radix. See Section 14.2.14 [Numeric Input Conversions], page 324.

'%e', '%f', '%g', '%E', '%G'
Matches an optionally signed floating-point number. See Section 14.2.14 [Numeric Input Conversions], page 324.

'%s' Matches a string containing only non-whitespace characters. See Section 14.2.15 [String Input Conversions], page 324.

'%c' Matches a string of one or more characters; the number of characters read is controlled by the maximum field width given for the conversion. See Section 14.2.15 [String Input Conversions], page 324.

'%%' This matches a literal '%' character in the input stream. No corresponding argument is used.

If the syntax of a conversion specification is invalid, the behavior is undefined. If there aren't enough function arguments provided to supply addresses for all the conversion specifications in the template strings that perform assignments, or if the arguments are not of the correct types, the behavior is also undefined. On the other hand, extra arguments are simply ignored.

### 14.2.14 Numeric Input Conversions

This section describes the `scanf` conversions for reading numeric values.

The '%d' conversion matches an optionally signed integer in decimal radix.

The '%i' conversion matches an optionally signed integer in any of the formats that the C language defines for specifying an integer constant.

For example, any of the strings '10', '0xa', or '012' could be read in as integers under the '%i' conversion. Each of these specifies a number with decimal value 10.

The '%o', '%u', and '%x' conversions match unsigned integers in octal, decimal, and hexadecimal radices, respectively.

The '%X' conversion is identical to the '%x' conversion. They both permit either uppercase or lowercase letters to be used as digits.

By default, integers are read as 32-bit quantities. With the 'h' modifier, 16-bit integers are used, and with the 'l' modifier, 64-bit integers are used.

The '%e', '%f', '%g', '%E', and '%G' conversions match optionally signed floating-point numbers. All five conversion specifications behave identically, and will read in numerical values of any floating point display style.

### 14.2.15 String Input Conversions

This section describes the `scanf` input conversions for reading string and character values: '%s' and '%c'.

The '%c' conversion is the simplest: it matches a fixed number of characters, always. The maximum field with says how many characters to read; if you don't specify the maximum,

the default is 1. This conversion does not skip over initial whitespace characters. It reads precisely the next *n* characters, and fails if it cannot get that many.

The '`%s`' conversion matches a string of non-whitespace characters. It skips and discards initial whitespace, but stops when it encounters more whitespace after having read something.

For example, reading the input:

```
hello, world
```

with the conversion '`%10c`' produces `" hello, wo"`, but reading the same input with the conversion '`%10s`' produces `"hello,"`.

### 14.2.16 Binary I/O

Octave can read and write binary data using the functions `fread` and `fwrite`, which are patterned after the standard C functions with the same names. They are able to automatically swap the byte order of integer data and convert among the supported floating point formats as the data are read.

*val* = fread (*fid*)
*val* = fread (*fid*, *size*)
*val* = fread (*fid*, *size*, *precision*)
*val* = fread (*fid*, *size*, *precision*, *skip*)
*val* = fread (*fid*, *size*, *precision*, *skip*, *arch*)
[*val*, *count*] = fread (...)

Read binary data from the file specified by the file descriptor *fid*.

The optional argument *size* specifies the amount of data to read and may be one of

`Inf` Read as much as possible, returning a column vector.

*nr* Read up to *nr* elements, returning a column vector.

`[`*nr*`, Inf]` Read as much as possible, returning a matrix with *nr* rows. If the number of elements read is not an exact multiple of *nr*, the last column is padded with zeros.

`[`*nr*`,` *nc*`]` Read up to *nr* * *nc* elements, returning a matrix with *nr* rows. If the number of elements read is not an exact multiple of *nr*, the last column is padded with zeros.

If *size* is omitted, a value of `Inf` is assumed.

The optional argument *precision* is a string specifying the type of data to read and may be one of

`"uint8"` (default)
8-bit unsigned integer.

`"int8"`
`"integer*1"`
8-bit signed integer.

`"uint16"`
`"ushort"`
`"unsigned short"`
16-bit unsigned integer.

`"int16"`
`"integer*2"`
`"short"` 16-bit signed integer.

`"uint"`
`"uint32"`
`"unsigned int"`
`"ulong"`
`"unsigned long"`
32-bit unsigned integer.

`"int"`
`"int32"`
`"integer*4"`
`"long"` 32-bit signed integer.

`"uint64"` 64-bit unsigned integer.

`"int64"`
`"integer*8"`
64-bit signed integer.

`"single"`
`"float"`
`"float32"`
`"real*4"` 32-bit floating point number.

`"double"`
`"float64"`
`"real*8"` 64-bit floating point number.

`"char"`
`"char*1"` 8-bit single character.

`"uchar"`
`"unsigned char"`
8-bit unsigned character.

`"schar"`
`"signed char"`
8-bit signed character.

The default precision is `"uint8"`.

The *precision* argument may also specify an optional repeat count. For example, '`32*single`' causes `fread` to read a block of 32 single precision floating point numbers. Reading in blocks is useful in combination with the *skip* argument.

The *precision* argument may also specify a type conversion. For example, '`int16=>int32`' causes `fread` to read 16-bit integer values and return an array of 32-bit integer values. By default, `fread` returns a double precision array. The special form '`*TYPE`' is shorthand for '`TYPE=>TYPE`'.

The conversion and repeat counts may be combined. For example, the specification '`32*single=>single`' causes `fread` to read blocks of single precision floating point

values and return an array of single precision values instead of the default array of double precision values.

The optional argument *skip* specifies the number of bytes to skip after each element (or block of elements) is read. If it is not specified, a value of 0 is assumed. If the final block read is not complete, the final skip is omitted. For example,

```
fread (f, 10, "3*single=>single", 8)
```

will omit the final 8-byte skip because the last read will not be a complete block of 3 values.

The optional argument *arch* is a string specifying the data format for the file. Valid values are

`"native"` or `"n"`
: The format of the current machine.

`"ieee-be"` or `"b"`
: IEEE big endian.

`"ieee-le"` or `"l"`
: IEEE little endian.

If no *arch* is given the value used in the call to `fopen` which created the file descriptor is used. Otherwise, the value specified with `fread` overrides that of `fopen` and determines the data format.

The output argument *val* contains the data read from the file.

The optional return value *count* contains the number of elements read.

**See also:** [fwrite], page 327, [fgets], page 314, [fgetl], page 314, [fscanf], page 321, [fopen], page 311.

*count* = `fwrite` (*fid*, *data*)
*count* = `fwrite` (*fid*, *data*, *precision*)
*count* = `fwrite` (*fid*, *data*, *precision*, *skip*)
*count* = `fwrite` (*fid*, *data*, *precision*, *skip*, *arch*)

Write data in binary form to the file specified by the file descriptor *fid*.

The argument *data* is a matrix of values that are to be written to the file. The values are extracted in column-major order.

The remaining arguments *precision*, *skip*, and *arch* are optional, and are interpreted as described for `fread`.

The output *count* is the number of data items successfully written.

Programming Note: The behavior of `fwrite` is undefined if the values in *data* are too large to fit in the specified precision.

**See also:** [fread], page 325, [fputs], page 313, [fprintf], page 315, [fopen], page 311.

### 14.2.17 Temporary Files

Sometimes one needs to write data to a file that is only temporary. This is most commonly used when an external program launched from within Octave needs to access data. When Octave exits all temporary files will be deleted, so this step need not be executed manually.

`[fid, name, msg] = mkstemp ("template")`
`[fid, name, msg] = mkstemp ("template", delete)`

Return the file descriptor *fid* corresponding to a new temporary file with a unique name created from *template*.

The last six characters of *template* must be `"XXXXXX"` and these are replaced with a string that makes the filename unique. The file is then created with mode read/write and permissions that are system dependent (on GNU/Linux systems, the permissions will be 0600 for versions of glibc 2.0.7 and later). The file is opened in binary mode and with the `O_EXCL` flag.

If the optional argument *delete* is supplied and is true, the file will be deleted automatically when Octave exits.

If successful, *fid* is a valid file ID, *name* is the name of the file, and *msg* is an empty string. Otherwise, *fid* is -1, *name* is empty, and *msg* contains a system-dependent error message.

**See also:** [tempname], page 328, [tempdir], page 328, [P_tmpdir], page 329, [tmpfile], page 328, [fopen], page 311.

`[fid, msg] = tmpfile ()`

Return the file ID corresponding to a new temporary file with a unique name.

The file is opened in binary read/write (`"w+b"`) mode and will be deleted automatically when it is closed or when Octave exits.

If successful, *fid* is a valid file ID and *msg* is an empty string. Otherwise, *fid* is -1 and *msg* contains a system-dependent error message.

**See also:** [tempname], page 328, [mkstemp], page 327, [tempdir], page 328, [P_tmpdir], page 329.

`fname = tempname ()`
`fname = tempname (dir)`
`fname = tempname (dir, prefix)`

Return a unique temporary filename as a string.

If *prefix* is omitted, a value of `"oct-"` is used.

If *dir* is also omitted, the default directory for temporary files (`P_tmpdir`) is used. If *dir* is provided, it must exist, otherwise the default directory for temporary files is used.

Programming Note: Because the named file is not opened by `tempname`, it is possible, though relatively unlikely, that it will not be available by the time your program attempts to open it. If this is a concern, see [`tmpfile`], page 328.

**See also:** [mkstemp], page 327, [tempdir], page 328, [P_tmpdir], page 329, [tmpfile], page 328.

`dir = tempdir ()`

Return the name of the host system's directory for temporary files.

The directory name is taken first from the environment variable `TMPDIR`. If that does not exist, the environment variable `TMP` (and on Windows platforms also with higher

priority the environment variable `TEMP`) is checked. If none of those are set, the system default returned by `P_tmpdir` is used.

**See also:** [P_tmpdir], page 329, [tempname], page 328, [mkstemp], page 327, [tmpfile], page 328.

*sys_tmpdir* = `P_tmpdir ()`

Return the name of the host system's **default** directory for temporary files.

Programming Note: The value returned by `P_tmpdir` is always the default location. This value may not agree with that returned from `tempdir` if the user has overridden the default with the `TMPDIR` environment variable.

**See also:** [tempdir], page 328, [tempname], page 328, [mkstemp], page 327, [tmpfile], page 328.

### 14.2.18 EOF (End of File) and Errors

Once a file has been opened its status can be acquired. As an example the `feof` functions determines if the end of the file has been reached. This can be very useful when reading small parts of a file at a time. The following example shows how to read one line at a time from a file until the end has been reached.

```
filename = "myfile.txt";
fid = fopen (filename, "r");
while (! feof (fid) )
  text_line = fgetl (fid);
endwhile
fclose (fid);
```

Note that in some situations it is more efficient to read the entire contents of a file and then process it, than it is to read it line by line. This has the potential advantage of removing the loop in the above code.

*status* = `feof (`*fid*`)`

Return 1 if an end-of-file condition has been encountered for the file specified by file descriptor *fid* and 0 otherwise.

Note that `feof` will only return 1 if the end of the file has already been encountered, not if the next read operation will result in an end-of-file condition.

**See also:** [fread], page 325, [frewind], page 331, [fseek], page 330, [fclear], page 330, [fopen], page 311.

*msg* = `ferror (`*fid*`)`
[*msg*, *err*] = `ferror (`*fid*`)`
[...] = `ferror (`*fid*, "*clear*"`)`

Query the error status of the stream specified by file descriptor *fid*.

If an error condition exists then return a string *msg* describing the error. Otherwise, return an empty string `""`.

The second input `"clear"` is optional. If supplied, the error state on the stream will be cleared.

The optional second output is a numeric indication of the error status. *err* is 1 if an error condition has been encountered and 0 otherwise.

Note that `ferror` indicates if an error has already occurred, not whether the next operation will result in an error condition.

**See also:** [fclear], page 330, [fopen], page 311.

`fclear (`*fid*`)`

Clear the stream state for the file specified by the file descriptor *fid*.

**See also:** [ferror], page 329, [fopen], page 311.

`freport ()`

Print a list of which files have been opened, and whether they are open for reading, writing, or both.

For example:

```
freport ()
```

```
⊣  number  mode  arch       name
⊣  ------  ----  ----       ----
⊣     0     r    ieee-le    stdin
⊣     1     w    ieee-le    stdout
⊣     2     w    ieee-le    stderr
⊣     3     r    ieee-le    myfile
```

**See also:** [fopen], page 311, [fclose], page 313, [is_valid_file_id], page 313.

### 14.2.19 File Positioning

Three functions are available for setting and determining the position of the file pointer for a given file.

*pos* `= ftell (`*fid*`)`

Return the position of the file pointer as the number of characters from the beginning of the file specified by file descriptor *fid*.

**See also:** [fseek], page 330, [frewind], page 331, [feof], page 329, [fopen], page 311.

*status* `= fseek (`*fid*`, `*offset*`)`
*status* `= fseek (`*fid*`, `*offset*`, `*origin*`)`

Set the file pointer to the location *offset* within the file *fid*.

The pointer is positioned *offset* characters from the *origin*, which may be one of the predefined variables `SEEK_SET` (beginning), `SEEK_CUR` (current position), or `SEEK_END` (end of file) or strings `"bof"`, `"cof"`, or `"eof"`. If *origin* is omitted, `SEEK_SET` is assumed. *offset* may be positive, negative, or zero but not all combinations of *origin* and *offset* can be realized.

`fseek` returns 0 on success and -1 on error.

**See also:** [fskipl], page 315, [frewind], page 331, [ftell], page 330, [fopen], page 311, [SEEK_SET], page 331, [SEEK_CUR], page 331, [SEEK_END], page 331.

*fseek_origin* = SEEK_SET ()
: Return the numerical value to pass to `fseek` to position the file pointer relative to the beginning of the file.

  **See also:** [SEEK_CUR], page 331, [SEEK_END], page 331, [fseek], page 330.

*fseek_origin* = SEEK_CUR ()
: Return the numerical value to pass to `fseek` to position the file pointer relative to the current position.

  **See also:** [SEEK_SET], page 331, [SEEK_END], page 331, [fseek], page 330.

*fseek_origin* = SEEK_END ()
: Return the numerical value to pass to `fseek` to position the file pointer relative to the end of the file.

  **See also:** [SEEK_SET], page 331, [SEEK_CUR], page 331, [fseek], page 330.

frewind (*fid*)
*status* = frewind (*fid*)
: Move the file pointer to the beginning of the file specified by file descriptor *fid*.

  If an output *status* is requested then `frewind` returns 0 for success, and -1 if an error is encountered.

  Programming Note: `frewind` is equivalent to `fseek (fid, 0, SEEK_SET)`.

  **See also:** [fseek], page 330, [ftell], page 330, [fopen], page 311.

The following example stores the current file position in the variable `marker`, moves the pointer to the beginning of the file, reads four characters, and then returns to the original position.

```
marker = ftell (myfile);
frewind (myfile);
fourch = fgets (myfile, 4);
fseek (myfile, marker, SEEK_SET);
```

# 15 Plotting

## 15.1 Introduction to Plotting

Earlier versions of Octave provided plotting through the use of gnuplot. This capability is still available. But, newer versions of Octave offer more modern plotting capabilities using OpenGL. Which plotting system is used is controlled by the `graphics_toolkit` function. See Section 15.4.8 [Graphics Toolkits], page 561.

The function call `graphics_toolkit ("qt")` selects the Qt/OpenGL system, `graphics_toolkit ("fltk")` selects the FLTK/OpenGL system, and `graphics_toolkit ("gnuplot")` selects the gnuplot system. The three systems may be used selectively through the use of the `graphics_toolkit` property of the graphics handle for each figure. This is explained in Section 15.3 [Graphics Data Structures], page 448.

**Caution:** The OpenGL-based toolkits use single precision variables internally which limits the maximum value that can be displayed to approximately $10^{38}$. If your data contains larger values you must use the gnuplot toolkit which supports values up to $10^{308}$. Similarly, single precision variables can accurately represent only 6-9 base10 digits. If your data contains very fine differences (approximately 1e-8) these cannot be resolved with the OpenGL-based graphics toolkits and the gnuplot toolkit is required.

**Note:** The gnuplot graphics toolkit uses the third party program gnuplot for plotting. The communication from Octave to gnuplot is done via a one-way pipe. This has implications for performance and functionality. Performance is significantly slower because the entire data set, which could be many megabytes, must be passed to gnuplot over the pipe. Functionality is negatively affected because the pipe is one-way from Octave to gnuplot. Octave has no way of knowing about user interactions with the plot window (be it resizing, moving, closing, or anything else). It is recommended not to interact with (or close) a gnuplot window if you will access the figure from Octave later on.

## 15.2 High-Level Plotting

Octave provides simple means to create many different types of two- and three-dimensional plots using high-level functions.

If you need more detailed control, see Section 15.3 [Graphics Data Structures], page 448, and Section 15.4 [Advanced Plotting], page 545.

### 15.2.1 Two-Dimensional Plots

The `plot` function allows you to create simple x-y plots with linear axes. For example,

```
x = -10:0.1:10;
plot (x, sin (x));
xlabel ("x");
ylabel ("sin (x)");
title ("Simple 2-D Plot");
```

displays a sine wave shown in Figure 15.1. On most systems, this command will open a separate plot window to display the graph.

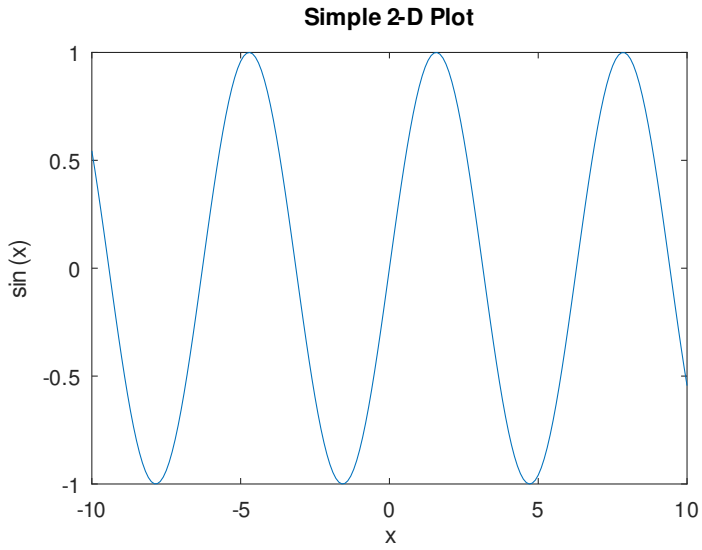



Figure 15.1: Simple Two-Dimensional Plot.

`plot (`*y*`)`
`plot (`*x*`, `*y*`)`
`plot (`*x*`, `*y*`, `*fmt*`)`
`plot (...,` *property*`,` *value*`, ...)`
`plot (`*x1*`, `*y1*`, ..., `*xn*`, `*yn*`)`
`plot (`*hax*`, ...)`
*h* `= plot (...)`

Produce 2-D plots.

Many different combinations of arguments are possible. The simplest form is

```
plot (y)
```

where the argument is taken as the set of $y$ coordinates and the $x$ coordinates are taken to be the range `1:numel (`*y*`)`.

If more than one argument is given, they are interpreted as

```
plot (y, property, value, ...)
```

or

```
plot (x, y, property, value, ...)
```

or

```
plot (x, y, fmt, ...)
```

and so on. Any number of argument sets may appear. The $x$ and $y$ values are interpreted as follows:

- If a single data argument is supplied, it is taken as the set of $y$ coordinates and the $x$ coordinates are taken to be the indices of the elements, starting with 1.
- If $x$ and $y$ are scalars, a single point is plotted.
- `squeeze()` is applied to arguments with more than two dimensions, but no more than two singleton dimensions.

- If either $x$ or $y$ is a scalar and the other is a vector, a series of points are plotted at the coordinates defined by the scalar and each element of the vector.
- If both arguments are vectors, the elements of $y$ are plotted versus the elements of $x$.
- If $x$ is a vector and $y$ is a matrix, then the columns (or rows) of $y$ are plotted versus $x$. (using whichever combination matches, with columns tried first.)
- If the $x$ is a matrix and $y$ is a vector, $y$ is plotted versus the columns (or rows) of $x$. (using whichever combination matches, with columns tried first.)
- If both arguments are matrices, the columns of $y$ are plotted versus the columns of $x$. In this case, both matrices must have the same number of rows and columns and no attempt is made to transpose the arguments to make the number of rows match.

Multiple property-value pairs may be specified, but they must appear in pairs. These arguments are applied to the line objects drawn by `plot`. Useful properties to modify are `"linestyle"`, `"linewidth"`, `"color"`, `"marker"`, `"markersize"`, `"markeredgecolor"`, `"markerfacecolor"`. The full list of properties is documented at Section 15.3.3.5 [Line Properties], page 485.

The *fmt* format argument can also be used to control the plot style. It is a string composed of four optional parts: `"<linestyle><marker><color><;displayname;>"`. When a marker is specified, but no linestyle, only the markers are plotted. Similarly, if a linestyle is specified, but no marker, then only lines are drawn. If both are specified then lines and markers will be plotted. If no *fmt* and no *property/value* pairs are given, then the default plot style is solid lines with no markers and the color determined by the `"colororder"` property of the current axes.

Format arguments:

linestyle

| | |
|---|---|
| '`-`' | Use solid lines (default). |
| '`--`' | Use dashed lines. |
| '`:`' | Use dotted lines. |
| '`-.`' | Use dash-dotted lines. |

marker

| | |
|---|---|
| '`+`' | crosshair |
| '`o`' | circle |
| '`*`' | star |
| '`.`' | point |
| '`x`' | cross |
| '`\|`' | vertical line |
| '`_`' | horizontal line |
| '`s`' | square |
| '`d`' | diamond |
| '`^`' | upward-facing triangle |
| '`v`' | downward-facing triangle |
| '`>`' | right-facing triangle |
| '`<`' | left-facing triangle |

| | |
|---|---|
| 'p' | pentagram |
| 'h' | hexagram |

color

| | |
|---|---|
| 'k', "black" | blacK |
| 'r', "red" | Red |
| 'g', "green" | Green |
| 'b', "blue" | Blue |
| 'y', "yellow" | Yellow |
| 'm', "magenta" | Magenta |
| 'c', "cyan" | Cyan |
| 'w', "white" | White |

";displayname;"
: The text between semicolons is used to set the "displayname" property which determines the label used for the plot legend.

The *fmt* argument may also be used to assign legend labels. To do so, include the desired label between semicolons after the formatting sequence described above, e.g., "+b;Data Series 3;". Note that the last semicolon is required and Octave will generate an error if it is left out.

Here are some plot examples:

```
plot (x, y, "or", x, y2, x, y3, "m", x, y4, "+")
```

This command will plot y with red circles, y2 with solid lines, y3 with solid magenta lines, and y4 with points displayed as '+'.

```
plot (b, "*", "markersize", 10)
```

This command will plot the data in the variable b, with points displayed as '*' and a marker size of 10.

```
t = 0:0.1:6.3;
plot (t, cos(t), "-;cos(t);", t, sin(t), "-b;sin(t);");
```

This will plot the cosine and sine functions and label them accordingly in the legend.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by gca.

The optional return value *h* is a vector of graphics handles to the created line objects.

To save a plot, in one of several image formats such as PostScript or PNG, use the print command.

**See also:** [axis], page 366, [box], page 420, [grid], page 420, [hold], page 429, [legend], page 415, [title], page 414, [xlabel], page 418, [ylabel], page 418, [xlim], page 369, [ylim], page 370, [ezplot], page 379, [errorbar], page 351, [fplot], page 378, [line], page 450, [plot3], page 399, [polar], page 355, [loglog], page 338, [semilogx], page 337, [semilogy], page 338, [subplot], page 423.

The plotyy function may be used to create a plot with two independent y axes.

plotyy (*x1*, *y1*, *x2*, *y2*)
plotyy (..., *fcn*)

`plotyy (...,` *fun1*, *fun2*`)`
`plotyy (`*hax*, `...)`
`[`*ax*, *h1*, *h2*`] = plotyy (...)`

Plot two sets of data with independent y-axes and a common x-axis.

The arguments *x1* and *y1* define the arguments for the first plot and *x1* and *y2* for the second.

By default the arguments are evaluated with `feval (@plot,` *x*, *y*`)`. However the type of plot can be modified with the *fcn* argument, in which case the plots are generated by `feval (`*fcn*, *x*, *y*`)`. *fcn* can be a function handle, an inline function, or a string of a function name.

The function to use for each of the plots can be independently defined with *fun1* and *fun2*.

The first argument *hax* can be an axes handle to the principal axes in which to plot the *x1* and *y1* data. It can also be a two-element vector with the axes handles to the primary and secondary axes (see output *ax*).

The return value *ax* is a vector with the axes handles of the two y-axes. *h1* and *h2* are handles to the objects generated by the plot commands.

```
x = 0:0.1:2*pi;
y1 = sin (x);
y2 = exp (x - 1);
ax = plotyy (x, y1, x - 1, y2, @plot, @semilogy);
xlabel ("X");
ylabel (ax(1), "Axis 1");
ylabel (ax(2), "Axis 2");
```

When using `plotyy` in conjunction with `subplot` make sure to call `subplot` first and pass the resulting axes handle to `plotyy`. Do not call `subplot` with any of the axes handles returned by `plotyy` or the other axes will be removed.

**See also:** [plot], page 334, [subplot], page 423.

The functions `semilogx`, `semilogy`, and `loglog` are similar to the `plot` function, but produce plots in which one or both of the axes use log scales.

`semilogx (`*y*`)`
`semilogx (`*x*, *y*`)`
`semilogx (`*x*, *y*, *property*, *value*, `...)`
`semilogx (`*x*, *y*, *fmt*`)`
`semilogx (`*hax*, `...)`
*h* `= semilogx (...)`

Produce a 2-D plot using a logarithmic scale for the x-axis.

See the documentation of `plot` for a description of the arguments that `semilogx` will accept.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created plot.

**See also:** [plot], page 334, [semilogy], page 338, [loglog], page 338.

`semilogy (`*y*`)`
`semilogy (`*x*`, `*y*`)`
`semilogy (`*x*`, `*y*`, `*property*`, `*value*`, ...)`
`semilogy (`*x*`, `*y*`, `*fmt*`)`
`semilogy (`*h*`, ...)`
*h* `= semilogy (...)`

Produce a 2-D plot using a logarithmic scale for the y-axis.

See the documentation of `plot` for a description of the arguments that `semilogy` will accept.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created plot.

**See also:** [plot], page 334, [semilogx], page 337, [loglog], page 338.

`loglog (`*y*`)`
`loglog (`*x*`, `*y*`)`
`loglog (`*x*`, `*y*`, `*prop*`, `*value*`, ...)`
`loglog (`*x*`, `*y*`, `*fmt*`)`
`loglog (`*hax*`, ...)`
*h* `= loglog (...)`

Produce a 2-D plot using logarithmic scales for both axes.

See the documentation of `plot` for a description of the arguments that `loglog` will accept.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created plot.

**See also:** [plot], page 334, [semilogx], page 337, [semilogy], page 338.

The functions `bar`, `barh`, `stairs`, and `stem` are useful for displaying discrete data. For example,

```
randn ("state", 1);
hist (randn (10000, 1), 30);
xlabel ("Value");
ylabel ("Count");
title ("Histogram of 10,000 normally distributed random numbers");
```

produces the histogram of 10,000 normally distributed random numbers shown in Figure 15.2. Note that, `randn ("state", 1);`, initializes the random number generator for `randn` to a known value so that the returned values are reproducible; This guarantees that the figure produced is identical to the one in this manual.

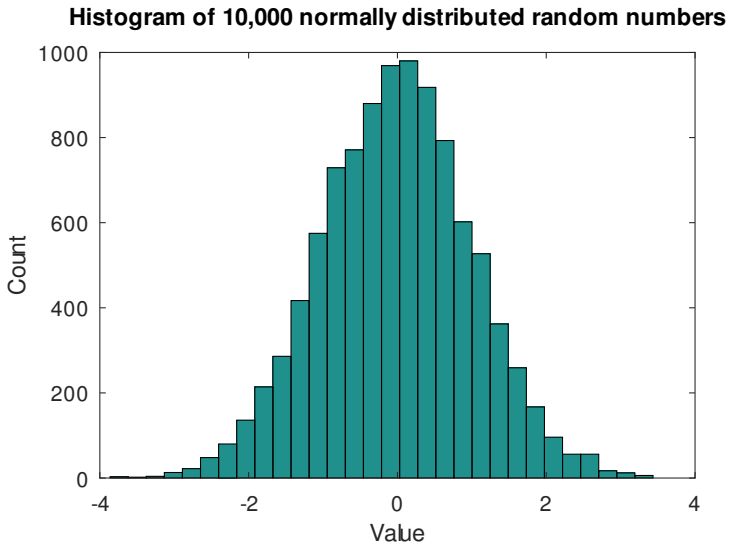



Figure 15.2: Histogram.

`bar` (*y*)
`bar` (*x*, *y*)
`bar` (. . . , *w*)
`bar` (. . . , *style*)
`bar` (. . . , *prop*, *val*, . . . )
`bar` (*hax*, . . . )
*h* = `bar` (. . . , *prop*, *val*, . . . )

Produce a bar graph from two vectors of X-Y data.

If only one argument is given, *y*, it is taken as a vector of Y values and the X coordinates are the range `1:numel (y)`.

The optional input *w* controls the width of the bars. A value of 1.0 will cause each bar to exactly touch any adjacent bars. The default width is 0.8.

If *y* is a matrix, then each column of *y* is taken to be a separate bar graph plotted on the same graph. By default the columns are plotted side-by-side. This behavior can be changed by the *style* argument which can take the following values:

`"grouped"` (default)
: Side-by-side bars with a gap between bars and centered over the X-coordinate.

`"stacked"`
: Bars are stacked so that each X value has a single bar composed of multiple segments.

`"hist"`
: Side-by-side bars with no gap between bars and centered over the X-coordinate.

`"histc"`
: Side-by-side bars with no gap between bars and left-aligned to the X-coordinate.

Optional property/value pairs are passed directly to the underlying patch objects. The full list of properties is documented at Section 15.3.3.8 [Patch Properties], page 495.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a vector of handles to the created `"bar series"` hggroups with one handle per column of the variable *y*. This series makes it possible to change a common element in one bar series object and have the change reflected in the other `"bar series"`. For example,

```
h = bar (rand (5, 10));
set (h(1), "basevalue", 0.5);
```

changes the position on the base of all of the bar series.

The following example modifies the face and edge colors using property/value pairs.

```
bar (randn (1, 100), "facecolor", "r", "edgecolor", "b");
```

The default color for bars is taken from the axes' `"ColorOrder"` property. The default color for bars when a histogram option (`"hist"`, `"histc"` is used is the `"Colormap"` property of either the axes or figure. The color of bars can also be set manually using the `"facecolor"` property as shown below.

```
h = bar (rand (10, 3));
set (h(1), "facecolor", "r")
set (h(2), "facecolor", "g")
set (h(3), "facecolor", "b")
```

**See also:** [barh], page 340, [hist], page 341, [pie], page 356, [plot], page 334, [patch], page 451.

`barh` (*y*)
`barh` (*x*, *y*)
`barh` (*...*, *w*)
`barh` (*...*, *style*)
`barh` (*...*, *prop*, *val*, *...*)
`barh` (*hax*, *...*)
*h* = `barh` (*...*, *prop*, *val*, *...*)

Produce a horizontal bar graph from two vectors of X-Y data.

If only one argument is given, it is taken as a vector of Y values and the X coordinates are the range `1:numel (y)`.

The optional input *w* controls the width of the bars. A value of 1.0 will cause each bar to exactly touch any adjacent bars. The default width is 0.8.

If *y* is a matrix, then each column of *y* is taken to be a separate bar graph plotted on the same graph. By default the columns are plotted side-by-side. This behavior can be changed by the *style* argument which can take the following values:

`"grouped"` (default)
: Side-by-side bars with a gap between bars and centered over the Y-coordinate.

`"stacked"`
: Bars are stacked so that each Y value has a single bar composed of multiple segments.

`"hist"`
: Side-by-side bars with no gap between bars and centered over the Y-coordinate.

`"histc"`
: Side-by-side bars with no gap between bars and left-aligned to the Y-coordinate.

Optional property/value pairs are passed directly to the underlying patch objects. The full list of properties is documented at Section 15.3.3.8 [Patch Properties], page 495.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created bar series hggroup. For a description of the use of the bar series, see [bar], page 339.

**See also:** [bar], page 339, [hist], page 341, [pie], page 356, [plot], page 334, [patch], page 451.

`hist (`*y*`)`
`hist (`*y*`, `*nbins*`)`
`hist (`*y*`, `*x*`)`
`hist (`*y*`, `*x*`, `*norm*`)`
`hist (...,` *prop*`, `*val*`, ...)`
`hist (`*hax*`, ...)`
`[`*nn*`, `*xx*`] = hist (...)`

Produce histogram counts or plots.

With one vector input argument, *y*, plot a histogram of the values with 10 bins. The range of the histogram bins is determined by the range of the data (difference between maximum and minimum value in *y*). Extreme values are lumped into the first and last bins. If *y* is a matrix then plot a histogram where each bin contains one bar per input column of *y*.

If the optional second argument is a scalar, *nbins*, it defines the number of bins.

If the optional second argument is a vector, *x*, it defines the centers of the bins. The width of the bins is determined from the adjacent values in the vector. The total number of bins is `numel (`*x*`)`.

If a third argument *norm* is provided, the histogram is normalized. In case *norm* is a positive scalar, the resulting bars are normalized to *norm*. If *norm* is a vector of positive scalars of length `columns (`*y*`)`, then the resulting bar of `y(:,i)` is normalized to *norm*`(i)`.

```
[nn, xx] = hist (rand (10, 3), 5, [1 2 3]);
sum (nn, 1)
⇒ ans =
      1   2   3
```

The histogram's appearance may be modified by specifying property/value pairs to the underlying patch object. For example, the face and edge color may be modified:

```
hist (randn (1, 100), 25, "facecolor", "r", "edgecolor", "b");
```

The full list of patch properties is documented at Section 15.3.3.8 [Patch Properties], page 495. property. If not specified, the default colors for the histogram are taken from the `"Colormap"` property of the axes or figure.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

If an output is requested then no plot is made. Instead, return the values *nn* (numbers of elements) and *xx* (bin centers) such that `bar (xx, nn)` will plot the histogram. If *y* is a vector, *nn* and *xx* will be row vectors. If *y* is an array, *nn* will be a 2-D array with one column of element counts for each column in *y*, and *xx* will be a column vector of bin centers.

**See also:** [histc], page 860, [bar], page 339, [pie], page 356, [rose], page 347.

`stemleaf (`*x*`, `*caption*`)`
`stemleaf (`*x*`, `*caption*`, `*stem_sz*`)`
*plotstr* `= stemleaf (...)`

Compute and display a stem and leaf plot of the vector *x*.

The input *x* should be a vector of integers. Any non-integer values will be converted to integer by `x = fix (x)`. By default each element of *x* will be plotted with the last digit of the element as a leaf value and the remaining digits as the stem. For example, 123 will be plotted with the stem '12' and the leaf '3'. The second argument, *caption*, should be a character array which provides a description of the data. It is included as a heading for the output.

The optional input *stem_sz* sets the width of each stem. The stem width is determined by `10^(stem_sz + 1)`. The default stem width is 10.

The output of `stemleaf` is composed of two parts: a "Fenced Letter Display", followed by the stem-and-leaf plot itself. The Fenced Letter Display is described in *Exploratory Data Analysis*. Briefly, the entries are as shown:

```
        Fenced Letter Display
#% nx|____________________     nx = numel (x)
M% mi|       md          |     mi median index, md median
H% hi|hl              hu| hs  hi lower hinge index, hl,hu hinges,
1    |x(1)         x(nx)|     hs h_spreadx(1), x(nx) first
           _______             and last data value.
     ______|step |_______      step 1.5*h_spread
    f|ifl            ifh|     inner fence, lower and higher
     |nfl            nfh|     no.\ of data points within fences
    F|ofl            ofh|     outer fence, lower and higher
     |nFl            nFh|     no.\ of data points outside outer
                               fences
```

The stem-and-leaf plot shows on each line the stem value followed by the string made up of the leaf digits. If the *stem_sz* is not 1 the successive leaf values are separated by ",".

With no return argument, the plot is immediately displayed. If an output argument is provided, the plot is returned as an array of strings.

The leaf digits are not sorted. If sorted leaf values are desired, use `xs = sort (x)` before calling `stemleaf (xs)`.

The stem and leaf plot and associated displays are described in: Chapter 3, *Exploratory Data Analysis* by J. W. Tukey, Addison-Wesley, 1977.

**See also:** [hist], page 341, [printd], page 343.

`printd (obj, filename)`
`out_file = printd (...)`

Convert any object acceptable to `disp` into the format selected by the suffix of *filename*.

If the optional output *out_file* is requested, the name of the created file is returned.

This function is intended to facilitate manipulation of the output of functions such as `stemleaf`.

**See also:** [stemleaf], page 342.

`stairs (y)`
`stairs (x, y)`
`stairs (..., style)`
`stairs (..., prop, val, ...)`
`stairs (hax, ...)`
`h = stairs (...)`
`[xstep, ystep] = stairs (...)`

Produce a stairstep plot.

The arguments *x* and *y* may be vectors or matrices. If only one argument is given, it is taken as a vector of Y values and the X coordinates are taken to be the indices of the elements (`x = 1:numel (y)`).

The style to use for the plot can be defined with a line style *style* of the same format as the `plot` command.

Multiple property/value pairs may be specified, but they must appear in pairs. The full list of properties is documented at Section 15.3.3.5 [Line Properties], page 485.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

If one output argument is requested, return a graphics handle to the created plot. If two output arguments are specified, the data are generated but not plotted. For example,

```
stairs (x, y);
```

and

```
[xs, ys] = stairs (x, y);
plot (xs, ys);
```

are equivalent.

**See also:** [bar], page 339, [hist], page 341, [plot], page 334, [stem], page 344.

`stem` (*y*)
`stem` (*x*, *y*)
`stem` (..., *linespec*)
`stem` (..., "*filled*")
`stem` (..., *prop*, *val*, ...)
`stem` (*hax*, ...)
*h* = `stem` (...)

Plot a 2-D stem graph.

If only one argument is given, it is taken as the y-values and the x-coordinates are taken from the indices of the elements.

If *y* is a matrix, then each column of the matrix is plotted as a separate stem graph. In this case *x* can either be a vector, the same length as the number of rows in *y*, or it can be a matrix of the same size as *y*.

The default color is `"b"` (blue), the default line style is `"-"`, and the default marker is `"o"`. The line style can be altered by the *linespec* argument in the same manner as the `plot` command. If the `"filled"` argument is present the markers at the top of the stems will be filled in. For example,

```
x = 1:10;
y = 2*x;
stem (x, y, "r");
```

plots 10 stems with heights from 2 to 20 in red;

Optional property/value pairs may be specified to control the appearance of the plot. The full list of properties is documented at Section 15.3.3.5 [Line Properties], page 485.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a handle to a "stem series" hggroup. The single hggroup handle has all of the graphical elements comprising the plot as its children; This allows the properties of multiple graphics objects to be changed by modifying just a single property of the "stem series" hggroup.

For example,

```
x = [0:10]';
y = [sin(x), cos(x)]
h = stem (x, y);
set (h(2), "color", "g");
set (h(1), "basevalue", -1)
```

changes the color of the second "stem series" and moves the base line of the first.

Stem Series Properties

| | |
|---|---|
| linestyle | The linestyle of the stem. (Default: `"-"`) |
| linewidth | The width of the stem. (Default: 0.5) |
| color | The color of the stem, and if not separately specified, the marker. (Default: `"b"` [blue]) |
| marker | The marker symbol to use at the top of each stem. (Default: `"o"`) |

markeredgecolor
: The edge color of the marker. (Default: `"color"` property)

markerfacecolor
: The color to use for "filling" the marker. (Default: `"none"` [unfilled])

markersize
: The size of the marker. (Default: 6)

baseline
: The handle of the line object which implements the baseline. Use `set` with the returned handle to change graphic properties of the baseline.

basevalue
: The y-value where the baseline is drawn. (Default: 0)

**See also:** [stem3], page 345, [bar], page 339, [hist], page 341, [plot], page 334, [stairs], page 343.

`stem3 (`*x*, *y*, *z*`)`
`stem3 (...,` *linespec*`)`
`stem3 (..., "`*filled*`")`
`stem3 (...,` *prop*, *val*, `...)`
`stem3 (`*hax*, `...)`
*h* `= stem3 (...)`

Plot a 3-D stem graph.

Stems are drawn from the height *z* to the location in the x-y plane determined by *x* and *y*. The default color is `"b"` (blue), the default line style is `"-"`, and the default marker is `"o"`.

The line style can be altered by the *linespec* argument in the same manner as the `plot` command. If the `"filled"` argument is present the markers at the top of the stems will be filled in.

Optional property/value pairs may be specified to control the appearance of the plot. The full list of properties is documented at Section 15.3.3.5 [Line Properties], page 485.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a handle to the "stem series" hggroup containing the line and marker objects used for the plot. See [stem], page 344, for a description of the "stem series" object.

Example:

```
theta = 0:0.2:6;
stem3 (cos (theta), sin (theta), theta);
```

plots 31 stems with heights from 0 to 6 lying on a circle.

Implementation Note: Color definitions with RGB-triples are not valid.

**See also:** [stem], page 344, [bar], page 339, [hist], page 341, [plot], page 334.

`scatter (`*x*, *y*`)`
`scatter (`*x*, *y*, *s*`)`
`scatter (`*x*, *y*, *s*, *c*`)`
`scatter (...,` *style*`)`

`scatter (...,` "*filled*")
`scatter (...,` *prop*, *val*, ...)
`scatter (`*hax*, ...)
*h* = `scatter (...)`

Draw a 2-D scatter plot.

A marker is plotted at each point defined by the coordinates in the vectors *x* and *y*.

The size of the markers is determined by *s*, which can be a scalar or a vector of the same length as *x* and *y*. If *s* is not given, or is an empty matrix, then a default value of 36 square points is used (The marker size itself is `sqrt (s)`).

The color of the markers is determined by *c*, which can be a string defining a fixed color; a 3-element vector giving the red, green, and blue components of the color; a vector of the same length as *x* that gives a scaled index into the current colormap; or an Nx3 matrix defining the RGB color of each marker individually.

The marker to use can be changed with the *style* argument; it is a string defining a marker in the same manner as the `plot` command. If no marker is specified it defaults to `"o"` or circles. If the argument `"filled"` is given then the markers are filled.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created scatter object.

Example:

```
x = randn (100, 1);
y = randn (100, 1);
scatter (x, y, [], sqrt (x.^2 + y.^2));
```

Programming Note: The full list of properties is documented at Section 15.3.3.9 [Scatter Properties], page 501.

**See also:** [scatter3], page 407, [patch], page 451, [plot], page 334.

`plotmatrix (`*x*, *y*`)`
`plotmatrix (`*x*`)`
`plotmatrix (...,` *style*`)`
`plotmatrix (`*hax*, ...`)`
`[`*h*, *ax*, *bigax*, *p*, *pax*`] = plotmatrix (...)`

Scatter plot of the columns of one matrix against another.

Given the arguments *x* and *y* that have a matching number of rows, `plotmatrix` plots a set of axes corresponding to

```
plot (x(:, i), y(:, j))
```

When called with a single argument *x* this is equivalent to

```
plotmatrix (x, x)
```

except that the diagonal of the set of axes will be replaced with the histogram `hist (x(:, i))`.

The marker to use can be changed with the *style* argument, that is a string defining a marker in the same manner as the `plot` command.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* provides handles to the individual graphics objects in the scatter plots, whereas *ax* returns the handles to the scatter plot axes objects.

*bigax* is a hidden axes object that surrounds the other axes, such that the commands `xlabel`, `title`, etc., will be associated with this hidden axes.

Finally, *p* returns the graphics objects associated with the histogram and *pax* the corresponding axes objects.

Example:

```
plotmatrix (randn (100, 3), "g+")
```

**See also:** [scatter], page 345, [plot], page 334.

`pareto (`*y*`)`
`pareto (`*y*`, `*x*`)`
`pareto (`*hax*`, ...)`
*h* `= pareto (...)`

Draw a Pareto chart.

A Pareto chart is a bar graph that arranges information in such a way that priorities for process improvement can be established; It organizes and displays information to show the relative importance of data. The chart is similar to the histogram or bar chart, except that the bars are arranged in decreasing magnitude from left to right along the x-axis.

The fundamental idea (Pareto principle) behind the use of Pareto diagrams is that the majority of an effect is due to a small subset of the causes. For quality improvement, the first few contributing causes (leftmost bars as presented on the diagram) to a problem usually account for the majority of the result. Thus, targeting these "major causes" for elimination results in the most cost-effective improvement scheme.

Typically only the magnitude data *y* is present in which case *x* is taken to be the range `1 : length (`*y*`)`. If *x* is given it may be a string array, a cell array of strings, or a numerical vector.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a 2-element vector with a graphics handle for the created bar plot and a second handle for the created line plot.

An example of the use of `pareto` is

```
Cheese = {"Cheddar", "Swiss", "Camembert", ...
          "Munster", "Stilton", "Blue"};
Sold = [105, 30, 70, 10, 15, 20];
pareto (Sold, Cheese);
```

**See also:** [bar], page 339, [barh], page 340, [hist], page 341, [pie], page 356, [plot], page 334.

`rose (`*theta*`)`
`rose (`*theta*`, `*nbins*`)`
`rose (`*theta*`, `*bins*`)`
`rose (`*hax*`, ...)`

`h = rose (...)`
`[th, r] = rose (...)`

Plot an angular histogram.

With one vector argument, *th*, plot the histogram with 20 angular bins. If *th* is a matrix then each column of *th* produces a separate histogram.

If *nbins* is given and is a scalar, then the histogram is produced with *nbin* bins. If *bins* is a vector, then the center of each bin is defined by the values in *bins* and the number of bins is given by the number of elements in *bins*.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a vector of graphics handles to the line objects representing each histogram.

If two output arguments are requested then no plot is made and the polar vectors necessary to plot the histogram are returned instead.

Example

```
[th, r] = rose ([2*randn(1e5,1), pi + 2*randn(1e5,1)]);
polar (th, r);
```

**See also:** [hist], page 341, [polar], page 355.

The `contour`, `contourf` and `contourc` functions produce two-dimensional contour plots from three-dimensional data.

`contour (z)`
`contour (z, vn)`
`contour (x, y, z)`
`contour (x, y, z, vn)`
`contour (..., style)`
`contour (hax, ...)`
`[c, h] = contour (...)`

Create a 2-D contour plot.

Plot level curves (contour lines) of the matrix *z*, using the contour matrix *c* computed by `contourc` from the same arguments; see the latter for their interpretation.

The appearance of contour lines can be defined with a line style *style* in the same manner as `plot`. Only line style and color are used; Any markers defined by *style* are ignored.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional output *c* contains the contour levels in `contourc` format.

The optional return value *h* is a graphics handle to the hggroup comprising the contour lines.

Example:

```
x = 0:3;
y = 0:2;
z = y' * x;
contour (x, y, z, 2:3)
```

**See also:** [ezcontour], page 380, [contourc], page 349, [contourf], page 349, [contour3], page 350, [clabel], page 419, [meshc], page 384, [surfc], page 386, [clim], page 368, [colormap], page 951, [plot], page 334.

`contourf (z)`
`contourf (z, vn)`
`contourf (x, y, z)`
`contourf (x, y, z, vn)`
`contourf (..., style)`
`contourf (hax, ...)`
`[c, h] = contourf (...)`

Create a 2-D contour plot with filled intervals.

Plot level curves (contour lines) of the matrix *z* and fill the region between lines with colors from the current colormap.

The level curves are taken from the contour matrix *c* computed by `contourc` for the same arguments; see the latter for their interpretation.

The appearance of contour lines can be defined with a line style *style* in the same manner as `plot`. Only line style and color are used; Any markers defined by *style* are ignored.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional output *c* contains the contour levels in `contourc` format.

The optional return value *h* is a graphics handle to the hggroup comprising the contour lines.

The following example plots filled contours of the `peaks` function.

```
[x, y, z] = peaks (50);
contourf (x, y, z, -7:9)
```

**See also:** [ezcontourf], page 380, [contour], page 348, [contourc], page 349, [contour3], page 350, [clabel], page 419, [meshc], page 384, [surfc], page 386, [clim], page 368, [colormap], page 951, [plot], page 334.

`c = contourc (z)`
`c = contourc (z, vn)`
`c = contourc (x, y, z)`
`c = contourc (x, y, z, vn)`
`[c, lev] = contourc (...)`

Compute contour lines (isolines of constant Z value).

The matrix *z* contains height values above the rectangular grid determined by *x* and *y*. If only a single input *z* is provided then *x* is taken to be `1:columns (z)` and *y* is taken to be `1:rows (z)`. The minimum data size is 2x2.

The optional input *vn* is either a scalar denoting the number of contour lines to compute or a vector containing the Z values where lines will be computed. When *vn* is a vector the number of contour lines is `numel (vn)`. However, to compute a single contour line at a given value use `vn = [val, val]`. If *vn* is omitted it defaults to 10.

The return value *c* is a 2x*n* matrix containing the contour lines in the following format

```
c = [lev1, x1, x2, ..., levn, x1, x2, ...
     len1, y1, y2, ..., lenn, y1, y2, ...]
```

in which contour line *n* has a level (height) of *levn* and length of *lenn*.

The optional return value *lev* is a vector with the Z values of the contour levels.

Example:

```
x = 0:2;
y = x;
z = x' * y;
c = contourc (x, y, z, 2:3)
⇒ c =
        2.0000   1.0000   1.0000   2.0000   2.0000   3.0000   1.5000   2.0000
        4.0000   2.0000   2.0000   1.0000   1.0000   2.0000   2.0000   1.5000
```

**See also:** [contour], page 348, [contourf], page 349, [contour3], page 350, [clabel], page 419.

`contour3 (`*z*`)`
`contour3 (`*z*`, `*vn*`)`
`contour3 (`*x*`, `*y*`, `*z*`)`
`contour3 (`*x*`, `*y*`, `*z*`, `*vn*`)`
`contour3 (...,` *style*`)`
`contour3 (`*hax*`, ...)`
`[`*c*`, `*h*`] = contour3 (...)`

Create a 3-D contour plot.

`contour3` plots level curves (contour lines) of the matrix *z* at a Z level corresponding to each contour. This is in contrast to `contour` which plots all of the contour lines at the same Z level and produces a 2-D plot.

The level curves are taken from the contour matrix *c* computed by `contourc` for the same arguments; see the latter for their interpretation.

The appearance of contour lines can be defined with a line style *style* in the same manner as `plot`. Only line style and color are used; Any markers defined by *style* are ignored.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional output *c* are the contour levels in `contourc` format.

The optional return value *h* is a graphics handle to the hggroup comprising the contour lines.

Example:

```
contour3 (peaks (19));
colormap cool;
hold on;
surf (peaks (19), "facecolor", "none", "edgecolor", "black");
```

**See also:** [contour], page 348, [contourc], page 349, [contourf], page 349, [clabel], page 419, [meshc], page 384, [surfc], page 386, [clim], page 368, [colormap], page 951, [plot], page 334.

The `errorbar`, `semilogxerr`, `semilogyerr`, and `loglogerr` functions produce plots with error bar markers. For example,

```
rand ("state", 2);
x = 0:0.1:10;
y = sin (x);
lerr = 0.1 .* rand (size (x));
uerr = 0.1 .* rand (size (x));
errorbar (x, y, lerr, uerr);
axis ([0, 10, -1.1, 1.1]);
xlabel ("x");
ylabel ("sin (x)");
title ("Errorbar plot of sin (x)");
```

produces the figure shown in Figure 15.3.

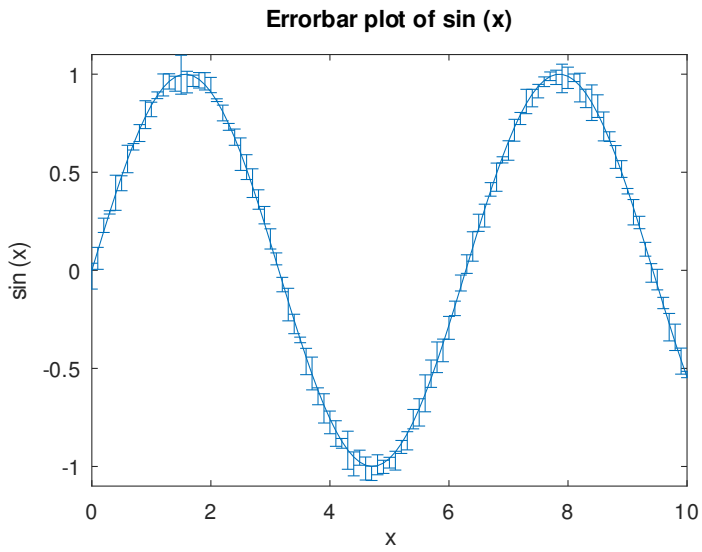



Figure 15.3: Errorbar plot.

`errorbar` (*y*, *ey*)
`errorbar` (*y*, ..., *fmt*)
`errorbar` (*x*, *y*, *ey*)
`errorbar` (*x*, *y*, *err*, *fmt*)
`errorbar` (*x*, *y*, *lerr*, *uerr*, *fmt*)
`errorbar` (*x*, *y*, *ex*, *ey*, *fmt*)
`errorbar` (*x*, *y*, *lx*, *ux*, *ly*, *uy*, *fmt*)
`errorbar` (*x1*, *y1*, ..., *fmt*, *xn*, *yn*, ...)
`errorbar` (*hax*, ...)
*h* = `errorbar` (...)

Create a 2-D plot with errorbars.

Many different combinations of arguments are possible. The simplest form is

```
errorbar (y, ey)
```

where the first argument is taken as the set of *y* coordinates, the second argument *ey* are the errors around the *y* values, and the *x* coordinates are taken to be the indices of the elements (`1:numel (y)`).

The general form of the function is

```
errorbar (x, y, err1, ..., fmt, ...)
```

After the *x* and *y* arguments there can be 1, 2, or 4 parameters specifying the error values depending on the nature of the error values and the plot format *fmt*.

*err* (scalar)
: When the error is a scalar all points share the same error value. The errorbars are symmetric and are drawn from *data*-*err* to *data*+*err*. The *fmt* argument determines whether *err* is in the x-direction, y-direction (default), or both.

*err* (vector or matrix)
: Each data point has a particular error value. The errorbars are symmetric and are drawn from *data*(n)-*err*(n) to *data*(n)+*err*(n).

*lerr*, *uerr* (scalar)
: The errors have a single low-side value and a single upper-side value. The errorbars are not symmetric and are drawn from *data*-*lerr* to *data*+*uerr*.

*lerr*, *uerr* (vector or matrix)
: Each data point has a low-side error and an upper-side error. The errorbars are not symmetric and are drawn from *data*(n)-*lerr*(n) to *data*(n)+*uerr*(n).

Any number of data sets (*x1*,*y1*, *x2*,*y2*, . . . ) may appear as long as they are separated by a format string *fmt*.

If *y* is a matrix, *x* and the error parameters must also be matrices having the same dimensions. The columns of *y* are plotted versus the corresponding columns of *x* and errorbars are taken from the corresponding columns of the error parameters.

If *fmt* is missing, the yerrorbars (`"~"`) plot style is assumed.

If the *fmt* argument is supplied then it is interpreted, as in normal plots, to specify the line style, marker, and color. In addition, *fmt* may include an errorbar style which **must precede** the ordinary format codes. The following errorbar styles are supported:

‘`~`’ Set yerrorbars plot style (default).

‘`>`’ Set xerrorbars plot style.

‘`~>`’ Set xyerrorbars plot style.

‘`#~`’ Set yboxes plot style.

‘`#`’ Set xboxes plot style.

‘`#~>`’ Set xyboxes plot style.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a handle to the hggroup object representing the data plot and errorbars.

Note: For compatibility with MATLAB a line is drawn through all data points. However, most scientific errorbar plots are a scatter plot of points with errorbars. To accomplish this, add a marker style to the *fmt* argument such as `"."`. Alternatively, remove the line by modifying the returned graphic handle with `set (h, "linestyle", "none")`.

Examples:

```
errorbar (x, y, ex, ">.r")
```

produces an xerrorbar plot of *y* versus *x* with *x* errorbars drawn from *x*-*ex* to *x*+*ex*. The marker `"."` is used so no connecting line is drawn and the errorbars appear in red.

```
errorbar (x, y1, ey, "~",
          x, y2, ly, uy)
```

produces yerrorbar plots with *y1* and *y2* versus *x*. Errorbars for *y1* are drawn from *y1*-*ey* to *y1*+*ey*, errorbars for *y2* from *y2*-*ly* to *y2*+*uy*.

```
errorbar (x, y, lx, ux,
          ly, uy, "~>")
```

produces an xyerrorbar plot of *y* versus *x* in which *x* errorbars are drawn from *x*-*lx* to *x*+*ux* and *y* errorbars from *y*-*ly* to *y*+*uy*.

**See also:** [semilogxerr], page 353, [semilogyerr], page 354, [loglogerr], page 354, [plot], page 334.

`semilogxerr (`*y*, *ey*`)`
`semilogxerr (`*y*, *...*, *fmt*`)`
`semilogxerr (`*x*, *y*, *ey*`)`
`semilogxerr (`*x*, *y*, *err*, *fmt*`)`
`semilogxerr (`*x*, *y*, *lerr*, *uerr*, *fmt*`)`
`semilogxerr (`*x*, *y*, *ex*, *ey*, *fmt*`)`
`semilogxerr (`*x*, *y*, *lx*, *ux*, *ly*, *uy*, *fmt*`)`
`semilogxerr (`*x1*, *y1*, *...*, *fmt*, *xn*, *yn*, *...*`)`
`semilogxerr (`*hax*, *...*`)`
*h* `= semilogxerr (...)`

Produce 2-D plots using a logarithmic scale for the x-axis and errorbars at each data point.

Many different combinations of arguments are possible. The most common form is

```
semilogxerr (x, y, ey, fmt)
```

which produces a semi-logarithmic plot of *y* versus *x* with errors in the *y*-scale defined by *ey* and the plot format defined by *fmt*. See [errorbar], page 351, for available formats and additional information.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

**See also:** [errorbar], page 351, [semilogyerr], page 354, [loglogerr], page 354.

`semilogyerr (`*y*, *ey*`)`
`semilogyerr (`*y*, ..., *fmt*`)`
`semilogyerr (`*x*, *y*, *ey*`)`
`semilogyerr (`*x*, *y*, *err*, *fmt*`)`
`semilogyerr (`*x*, *y*, *lerr*, *uerr*, *fmt*`)`
`semilogyerr (`*x*, *y*, *ex*, *ey*, *fmt*`)`
`semilogyerr (`*x*, *y*, *lx*, *ux*, *ly*, *uy*, *fmt*`)`
`semilogyerr (`*x1*, *y1*, ..., *fmt*, *xn*, *yn*, ...`)`
`semilogyerr (`*hax*, ...`)`
*h* `= semilogyerr (...)`

Produce 2-D plots using a logarithmic scale for the y-axis and errorbars at each data point.

Many different combinations of arguments are possible. The most common form is

```
semilogyerr (x, y, ey, fmt)
```

which produces a semi-logarithmic plot of *y* versus *x* with errors in the *y*-scale defined by *ey* and the plot format defined by *fmt*. See [errorbar], page 351, for available formats and additional information.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

**See also:** [errorbar], page 351, [semilogxerr], page 353, [loglogerr], page 354.

`loglogerr (`*y*, *ey*`)`
`loglogerr (`*y*, ..., *fmt*`)`
`loglogerr (`*x*, *y*, *ey*`)`
`loglogerr (`*x*, *y*, *err*, *fmt*`)`
`loglogerr (`*x*, *y*, *lerr*, *uerr*, *fmt*`)`
`loglogerr (`*x*, *y*, *ex*, *ey*, *fmt*`)`
`loglogerr (`*x*, *y*, *lx*, *ux*, *ly*, *uy*, *fmt*`)`
`loglogerr (`*x1*, *y1*, ..., *fmt*, *xn*, *yn*, ...`)`
`loglogerr (`*hax*, ...`)`
*h* `= loglogerr (...)`

Produce 2-D plots on a double logarithm axis with errorbars.

Many different combinations of arguments are possible. The most common form is

```
loglogerr (x, y, ey, fmt)
```

which produces a double logarithm plot of *y* versus *x* with errors in the *y*-scale defined by *ey* and the plot format defined by *fmt*. See [errorbar], page 351, for available formats and additional information.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

**See also:** [errorbar], page 351, [semilogxerr], page 353, [semilogyerr], page 354.

Finally, the `polar` function allows you to easily plot data in polar coordinates. However, the display coordinates remain rectangular and linear. For example,

```
polar (0:0.1:10*pi, 0:0.1:10*pi);
title ("Example polar plot from 0 to 10*pi");
```

produces the spiral plot shown in Figure 15.4.

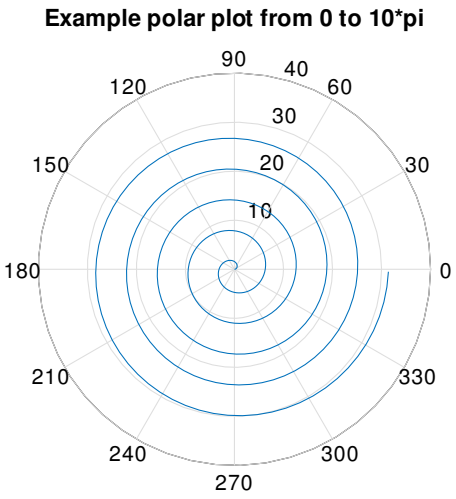



Figure 15.4: Polar plot.

`polar (`*`theta`*`, `*`rho`*`)`
`polar (`*`theta`*`, `*`rho`*`, `*`fmt`*`)`
`polar (`*`cplx`*`)`
`polar (`*`cplx`*`, `*`fmt`*`)`
`polar (`*`hax`*`, ...)`
*`h`* `= polar (...)`

Create a 2-D plot from polar coordinates *theta* and *rho*.

The input *theta* is assumed to be radians and is converted to degrees for plotting. If you have degrees then you must convert (see [cart2pol], page 641) to radians before passing the data to this function.

If a single complex input *cplx* is given then the real part is used for *theta* and the imaginary part is used for *rho*.

The optional argument *fmt* specifies the line format in the same way as `plot`.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created plot.

Implementation Note: The polar axis is drawn using line and text objects encapsulated in an hggroup. The hggroup properties are linked to the original axes object such that altering an appearance property, for example `fontname`, will update the polar axis. Two new properties are added to the original axes–`rtick`, `ttick`–which replace `xtick`, `ytick`. The first is a list of tick locations in the radial (rho) direction; The second is a list of tick locations in the angular (theta) direction specified in degrees, i.e., in the range 0–359.

**See also:** [rose], page 347, [compass], page 362, [plot], page 334, [cart2pol], page 641.

`pie (`*x*`)`
`pie (...,` *explode*`)`
`pie (...,` *labels*`)`
`pie (`*hax*`, ...)`
*h* `= pie (...)`

Plot a 2-D pie chart.

When called with a single vector argument, produce a pie chart of the elements in *x*. The size of the ith slice is the percentage that the element *x*i represents of the total sum of *x*: `pct = x(i) / sum (x)`.

The optional input *explode* is a vector of the same length as *x* that, if nonzero, "explodes" the slice from the pie chart.

The optional input *labels* is a cell array of strings of the same length as *x* specifying the label for each slice.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a list of handles to the patch and text objects generating the plot.

Note: If `sum (x)` $\leq$ `1` then the elements of *x* are interpreted as percentages directly and are not normalized by `sum (x)`. Furthermore, if the sum is less than 1 then there will be a missing slice in the pie plot to represent the missing, unspecified percentage.

**See also:** [pie3], page 356, [bar], page 339, [hist], page 341, [rose], page 347.

`pie3 (`*x*`)`
`pie3 (...,` *explode*`)`
`pie3 (...,` *labels*`)`
`pie3 (`*hax*`, ...)`
*h* `= pie3 (...)`

Plot a 3-D pie chart.

Called with a single vector argument, produces a 3-D pie chart of the elements in *x*. The size of the ith slice is the percentage that the element *x*i represents of the total sum of *x*: `pct = x(i) / sum (x)`.

The optional input *explode* is a vector of the same length as *x* that, if nonzero, "explodes" the slice from the pie chart.

The optional input *labels* is a cell array of strings of the same length as *x* specifying the label for each slice.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a list of graphics handles to the patch, surface, and text objects generating the plot.

Note: If `sum (x)` $\leq$ `1` then the elements of *x* are interpreted as percentages directly and are not normalized by `sum (x)`. Furthermore, if the sum is less than 1 then there will be a missing slice in the pie plot to represent the missing, unspecified percentage.

**See also:** [pie], page 356, [bar], page 339, [hist], page 341, [rose], page 347.

`quiver (u, v)`
`quiver (x, y, u, v)`
`quiver (..., s)`
`quiver (..., style)`
`quiver (..., "filled")`
`quiver (hax, ...)`
`h = quiver (...)`

Plot a 2-D vector field with arrows.

Plot the ($u$, $v$) components of a vector field at the grid points defined by ($x$, $y$). If the grid is uniform then $x$ and $y$ can be specified as grid vectors and `meshgrid` is used to create the 2-D grid.

If $x$ and $y$ are not given they are assumed to be `(1:m, 1:n)` where `[m, n] = size (u)`.

The optional input $s$ is a scalar defining a scaling factor to use for the arrows of the field relative to the mesh spacing. A value of 1.0 will result in the longest vector exactly filling one grid square. A value of 0 or `"off"` disables all scaling. The default value is 0.9.

The style to use for the plot can be defined with a line style, *style*, of the same format as the `plot` command. If a marker is specified then the markers are drawn at the origin of the vectors (which are the grid points defined by $x$ and $y$). When a marker is specified, the arrowhead is not drawn. If the argument `"filled"` is given then the markers are filled. If name-value plot style properties are used, they must appear in pairs and follow any other plot style arguments.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value $h$ is a graphics handle to a quiver object. A quiver object regroups the components of the quiver plot (body, arrow, and marker), and allows them to be changed together.

Example:

```
[x, y] = meshgrid (1:2:20);
h = quiver (x, y, sin (2*pi*x/10), sin (2*pi*y/10));
set (h, "maxheadsize", 0.33);
```

**See also:** [quiver3], page 357, [compass], page 362, [feather], page 362, [plot], page 334.

`quiver3 (x, y, z, u, v, w)`
`quiver3 (z, u, v, w)`
`quiver3 (..., s)`
`quiver3 (..., style)`
`quiver3 (..., "filled")`
`quiver3 (hax, ...)`
`h = quiver3 (...)`

Plot a 3-D vector field with arrows.

Plot the ($u$, $v$, $w$) components of a vector field at the grid points defined by ($x$, $y$, $z$). If the grid is uniform then $x$, $y$, and $z$ can be specified as grid vectors and `meshgrid` is used to create the 3-D grid.

If *x* and *y* are not given they are assumed to be `(1:m, 1:n)` where `[m, n] = size (u)`.

The optional input *s* is a scalar defining a scaling factor to use for the arrows of the field relative to the mesh spacing. A value of 1.0 will result in the longest vector exactly filling one grid cube. A value of 0 or `"off"` disables all scaling. The default value is 0.9.

The style to use for the plot can be defined with a line style *style* of the same format as the `plot` command. If a marker is specified then the markers are drawn at the origin of the vectors (which are the grid points defined by *x*, *y*, *z*). When a marker is specified, the arrowhead is not drawn. If the argument `"filled"` is given then the markers are filled. If name-value plot style properties are used, they must appear in pairs and follow any other plot style arguments.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to a quiver object. A quiver object regroups the components of the quiver plot (body, arrow, and marker), and allows them to be changed together.

```
[x, y, z] = peaks (25);
surf (x, y, z);
hold on;
[u, v, w] = surfnorm (x, y, z / 10);
h = quiver3 (x, y, z, u, v, w);
set (h, "maxheadsize", 0.33);
```

**See also:** [quiver], page 357, [compass], page 362, [feather], page 362, [plot], page 334.

`streamribbon (`*x*, *y*, *z*, *u*, *v*, *w*, *sx*, *sy*, *sz*`)`
`streamribbon (`*u*, *v*, *w*, *sx*, *sy*, *sz*`)`
`streamribbon (`*xyz*, *x*, *y*, *z*, *anlr_spd*, *lin_spd*`)`
`streamribbon (`*xyz*, *anlr_spd*, *lin_spd*`)`
`streamribbon (`*xyz*, *anlr_rot*`)`
`streamribbon (...,` *width*`)`
`streamribbon (`*hax*, `...)`
*h* `= streamribbon (...)`

Calculate and display streamribbons.

The streamribbon is constructed by rotating a normal vector around a streamline according to the angular rotation of the vector field.

The vector field is given by `[u, v, w]` and is defined over a rectangular grid given by `[x, y, z]`. The streamribbons start at the seed points `[sx, sy, sz]`.

`streamribbon` can be called with a cell array that contains pre-computed streamline data. To do this, *xyz* must be created with the `stream3` function. *lin_spd* is the linear speed of the vector field and can be calculated from `[u, v, w]` by the square root of the sum of the squares. The angular speed *anlr_spd* is the projection of the angular velocity onto the velocity of the normalized vector field and can be calculated with the `curl` command. This option is useful if you need to alter the integrator step size or the maximum number of streamline vertices.

Alternatively, ribbons can be created from an array of vertices *xyz* of a path curve. *anlr_rot* contains the angles of rotation around the edges between adjacent vertices of the path curve.

The input parameter *width* sets the width of the streamribbons.

Streamribbons are colored according to the total angle of rotation along the ribbon.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the plot objects created for each streamribbon.

Example:

```
[x, y, z] = meshgrid (0:0.2:4, -1:0.2:1, -1:0.2:1);
u = - x + 10;
v = 10 * z.*x;
w = - 10 * y.*x;
streamribbon (x, y, z, u, v, w, [0, 0], [0, 0.6], [0, 0]);
view (3);
```

**See also:** [streamline], page 360, [stream3], page 361, [streamtube], page 359, [ostreamtube], page 360.

`streamtube (`*x*, *y*, *z*, *u*, *v*, *w*, *sx*, *sy*, *sz*`)`
`streamtube (`*u*, *v*, *w*, *sx*, *sy*, *sz*`)`
`streamtube (`*xyz*, *x*, *y*, *z*, *div*`)`
`streamtube (`*xyz*, *div*`)`
`streamtube (`*xyz*, *dia*`)`
`streamtube (...,` *options*`)`
`streamtube (`*hax*, `...)`
*h* `= streamtube (...)`

Plot tubes scaled by the divergence along streamlines.

`streamtube` draws tubes whose diameter is scaled by the divergence of a vector field. The vector field is given by `[`*u*`,` *v*`,` *w*`]` and is defined over a rectangular grid given by `[`*x*`,` *y*`,` *z*`]`. The tubes start at the seed points `[`*sx*`,` *sy*`,` *sz*`]` and are plot along streamlines.

`streamtube` can also be called with a cell array containing pre-computed streamline data. To do this, *xyz* must be created with the `stream3` command. *div* is used to scale the tubes. In order to plot tubes scaled by the vector field divergence, *div* must be calculated with the `divergence` command.

A tube diameter of zero corresponds to the smallest scaling value along the streamline and the largest tube diameter corresponds to the largest scaling value.

It is also possible to draw a tube along an arbitrary array of vertices *xyz*. The tube diameter can be specified by the vertex array *dia* or by a constant.

The input parameter *options* is a 2-D vector of the form `[`*scale*`,` *n*`]`. The first parameter scales the tube diameter (default 1). The second parameter specifies the number of vertices that are used to construct the tube circumference (default 20).

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the plot objects created for each tube.

**See also:** [stream3], page 361, [streamline], page 360, [streamribbon], page 358, [ostreamtube], page 360.

`ostreamtube (x, y, z, u, v, w, sx, sy, sz)`
`ostreamtube (u, v, w, sx, sy, sz)`
`ostreamtube (xyz, x, y, z, u, v, w)`
`ostreamtube (..., options)`
`ostreamtube (hax, ...)`
`h = ostreamtube (...)`

Calculate and display streamtubes.

Streamtubes are approximated by connecting circular crossflow areas along a streamline. The expansion of the flow is determined by the local crossflow divergence.

The vector field is given by `[u, v, w]` and is defined over a rectangular grid given by `[x, y, z]`. The streamtubes start at the seed points `[sx, sy, sz]`.

The tubes are colored based on the local vector field strength.

The input parameter *options* is a 2-D vector of the form `[scale, n]`. The first parameter scales the start radius of the streamtubes (default 1). The second parameter specifies the number of vertices that are used to construct the tube circumference (default 20).

`ostreamtube` can be called with a cell array containing pre-computed streamline data. To do this, *xyz* must be created with the `stream3` function. This option is useful if you need to alter the integrator step size or the maximum number of vertices of the streamline.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the plot objects created for each streamtube.

Example:

```
[x, y, z] = meshgrid (-1:0.1:1, -1:0.1:1, -3:0.1:0);
u = -x / 10 - y;
v = x - y / 10;
w = - ones (size (x)) / 10;
ostreamtube (x, y, z, u, v, w, 1, 0, 0);
```

**See also:** [stream3], page 361, [streamline], page 360, [streamribbon], page 358, [streamtube], page 359.

`streamline (x, y, z, u, v, w, sx, sy, sz)`
`streamline (u, v, w, sx, sy, sz)`
`streamline (..., options)`
`streamline (hax, ...)`
`h = streamline (...)`

Plot streamlines of 2-D or 3-D vector fields.

Plot streamlines of a 2-D or 3-D vector field given by `[u, v]` or `[u, v, w]`. The vector field is defined over a rectangular grid given by `[x, y]` or `[x, y, z]`. The streamlines start at the seed points `[sx, sy]` or `[sx, sy, sz]`.

The input parameter *options* is a 2-D vector of the form `[stepsize, max_vertices]`. The first parameter specifies the step size used for trajectory integration (default 0.1). A negative value is allowed which will reverse the direction of integration. The second parameter specifies the maximum number of segments used to create a streamline (default 10,000).

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the hggroup comprising the field lines.

Example:

```
[x, y] = meshgrid (-1.5:0.2:2, -1:0.2:2);
u = - x / 4 - y;
v = x - y / 4;
streamline (x, y, u, v, 1.7, 1.5);
```

**See also:** [stream2], page 361, [stream3], page 361, [streamribbon], page 358, [streamtube], page 359, [ostreamtube], page 360.

`xy = stream2 (x, y, u, v, sx, sy)`
`xy = stream2 (u, v, sx, sy)`
`xy = stream2 (..., options)`

Compute 2-D streamline data.

Calculates streamlines of a vector field given by `[u, v]`. The vector field is defined over a rectangular grid given by `[x, y]`. The streamlines start at the seed points `[sx, sy]`. The returned value *xy* contains a cell array of vertex arrays. If the starting point is outside the vector field, `[]` is returned.

The input parameter *options* is a 2-D vector of the form `[stepsize, max_vertices]`. The first parameter specifies the step size used for trajectory integration (default 0.1). A negative value is allowed which will reverse the direction of integration. The second parameter specifies the maximum number of segments used to create a streamline (default 10,000).

The return value *xy* is a nverts x 2 matrix containing the coordinates of the field line segments.

Example:

```
[x, y] = meshgrid (0:3);
u = 2 * x;
v = y;
xy = stream2 (x, y, u, v, 1.0, 0.5);
```

**See also:** [streamline], page 360, [stream3], page 361.

`xyz = stream3 (x, y, z, u, v, w, sx, sy, sz)`
`xyz = stream3 (u, v, w, sx, sy, sz)`

*xyz* = stream3 (*...*, *options*)

Compute 3-D streamline data.

Calculate streamlines of a vector field given by `[u, v, w]`. The vector field is defined over a rectangular grid given by `[x, y, z]`. The streamlines start at the seed points `[sx, sy, sz]`. The returned value *xyz* contains a cell array of vertex arrays. If the starting point is outside the vector field, `[]` is returned.

The input parameter *options* is a 2-D vector of the form `[stepsize, max_vertices]`. The first parameter specifies the step size used for trajectory integration (default 0.1). A negative value is allowed which will reverse the direction of integration. The second parameter specifies the maximum number of segments used to create a streamline (default 10,000).

The return value *xyz* is a nverts x 3 matrix containing the coordinates of the field line segments.

Example:

```
[x, y, z] = meshgrid (0:3);
u = 2 * x;
v = y;
w = 3 * z;
xyz = stream3 (x, y, z, u, v, w, 1.0, 0.5, 0.0);
```

**See also:** [stream2], page 361, [streamline], page 360, [streamribbon], page 358, [streamtube], page 359, [ostreamtube], page 360.

compass (*u*, *v*)
compass (*z*)
compass (*...*, *style*)
compass (*hax*, *...*)
*h* = compass (*...*)

Plot the `(u, v)` components of a vector field emanating from the origin of a polar plot.

The arrow representing each vector has one end at the origin and the tip at [*u*(i), *v*(i)]. If a single complex argument *z* is given, then `u = real (z)` and `v = imag (z)`.

The style to use for the plot can be defined with a line style *style* of the same format as the `plot` command.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a vector of graphics handles to the line objects representing the drawn vectors.

```
a = toeplitz ([1;randn(9,1)], [1,randn(1,9)]);
compass (eig (a));
```

**See also:** [polar], page 355, [feather], page 362, [quiver], page 357, [rose], page 347, [plot], page 334.

feather (*u*, *v*)
feather (*z*)
feather (*...*, *style*)

`feather (hax, ...)`
`h = feather (...)`

Plot the `(u, v)` components of a vector field emanating from equidistant points on the x-axis.

If a single complex argument $z$ is given, then `u = real (z)` and `v = imag (z)`.

The style to use for the plot can be defined with a line style *style* of the same format as the `plot` command.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value $h$ is a vector of graphics handles to the line objects representing the drawn vectors.

```
phi = [0 : 15 : 360] * pi/180;
feather (sin (phi), cos (phi));
```

**See also:** [plot], page 334, [quiver], page 357, [compass], page 362.

`pcolor (x, y, c)`
`pcolor (c)`
`pcolor (hax, ...)`
`h = pcolor (...)`

Produce a 2-D density plot.

A `pcolor` plot draws rectangles with colors from the matrix $c$ over the two-dimensional region represented by the matrices $x$ and $y$. $x$ and $y$ are the coordinates of the mesh's vertices and are typically the output of `meshgrid`. If $x$ and $y$ are vectors, then a typical vertex is ($x$(j), $y$(i), $c$(i,j)). Thus, columns of $c$ correspond to different $x$ values and rows of $c$ correspond to different $y$ values.

The values in $c$ are scaled to span the range of the current colormap. Limits may be placed on the color axis by the command `clim`, or by setting the `clim` property of the parent axis.

The face color of each cell of the mesh is determined by interpolating the values of $c$ for each of the cell's vertices; Contrast this with `imagesc` which renders one cell for each element of $c$.

`shading` modifies an attribute determining the manner by which the face color of each cell is interpolated from the values of $c$, and the visibility of the cells' edges. By default the attribute is `"faceted"`, which renders a single color for each cell's face with the edge visible.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value $h$ is a graphics handle to the created surface object.

**See also:** [clim], page 368, [shading], page 407, [meshgrid], page 398, [contour], page 348, [imagesc], page 946.

`area (y)`
`area (x, y)`
`area (..., lvl)`

`area (...,` *`prop`*`,` *`val`*`, ...)`
`area (`*`hax`*`, ...)`
*`h`* `= area (...)`

Area plot of the columns of *y*.

This plot shows the contributions of each column value to the row sum. It is functionally similar to `plot (x, cumsum (y, 2))`, except that the area under the curve is shaded.

If the *x* argument is omitted it defaults to `1:rows (y)`. A value *lvl* can be defined that determines where the base level of the shading under the curve should be defined. The default level is 0.

Additional property/value pairs are passed directly to the underlying patch object. The full list of properties is documented at Section 15.3.3.8 [Patch Properties], page 495.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the hggroup object comprising the area patch objects. The `"BaseValue"` property of the hggroup can be used to adjust the level where shading begins.

Example: Verify identity sin^2 + cos^2 = 1

```
t = linspace (0, 2*pi, 100)';
y = [sin(t).^2, cos(t).^2];
area (t, y);
legend ("sin^2", "cos^2", "location", "NorthEastOutside");
```

**See also:** [plot], page 334, [patch], page 451.

`fill (`*`x`*`,` *`y`*`,` *`c`*`)`
`fill (`*`x1`*`,` *`y1`*`,` *`c1`*`,` *`x2`*`,` *`y2`*`,` *`c2`*`)`
`fill (...,` *`prop`*`,` *`val`*`)`
`fill (`*`hax`*`, ...)`
*`h`* `= fill (...)`

Create one or more filled 2-D polygons.

The inputs *x* and *y* are the coordinates of the polygon vertices. If the inputs are matrices then the rows represent different vertices and each column produces a different polygon. `fill` will close any open polygons before plotting.

The input *c* determines the color of the polygon. The simplest form is a single color specification such as a `plot` format or an RGB-triple. In this case the polygon(s) will have one unique color. If *c* is a vector or matrix then the color data is first scaled using `clim` and then indexed into the current colormap. A vector will color each polygon (a column from matrices *x* and *y*) with a single computed color. A matrix *c* of the same size as *x* and *y* will compute the color of each vertex and then interpolate the face color between the vertices.

Multiple property/value pairs for the underlying patch object may be specified, but they must appear in pairs. The full list of properties is documented at Section 15.3.3.8 [Patch Properties], page 495.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a vector of graphics handles to the created patch objects.

Example: red square

```
vertices = [0 0
            1 0
            1 1
            0 1];
fill (vertices(:,1), vertices(:,2), "r");
axis ([-0.5 1.5, -0.5 1.5])
axis equal
```

**See also:** [patch], page 451, [fill3], page 365, [clim], page 368, [colormap], page 951.

`fill3` (*x*, *y*, *z*, *c*)
`fill3` (*x1*, *y1*, *z1*, *c1*, *x2*, *y2*, *z2*, *c2*)
`fill3` (..., *prop*, *val*)
`fill3` (*hax*, ...)
*h* = `fill3` (...)

Create one or more filled 3-D polygons.

The inputs *x*, *y*, and *z* are the coordinates of the polygon vertices. If the inputs are matrices then the rows represent different vertices and each column produces a different polygon. `fill3` will close any open polygons before plotting.

The input *c* determines the color of the polygon. The simplest form is a single color specification such as a `plot` format or an RGB-triple. In this case the polygon(s) will have one unique color. If *c* is a vector or matrix then the color data is first scaled using `clim` and then indexed into the current colormap. A vector will color each polygon (a column from matrices *x*, *y*, and *z*) with a single computed color. A matrix *c* of the same size as *x*, *y*, and *z* will compute the color of each vertex and then interpolate the face color between the vertices.

Multiple property/value pairs for the underlying patch object may be specified, but they must appear in pairs. The full list of properties is documented at Section 15.3.3.8 [Patch Properties], page 495.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a vector of graphics handles to the created patch objects.

Example: oblique red rectangle

```
vertices = [0 0 0
            1 1 0
            1 1 1
            0 0 1];
fill3 (vertices(:,1), vertices(:,2), vertices(:,3), "r");
axis ([-0.5 1.5, -0.5 1.5, -0.5 1.5]);
axis equal
grid on
view (-80, 25);
```

**See also:** [patch], page 451, [fill], page 364, [clim], page 368, [colormap], page 951.

`comet (`*y*`)`
`comet (`*x*`, `*y*`)`
`comet (`*x*`, `*y*`, `*p*`)`
`comet (`*hax*`, ...)`

Produce a simple comet style animation along the trajectory provided by the input coordinate vectors (*x*, *y*).

If *x* is not specified it defaults to the indices of *y*.

The speed of the comet may be controlled by *p*, which represents the time each point is displayed before moving to the next one. The default for *p* is `5 / numel (`*y*`)`.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

**See also:** [comet3], page 366.

`comet3 (`*z*`)`
`comet3 (`*x*`, `*y*`, `*z*`)`
`comet3 (`*x*`, `*y*`, `*z*`, `*p*`)`
`comet3 (`*hax*`, ...)`

Produce a simple comet style animation along the trajectory provided by the input coordinate vectors (*x*, *y*, *z*).

If only *z* is specified then *x*, *y* default to the indices of *z*.

The speed of the comet may be controlled by *p*, which represents the time each point is displayed before moving to the next one. The default for *p* is `5 / numel (`*z*`)`.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

**See also:** [comet], page 366.

#### 15.2.1.1 Axis Configuration

The axis function may be used to change the axis limits of an existing plot and various other axis properties, such as the aspect ratio and the appearance of tic marks. By default, high level plotting functions such as `plot` reset axes properties. Any customization of properties, for example by calling `axis`, `xlim`, etc., should happen after the plot is done or, alternatively, after calling the [hold function], page 429.

`axis` ()
`axis` ([*x_lo x_hi*])
`axis` ([*x_lo x_hi y_lo y_hi*])
`axis` ([*x_lo x_hi y_lo y_hi z_lo z_hi*])
`axis` ([*x_lo x_hi y_lo y_hi z_lo z_hi c_lo c_hi*])
`axis` (*option*)
`axis` (*option1*, *option2*, ...)
`axis` (*hax*, ...)
*limits* = `axis` ()

Set axis limits and appearance.

The argument *limits* should be a 2-, 4-, 6-, or 8-element vector. The first and second elements specify the lower and upper limits for the x-axis. The third and fourth specify the limits for the y-axis, the fifth and sixth specify the limits for the z-axis, and the seventh and eighth specify the limits for the color axis. The special values `-Inf` and `Inf` may be used to indicate that the limit should be automatically computed based on the data in the axes.

Without any arguments, `axis` turns autoscaling on.

With one output argument, *limits* = `axis` returns the current axis limits.

The vector argument specifying limits is optional, and additional string arguments may be used to specify various axis properties.

The following options control the aspect ratio of the axes.

`"equal"` Force x-axis unit distance to equal y-axis (and z-axis) unit distance.

`"square"` Force a square axis aspect ratio.

`"vis3d"` Set aspect ratio modes (`"DataAspectRatio"`, `"PlotBoxAspectRatio"`) to `"manual"` for rotation without stretching.

`"normal"`
`"fill"` Restore default automatically computed aspect ratios.

The following options control the way axis limits are interpreted.

`"auto"`
`"auto[xyz]"`
`"auto [xyz]"`
Set nice auto-computed limits around the data for all axes, or only the specified axes.

`"manual"` Fix the current axes limits.

`"tickaligned"`
Fix axes to the limits of the closest ticks.

`"tight"` Fix axes to the limits of the data.

`"padded"` Fix axes to the limits of the data plus margins of about 7% of the data extent.

`"image"` Equivalent to `"tight"` and `"equal"`.

The following options affect the appearance of tick marks.

`"tic"`
`"tic[xyz]"`
`"tic [xyz]"`
: Turn tick marks on for all axes, or turn them on for the specified axes and off for the remainder.

`"label"`
`"label[xyz]"`
`"label [xyz]"`
: Turn tick labels on for all axes, or turn them on for the specified axes and off for the remainder.

`"nolabel"`
: Turn tick labels off for all axes.

The following options affect the direction of increasing values on the axes.

`"xy"` Default y-axis, larger values are near the top.

`"ij"` Reverse y-axis, smaller values are near the top.

The following options affects the visibility of the axes.

`"on"` Make the axes visible.

`"off"` Hide the axes.

If the first argument *hax* is an axes handle, then operate on this axes rather than the current axes returned by `gca`.

Example 1: set X/Y limits and force a square aspect ratio

```
axis ([1, 2, 3, 4], "square");
```

Example 2: enable tick marks on all axes, enable tick mark labels only on the y-axis

```
axis ("tic", "labely");
```

**See also:** [xlim], page 369, [ylim], page 370, [zlim], page 371, [clim], page 368, [daspect], page 408, [pbaspect], page 409, [box], page 420, [grid], page 420.

Similarly the axis limits of the colormap can be changed with the `clim` function.

`clim` ([*cmin cmax*])
`clim` ("*auto*")
`clim` ("*manual*")
`clim` (*hax*, ...)
*limits* = `clim` ()
: Query or set color axis limits for plots.

The limits argument should be a 2-element vector specifying the lower and upper limits to assign to the first and last value in the colormap. Data values outside this range are clamped to the first and last colormap entries.

If the `"auto"` option is given then automatic colormap limits are applied. The automatic algorithm sets *cmin* to the minimum data value and *cmax* to the maximum

data value. If `"manual"` is specified then the `"climmode"` property is set to `"manual"` and the numeric values in the `"clim"` property are used for limits.

If the first argument *hax* is an axes handle, then operate on this axes rather than the current axes returned by `gca`.

Called without arguments the current color axis limits are returned.

Programming Note: The color axis affects the display of image, patch, and surface graphics objects, but **only** if the `"cdata"` property has indexed data and the `"cdatamapping"` property is set to `"scaled"`. Graphic objects with true color `cdata`, or `"direct"` `cdatamapping` are not affected.

**See also:** [colormap], page 951, [axis], page 366.

The `xlim`, `ylim`, and `zlim` functions may be used to get or set individual axis limits. Each has the same form.

*xlimits* = xlim ()
*xmode* = xlim ("*mode*")
*xmethod* = xlim ("*method*")
xlim ([*x_lo x_hi*])
xlim ("*mode*")
xlim ("*method*")
xlim (*hax*, ...)

Query or set the limits of the x-axis for the current plot.

Called without arguments `xlim` returns the x-axis limits of the current plot.

With the input query `"mode"`, return the current x-limit calculation mode which is either `"auto"` or `"manual"`.

With the input query `"method"`, return the current x-limit calculation method which is either `"tickaligned"`, `"tight"`, or `"padded"`.

If passed a 2-element vector [*x_lo x_hi*], the limits of the x-axis are set to these values and the mode is set to `"manual"`. The special values -Inf and Inf can be used to indicate that either the lower axis limit or upper axis limit should be automatically calculated.

The current limit calculation `"mode"` may be one of

`"auto"` (default)
: Automatically calculate limits based on the plot data and the currently specified limit calculation method.

`"manual"`
: Fix axis limits at current values.

The current limit calculation method may be one of

`"tickaligned"` (default)
: Calculate limits that encompass all of the data and extend outwards to the nearest tick mark.

`"tight"`
: Calculate limits that exactly fit the data range.

`"padded"`
: Calculate limits that leave a margin around the data of approximately 7% of the data range.

If the first argument *hax* is an axes handle, then operate on this axes rather than the current axes returned by `gca`.

Programming Note: The `xlim` function operates by modifying the `"xlim"`, `"xlimmode"`, and `"xlimitmethod"` properties of an axes object. These properties can be directly inspected and altered with `get`/`set`.

**See also:** [ylim], page 370, [zlim], page 371, [axis], page 366, [set], page 457, [get], page 457, [gca], page 455.

*ylimits* = ylim ()
*ymode* = ylim ("*mode*")
*ymethod* = ylim ("*method*")
ylim ([*y_lo y_hi*])
ylim ("*mode*")
ylim ("*method*")
ylim (*hax*, ...)

Query or set the limits of the y-axis for the current plot.

Called without arguments `ylim` returns the y-axis limits of the current plot.

With the input query `"mode"`, return the current y-limit calculation mode which is either `"auto"` or `"manual"`.

With the input query `"method"`, return the current y-limit calculation method which is either `"tickaligned"`, `"tight"`, or `"padded"`.

If passed a 2-element vector [*y_lo y_hi*], the limits of the y-axis are set to these values and the mode is set to `"manual"`. The special values -Inf and Inf can be used to indicate that either the lower axis limit or upper axis limit should be automatically calculated.

The current limit calculation `"mode"` may be one of

`"auto"` (default)
: Automatically calculate limits based on the plot data and the currently specified limit calculation method.

`"manual"`
: Fix axis limits at current values.

The current limit calculation method may be one of

`"tickaligned"` (default)
: Calculate limits that encompass all of the data and extend outwards to the nearest tick mark.

`"tight"`
: Calculate limits that exactly fit the data range.

`"padded"`
: Calculate limits that leave a margin around the data of approximately 7% of the data range.

If the first argument *hax* is an axes handle, then operate on this axes rather than the current axes returned by `gca`.

Programming Note: The `ylim` function operates by modifying the `"ylim"`, `"ylimmode"`, and `"ylimitmethod"` properties of an axes object. These properties can be directly inspected and altered with `get`/`set`.

**See also:** [xlim], page 369, [zlim], page 371, [axis], page 366, [set], page 457, [get], page 457, [gca], page 455.

*zlimits* = zlim ()
*zmode* = zlim ("*mode*")
*zmethod* = zlim ("*method*")
zlim ([*z_lo z_hi*])
zlim ("*mode*")
zlim ("*method*")
zlim (*hax*, ...)

Query or set the limits of the z-axis for the current plot.

Called without arguments `zlim` returns the z-axis limits of the current plot.

With the input query `"mode"`, return the current z-limit calculation mode which is either `"auto"` or `"manual"`.

With the input query `"method"`, return the current z-limit calculation method which is either `"tickaligned"`, `"tight"`, or `"padded"`.

If passed a 2-element vector [*z_lo z_hi*], the limits of the z-axis are set to these values and the mode is set to `"manual"`. The special values -Inf and Inf can be used to indicate that either the lower axis limit or upper axis limit should be automatically calculated.

The current limit calculation `"mode"` may be one of

`"auto"` (default)
: Automatically calculate limits based on the plot data and the currently specified limit calculation method.

`"manual"`
: Fix axis limits at current values.

The current limit calculation method may be one of

`"tickaligned"` (default)
: Calculate limits that encompass all of the data and extend outwards to the nearest tick mark.

`"tight"`
: Calculate limits that exactly fit the data range.

`"padded"`
: Calculate limits that leave a margin around the data of approximately 7% of the data range.

If the first argument *hax* is an axes handle, then operate on this axes rather than the current axes returned by `gca`.

Programming Note: The `zlim` function operates by modifying the `"zlim"`, `"zlimmode"`, and `"zlimitmethod"` properties of an axes object. These properties can be directly inspected and altered with `get`/`set`.

**See also:** [xlim], page 369, [ylim], page 370, [axis], page 366, [set], page 457, [get], page 457, [gca], page 455.

The `xticks`, `yticks`, `zticks`, `rticks`, and `thetaticks` functions may be used to get or set the tick mark locations and modes on the respective axis. Each has the same form, although mode options are not currently available for `rticks`, and `thetaticks`.

*tickval* = xticks
*mode* = xticks ("*mode*")
xticks (*tickval*)
xticks ("*auto*")
xticks ("*manual*")
... = xticks (*hax*, ...)

Query or set the tick values on the x-axis of the current axis.

When called without an argument, return the current tick locations as specified in the "xtick" axes property. These locations can be changed by calling xticks with a vector of tick values. Note: ascending order is not required.

When called with argument "mode", xticks returns the current value of the axes property "xtickmode". This property can be changed by calling xticks with either "auto" (algorithm determines tick positions) or "manual" (tick values remain fixed regardless of axes resizing or rotation). Note: Specifying xtick values will also set the property "xtickmode" to "manual".

If the first argument *hax* is an axes handle, then operate on this axis rather than the current axes returned by gca.

Requesting a return value when calling xticks to set a property value will result in an error.

**See also:** [xticklabels], page 374, [yticks], page 372, [zticks], page 373, [rticks], page 373, [thetaticks], page 373, [get], page 457, [set], page 457.

*tickval* = yticks
*mode* = yticks ("*mode*")
yticks (*tickval*)
yticks ("*auto*")
yticks ("*manual*")
... = yticks (*hax*, ...)

Query or set the tick values on the y-axis of the current axis.

When called without an argument, return the current tick locations as specified in the "ytick" axes property. These locations can be changed by calling yticks with a vector of tick values. Note: ascending order is not required.

When called with argument "mode", yticks returns the current value of the axes property "ytickmode". This property can be changed by calling yticks with either "auto" (algorithm determines tick positions) or "manual" (tick values remain fixed regardless of axes resizing or rotation). Note: Specifying ytick values will also set the property "ytickmode" to "manual".

If the first argument *hax* is an axes handle, then operate on this axis rather than the current axes returned by gca.

Requesting a return value when calling yticks to set a property value will result in an error.

**See also:** [yticklabels], page 374, [xticks], page 371, [zticks], page 373, [rticks], page 373, [thetaticks], page 373, [get], page 457, [set], page 457.

*tickval* = zticks
*mode* = zticks ("*mode*")
zticks (*tickval*)
zticks ("*auto*")
zticks ("*manual*")
... = zticks (*hax*, ...)

Query or set the tick values on the z-axis of the current axis.

When called without an argument, return the current tick locations as specified in the "ztick" axes property. These locations can be changed by calling zticks with a vector of tick values. Note: ascending order is not required.

When called with argument "mode", zticks returns the current value of the axes property "ztickmode". This property can be changed by calling zticks with either "auto" (algorithm determines tick positions) or "manual" (tick values remain fixed regardless of axes resizing or rotation). Note: Specifying ztick values will also set the property "ztickmode" to "manual".

If the first argument *hax* is an axes handle, then operate on this axis rather than the current axes returned by gca.

Requesting a return value when calling zticks to set a property value will result in an error.

**See also:** [zticklabels], page 375, [xticks], page 371, [yticks], page 372, [rticks], page 373, [thetaticks], page 373, [get], page 457, [set], page 457.

*tickval* = rticks
rticks (*tickval*)
... = rticks (*hax*, ...)

Query or set the tick values on the r-axis of the current axis.

When called without argument, return the current tick locations as specified in the "rtick" axes property. These locations can be changed by calling rticks with a vector of tick values. Note: ascending order is not required.

If the first argument *hax* is an axes handle, then operate on this axis rather than the current axes returned by gca.

Requesting a return value when calling rticks to set a property value will result in an error.

NOTE: Octave does not currently implement polaraxes objects. It is therefore not possible to query or set a "mode" for the "rtick" property as can be done with the equivalent functions for *x*, *y*, and *z* axes.

**See also:** [thetaticks], page 373, [xticks], page 371, [yticks], page 372, [zticks], page 373, [polar], page 355, [get], page 457, [set], page 457.

*tickval* = thetaticks
thetaticks (*tickval*)
... = thetaticks (*hax*, ...)

Query or set the tick values on the theta-axis of the current axis.

When called without argument, return the current tick locations as specified in the "ttick" axes property. These locations can be changed by calling thetaticks with a vector of tick values. Note: ascending order is not required.

If the first argument *hax* is an axes handle, then operate on this axis rather than the current axes returned by `gca`.

Requesting a return value when calling `thetaticks` to set a property value will result in an error.

NOTE: Octave does not currently implement polaraxes objects. It is therefore not possible to query or set a `"mode"` for the `"thetatick"` property as can be done with the equivalent functions for *x*, *y*, and *z* axes.

**See also:** [rticks], page 373, [xticks], page 371, [yticks], page 372, [zticks], page 373, [polar], page 355, [get], page 457, [set], page 457.

The `xticklabels`, `yticklabels`, and `zticklabels` functions may be used to get or set the label assigned to each tick location and the labeling mode on the respective axis. Each has the same form.

*labels* = xticklabels
*mode* = xticklabels ("*mode*")
xticklabels (*tickval*)
xticklabels ("*auto*")
xticklabels ("*manual*")
... = xticklabels (*hax*, ...)

Query or set the tick labels on the x-axis of the current axis.

When called without an argument, return a cell array of strings of the current tick labels as specified in the `"xticklabel"` axes property. These labels can be changed by calling `xticklabels` with a cell array of strings. Note: a vector of numbers will be mapped to a cell array of strings. If fewer labels are specified than the current number of ticks, blank labels will be appended to the array.

When called with argument `"mode"`, `xticklabels` returns the current value of the axes property `"xticklabelmode"`. This property can be changed by calling `xticklabels` with either `"auto"` (algorithm determines tick labels) or `"manual"` (tick labels remain fixed). Note: Specifying xticklabel values will also set the `"xticklabelmode"` and `"xticks"` properties to `"manual"`.

If the first argument *hax* is an axes handle, then operate on this axis rather than the current axes returned by `gca`.

Requesting a return value when calling `xticklabels` to set a property value will result in an error.

**See also:** [xticks], page 371, [yticklabels], page 374, [zticklabels], page 375, [get], page 457, [set], page 457.

*labels* = yticklabels
*mode* = yticklabels ("*mode*")
yticklabels (*tickval*)
yticklabels ("*auto*")
yticklabels ("*manual*")
... = yticklabels (*hax*, ...)

Query or set the tick labels on the x-axis of the current axis.

When called without an argument, return a cell array of strings of the current tick labels as specified in the `"yticklabel"` axes property. These labels can be changed by calling `yticklabels` with a cell array of strings. Note: a vector of numbers will be mapped to a cell array of strings. If fewer labels are specified than the current number of ticks, blank labels will be appended to the array.

When called with argument `"mode"`, `yticklabels` returns the current value of the axes property `"yticklabelmode"`. This property can be changed by calling `yticklabels` with either `"auto"` (algorithm determines tick labels) or `"manual"` (tick labels remain fixed). Note: Specifying yticklabel values will also set the `"yticklabelmode"` and `"yticks"` properties to `"manual"`.

If the first argument *hax* is an axes handle, then operate on this axis rather than the current axes returned by `gca`.

Requesting a return value when calling `xticklabels` to set a property value will result in an error.

**See also:** [yticks], page 372, [xticklabels], page 374, [zticklabels], page 375, [get], page 457, [set], page 457.

*labels* = zticklabels
*mode* = zticklabels ("*mode*")
zticklabels (*tickval*)
zticklabels ("*auto*")
zticklabels ("*manual*")
... = zticklabels (*hax*, ...)

Query or set the tick labels on the x-axis of the current axis.

When called without an argument, return a cell array of strings of the current tick labels as specified in the `"zticklabel"` axes property. These labels can be changed by calling `zticklabels` with a cell array of strings. Note: a vector of numbers will be mapped to a cell array of strings. If fewer labels are specified than the current number of ticks, blank labels will be appended to the array.

When called with argument `"mode"`, `zticklabels` returns the current value of the axes property `"zticklabelmode"`. This property can be changed by calling `zticklabels` with either `"auto"` (algorithm determines tick labels) or `"manual"` (tick labels remain fixed). Note: Specifying zticklabel values will also set the `"zticklabelmode"` and `"zticks"` properties to `"manual"`.

If the first argument *hax* is an axes handle, then operate on this axis rather than the current axes returned by `gca`.

Requesting a return value when calling `xticklabels` to set a property value will result in an error.

**See also:** [zticks], page 373, [xticklabels], page 374, [zticklabels], page 375, [get], page 457, [set], page 457.

*labels* = rticklabels
rticklabels (*tickval*)
... = rticklabels (*hax*, ...)

Query or set the tick labels on the r-axis of a polar plot.

When called without an argument, return a cell array of strings of the current rtick labels.

When called with the argument *tickval* being a vector of numbers or a cell array of strings and/or numbers, the labels will be changed to match these new values. Note that the center point of the plots made by `polar` are never labeled, so the first specified label will be applied to the second rtick location and subsequent labels will progress outward.

If fewer labels are specified than the current number of tick marks, those labels will be applied starting with the innermost tick labels, and blank labels will be appended to the remainder. If the specified label count exceeds the number of tick labels, the excess labels are ignored.

If the first argument *hax* is an axes handle, then operate on this axis rather than the current axes returned by `gca`.

Requesting a return value when calling `rticklabels` to set a property value will result in an error.

Compatibility Note: The `'mode'` property for rticklabels has not yet been implemented.

**See also:** [polar], page 355, [rticks], page 373, [thetaticklabels], page 376, [xticklabels], page 374, [yticklabels], page 374, [zticklabels], page 375, [get], page 457, [set], page 457.

*labels* = thetaticklabels
thetaticklabels (*tickval*)
... = thetaticklabels (*hax*, ...)

Query or set the tick labels on the theta-axis of a polar plot.

When called without an argument, return a cell array of strings of the current ttick labels.

When called with the argument *tickval* being a vector of numbers or a cell array of strings and/or numbers, the labels will be changed to match these new values. Values will be applied starting with the zero-degree tick mark and will progress counterclockwise.

If fewer labels are specified than the current number of theta tick marks, those labels will be applied starting at the zero-degree tick mark and blank labels will be appended to the remainder. If the specified label count exceeds the number of tick labels, the excess labels are ignored.

If the first argument *hax* is an axes handle, then operate on this axis rather than the current axes returned by `gca`.

Requesting a return value when calling `thetaticklabels` to set a property value will result in an error.

Compatibility Note: The `'mode'` property for thetaticklabels has not yet been implemented.

**See also:** [polar], page 355, [thetaticks], page 373, [rticklabels], page 375, [xticklabels], page 374, [yticklabels], page 374, [zticklabels], page 375, [get], page 457, [set], page 457.

The `xtickangle`, `ytickangle`, and `ztickangle` functions may be used to get or set the rotation angle of labels for the respective axis. Each has the same form.

*angle* = `xtickangle` ()
*angle* = `xtickangle` (*hax*)
`xtickangle` (*angle*)
`xtickangle` (*hax*, *angle*)

Query or set the rotation angle of the tick labels on the x-axis of the current axes.

When called without an argument, return the rotation angle in degrees of the tick labels as specified in the axes property `"XTickLabelRotation"`. When called with a numeric scalar *angle*, rotate the tick labels counterclockwise to *angle* degrees.

If the first argument *hax* is an axes handle, then operate on this axes rather than the current axes returned by `gca`.

Programming Notes:

1. The `"XTickLabelRotation"` property is *not* currently implemented in Octave. The property can be set and queried, but has no effect on the plot.
2. Requesting a return value while also setting a specified rotation will result in an error.

**See also:** [ytickangle], page 377, [ztickangle], page 377, [get], page 457, [set], page 457.

*angle* = `ytickangle` ()
*angle* = `ytickangle` (*hax*)
`ytickangle` (*angle*)
`ytickangle` (*hax*, *angle*)

Query or set the rotation angle of the tick labels on the y-axis of the current axes.

When called without an argument, return the rotation angle in degrees of the tick labels as specified in the axes property `"YTickLabelRotation"`. When called with a numeric scalar *angle*, rotate the tick labels counterclockwise to *angle* degrees.

If the first argument *hax* is an axes handle, then operate on this axes rather than the current axes returned by `gca`.

Programming Notes:

1. The `"YTickLabelRotation"` property is *not* currently implemented in Octave. The property can be set and queried, but has no effect on the plot.
2. Requesting a return value while also setting a specified rotation will result in an error.

**See also:** [xtickangle], page 377, [ztickangle], page 377, [get], page 457, [set], page 457.

*angle* = `ztickangle` ()
*angle* = `ztickangle` (*hax*)
`ztickangle` (*angle*)
`ztickangle` (*hax*, *angle*)

Query or set the rotation angle of the tick labels on the z-axis of the current axes.

When called without an argument, return the rotation angle in degrees of the tick labels as specified in the axes property `"ZTickLabelRotation"`. When called with a numeric scalar *angle*, rotate the tick labels counterclockwise to *angle* degrees.

If the first argument *hax* is an axes handle, then operate on this axes rather than the current axes returned by `gca`.

Programming Notes:

1. The `"ZTickLabelRotation"` property is *not* currently implemented in Octave. The property can be set and queried, but has no effect on the plot.
2. Requesting a return value while also setting a specified rotation will result in an error.

**See also:** [xtickangle], page 377, [ytickangle], page 377, [get], page 457, [set], page 457.

### 15.2.1.2 Two-dimensional Function Plotting

Octave can plot a function from a function handle or string defining the function without the user needing to explicitly create the data to be plotted. The function `fplot` also generates two-dimensional plots with linear axes using a function name and limits for the range of the x-coordinate instead of the x and y data. For example,

```
fplot (@sin, [-10, 10], 201);
```

produces a plot that is equivalent to the one above, but also includes a legend displaying the name of the plotted function.

`fplot (`*`fcn`*`)`
`fplot (`*`fcn`*`, `*`limits`*`)`
`fplot (...,` *`tol`*`)`
`fplot (...,` *`n`*`)`
`fplot (...,` *`fmt`*`)`
`fplot (...,` *`property`*`,` *`value`*`, ...)`
`fplot (`*`hax`*`, ...)`
`[`*`x`*`,` *`y`*`] = fplot (...)`

Plot a function *fcn* within the range defined by *limits*.

*fcn* is a function handle, inline function, or string containing the name of the function to evaluate.

The limits of the plot are of the form `[`*`xlo`*`,` *`xhi`*`]` or `[`*`xlo`*`,` *`xhi`*`,` *`ylo`*`,` *`yhi`*`]`. If no limits are specified the default is `[-5, 5]`.

The next three arguments are all optional and any number of them may be given in any order.

*tol* is the relative tolerance to use for the plot and defaults to 2e-3 (.2%).

*n* is the minimum number of points to use. When *n* is specified, the maximum stepsize will be `(`*`xhi`*` - `*`xlo`*`) / `*`n`*. More than *n* points may still be used in order to meet the relative tolerance requirement.

The *fmt* argument specifies the linestyle to be used by the plot command.

Multiple property-value pairs may also be specified, but they must appear in pairs. These arguments are applied to the line objects drawn by `plot`.

The full list of line properties is documented at Section 15.3.3.5 [Line Properties], page 485.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

With no output arguments, the results are immediately plotted. With two output arguments, the 2-D plot data is returned. The data can subsequently be plotted manually with `plot (x, y)`.

Example:

```
fplot (@cos, [0, 2*pi])
fplot ("[cos(x), sin(x)]", [0, 2*pi])
```

Programming Notes:

`fplot` works best with continuous functions. Functions with discontinuities are unlikely to plot well. This restriction may be removed in the future.

`fplot` performance is better when the function accepts and returns a vector argument. Consider this when writing user-defined functions and use element-by-element operators such as `.*`, `./`, etc.

**See also:** [ezplot], page 379, [plot], page 334.

Other functions that can create two-dimensional plots directly from a function include `ezplot`, `ezcontour`, `ezcontourf` and `ezpolar`.

`ezplot (f)`
`ezplot (f2v)`
`ezplot (fx, fy)`
`ezplot (..., dom)`
`ezplot (..., n)`
`ezplot (hax, ...)`
`h = ezplot (...)`

Plot the 2-D curve defined by the function *f*.

The function *f* may be a string, inline function, or function handle and can have either one or two variables. If *f* has one variable, then the function is plotted over the domain `-2*pi < x < 2*pi` with 500 points.

If *f2v* is a function of two variables then the implicit function `f(x,y) = 0` is calculated over the meshed domain `-2*pi <= x | y <= 2*pi` with 60 points in each dimension.

For example:

```
ezplot (@(x, y) x.^2 - y.^2 - 1)
```

If two functions are passed as inputs then the parametric function

```
x = fx (t)
y = fy (t)
```

is plotted over the domain `-2*pi <= t <= 2*pi` with 500 points.

If *dom* is a two element vector, it represents the minimum and maximum values of both *x* and *y*, or *t* for a parametric plot. If *dom* is a four element vector, then the minimum and maximum values are `[xmin xmax ymin ymax]`.

*n* is a scalar defining the number of points to use in plotting the function.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a vector of graphics handles to the created line objects.

**See also:** [plot], page 334, [ezplot3], page 409, [ezpolar], page 381, [ezcontour], page 380, [ezcontourf], page 380, [ezmesh], page 410, [ezmeshc], page 411, [ezsurf], page 411, [ezsurfc], page 412.

`ezcontour (`*f*`)`
`ezcontour (...,` *dom*`)`
`ezcontour (...,` *n*`)`
`ezcontour (`*hax*`, ...)`
*h* `= ezcontour (...)`

Plot the contour lines of a function.

*f* is a string, inline function, or function handle with two arguments defining the function. By default the plot is over the meshed domain `-2*pi <=` *x* `| y <= 2*pi` with 60 points in each dimension.

If *dom* is a two element vector, it represents the minimum and maximum values of both *x* and *y*. If *dom* is a four element vector, then the minimum and maximum values are `[xmin xmax ymin ymax]`.

*n* is a scalar defining the number of points to use in each dimension.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created plot.

Example:

```
f = @(x,y) sqrt (abs (x .* y)) ./ (1 + x.^2 + y.^2);
ezcontour (f, [-3, 3]);
```

**See also:** [contour], page 348, [ezcontourf], page 380, [ezplot], page 379, [ezmeshc], page 411, [ezsurfc], page 412.

`ezcontourf (`*f*`)`
`ezcontourf (...,` *dom*`)`
`ezcontourf (...,` *n*`)`
`ezcontourf (`*hax*`, ...)`
*h* `= ezcontourf (...)`

Plot the filled contour lines of a function.

*f* is a string, inline function, or function handle with two arguments defining the function. By default the plot is over the meshed domain `-2*pi <=` *x* `| y <= 2*pi` with 60 points in each dimension.

If *dom* is a two element vector, it represents the minimum and maximum values of both *x* and *y*. If *dom* is a four element vector, then the minimum and maximum values are `[xmin xmax ymin ymax]`.

*n* is a scalar defining the number of points to use in each dimension.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created plot.

Example:

```
f = @(x,y) sqrt (abs (x .* y)) ./ (1 + x.^2 + y.^2);
ezcontourf (f, [-3, 3]);
```

**See also:** [contourf], page 349, [ezcontour], page 380, [ezplot], page 379, [ezmeshc], page 411, [ezsurfc], page 412.

`ezpolar (`*`f`*`)`
`ezpolar (...,` *`dom`*`)`
`ezpolar (...,` *`n`*`)`
`ezpolar (`*`hax`*`, ...)`
*`h`* `= ezpolar (...)`

Plot a 2-D function in polar coordinates.

The function *f* is a string, inline function, or function handle with a single argument. The expected form of the function is `rho = f(theta)`. By default the plot is over the domain `0 <= theta <= 2*pi` with 500 points.

If *dom* is a two element vector, it represents the minimum and maximum values of *theta*.

*n* is a scalar defining the number of points to use in plotting the function.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created plot.

Example:

```
ezpolar (@(t) sin (5/4 * t), [0, 8*pi]);
```

**See also:** [polar], page 355, [ezplot], page 379.

### 15.2.1.3 Two-dimensional Geometric Shapes

`rectangle ()`
`rectangle (..., "`*Position*`",` *`pos`*`)`
`rectangle (..., "`*Curvature*`",` *`curv`*`)`
`rectangle (..., "`*EdgeColor*`",` *`ec`*`)`
`rectangle (..., "`*FaceColor*`",` *`fc`*`)`
`rectangle (`*`hax`*`, ...)`
*`h`* `= rectangle (...)`

Draw a rectangular patch defined by *pos* and *curv*.

The variable `pos(1:2)` defines the lower left-hand corner of the patch and `pos(3:4)` defines its width and height. By default, the value of *pos* is `[0, 0, 1, 1]`.

The variable *curv* defines the curvature of the sides of the rectangle and may be a scalar or two-element vector with values between 0 and 1. A value of 0 represents no curvature of the side, whereas a value of 1 means that the side is entirely curved into the arc of a circle. If *curv* is a two-element vector, then the first element is the curvature along the x-axis of the patch and the second along y-axis.

If *curv* is a scalar, it represents the curvature of the shorter of the two sides of the rectangle and the curvature of the other side is defined by

```
min (pos(1:2)) / max (pos(1:2)) * curv
```

Additional property/value pairs are passed to the underlying patch command. The full list of properties is documented at Section 15.3.3.8 [Patch Properties], page 495.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created rectangle object.

**See also:** [patch], page 451, [line], page 450, [cylinder], page 413, [ellipsoid], page 414, [sphere], page 413.

### 15.2.2 Three-Dimensional Plots

The function `mesh` produces mesh surface plots. For example,

```
tx = ty = linspace (-8, 8, 41)';
[xx, yy] = meshgrid (tx, ty);
r = sqrt (xx .^ 2 + yy .^ 2) + eps;
tz = sin (r) ./ r;
mesh (tx, ty, tz);
xlabel ("tx");
ylabel ("ty");
zlabel ("tz");
title ("3-D Sombrero plot");
```

produces the familiar "sombrero" plot shown in Figure 15.5. Note the use of the function `meshgrid` to create matrices of X and Y coordinates to use for plotting the Z data. The `ndgrid` function is similar to `meshgrid`, but works for N-dimensional matrices.

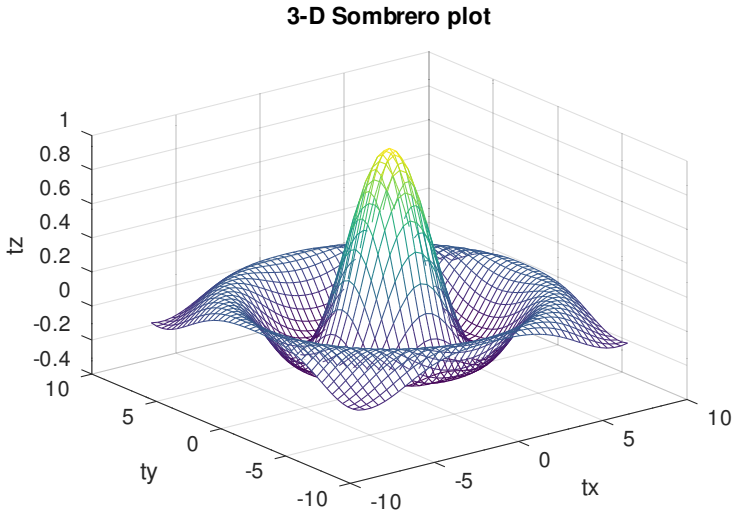



Figure 15.5: Mesh plot.

The `meshc` function is similar to `mesh`, but also produces a plot of contours for the surface.

The `plot3` function displays arbitrary three-dimensional data, without requiring it to form a surface. For example,

```
t = 0:0.1:10*pi;
r = linspace (0, 1, numel (t));
z = linspace (0, 1, numel (t));
plot3 (r.*sin (t), r.*cos (t), z);
xlabel ("r.*sin (t)");
ylabel ("r.*cos (t)");
zlabel ("z");
title ("plot3 display of 3-D helix");
```

displays the spiral in three dimensions shown in Figure 15.6.

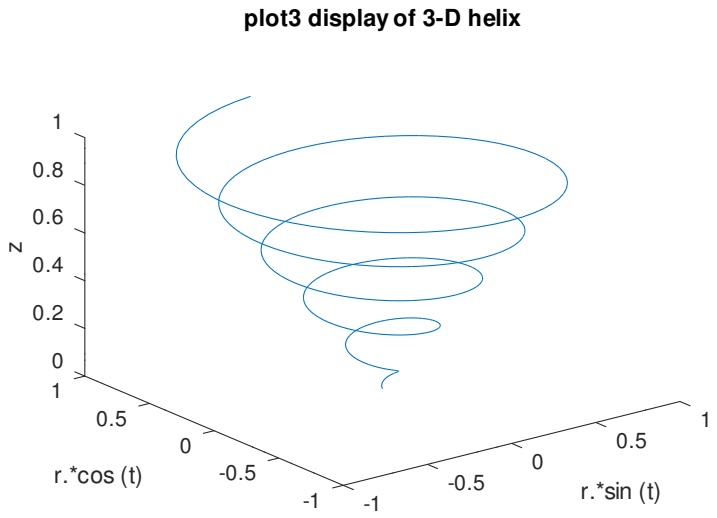



Figure 15.6: Three-dimensional spiral.

Finally, the `view` function changes the viewpoint for three-dimensional plots.

`mesh` (*x*, *y*, *z*)
`mesh` (*z*)
`mesh` (*...*, *c*)
`mesh` (*...*, *prop*, *val*, *...*)
`mesh` (*hax*, *...*)
*h* = `mesh` (*...*)

Plot a 3-D wireframe mesh.

The wireframe mesh is plotted using rectangles. The vertices of the rectangles [*x*, *y*] are typically the output of `meshgrid`. over a 2-D rectangular region in the x-y plane. *z* determines the height above the plane of each vertex. If only a single *z* matrix is given, then it is plotted over the meshgrid *x* = `1:columns (z)`, *y* = `1:rows (z)`. Thus, columns of *z* correspond to different *x* values and rows of *z* correspond to different *y* values.

The color of the mesh is computed by linearly scaling the *z* values to fit the range of the current colormap. Use `clim` and/or change the colormap to control the appearance.

Optionally, the color of the mesh can be specified independently of $z$ by supplying a color matrix, $c$.

Any property/value pairs are passed directly to the underlying surface object. The full list of properties is documented at Section 15.3.3.10 [Surface Properties], page 505.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value $h$ is a graphics handle to the created surface object.

**See also:** [ezmesh], page 410, [meshc], page 384, [meshz], page 384, [trimesh], page 908, [contour], page 348, [surf], page 385, [surface], page 452, [meshgrid], page 398, [hidden], page 385, [shading], page 407, [colormap], page 951, [clim], page 368.

`meshc (`*x*, *y*, *z*`)`
`meshc (`*z*`)`
`meshc (...,` *c*`)`
`meshc (...,` *prop*, *val*, `...)`
`meshc (`*hax*, `...)`
*h* `= meshc (...)`

Plot a 3-D wireframe mesh with underlying contour lines.

The wireframe mesh is plotted using rectangles. The vertices of the rectangles [x, y] are typically the output of `meshgrid`. over a 2-D rectangular region in the x-y plane. $z$ determines the height above the plane of each vertex. If only a single $z$ matrix is given, then it is plotted over the meshgrid *x* `= 1:columns (`*z*`),` *y* `= 1:rows (`*z*`)`. Thus, columns of $z$ correspond to different $x$ values and rows of $z$ correspond to different $y$ values.

The color of the mesh is computed by linearly scaling the $z$ values to fit the range of the current colormap. Use `clim` and/or change the colormap to control the appearance.

Optionally the color of the mesh can be specified independently of $z$ by supplying a color matrix, $c$.

Any property/value pairs are passed directly to the underlying surface object. The full list of properties is documented at Section 15.3.3.10 [Surface Properties], page 505.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value $h$ is a 2-element vector with a graphics handle to the created surface object and to the created contour plot.

**See also:** [ezmeshc], page 411, [mesh], page 383, [meshz], page 384, [contour], page 348, [surfc], page 386, [surface], page 452, [meshgrid], page 398, [hidden], page 385, [shading], page 407, [colormap], page 951, [clim], page 368.

`meshz (`*x*, *y*, *z*`)`
`meshz (`*z*`)`
`meshz (...,` *c*`)`
`meshz (...,` *prop*, *val*, `...)`
`meshz (`*hax*, `...)`

`h = meshz (...)`

Plot a 3-D wireframe mesh with a surrounding curtain.

The wireframe mesh is plotted using rectangles. The vertices of the rectangles [*x*, *y*] are typically the output of `meshgrid`. over a 2-D rectangular region in the x-y plane. *z* determines the height above the plane of each vertex. If only a single *z* matrix is given, then it is plotted over the meshgrid `x = 1:columns (z), y = 1:rows (z)`. Thus, columns of *z* correspond to different *x* values and rows of *z* correspond to different *y* values.

The color of the mesh is computed by linearly scaling the *z* values to fit the range of the current colormap. Use `clim` and/or change the colormap to control the appearance.

Optionally the color of the mesh can be specified independently of *z* by supplying a color matrix, *c*.

Any property/value pairs are passed directly to the underlying surface object. The full list of properties is documented at Section 15.3.3.10 [Surface Properties], page 505.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created surface object.

**See also:** [mesh], page 383, [meshc], page 384, [contour], page 348, [surf], page 385, [surface], page 452, [waterfall], page 408, [meshgrid], page 398, [hidden], page 385, [shading], page 407, [colormap], page 951, [clim], page 368.

`hidden`
`hidden` *on*
`hidden` *off*
`mode = hidden (...)`

Control mesh hidden line removal.

When called with no argument the hidden line removal state is toggled.

When called with one of the modes `"on"` or `"off"` the state is set accordingly.

The optional output argument *mode* is the current state.

Hidden Line Removal determines what graphic objects behind a mesh plot are visible. The default is for the mesh to be opaque and lines behind the mesh are not visible. If hidden line removal is turned off then objects behind the mesh can be seen through the faces (openings) of the mesh, although the mesh grid lines are still opaque.

**See also:** [mesh], page 383, [meshc], page 384, [meshz], page 384, [ezmesh], page 410, [ezmeshc], page 411, [trimesh], page 908, [waterfall], page 408.

`surf (x, y, z)`
`surf (z)`
`surf (..., c)`
`surf (..., prop, val, ...)`
`surf (hax, ...)`
`h = surf (...)`

Plot a 3-D surface mesh.

The surface mesh is plotted using shaded rectangles. The vertices of the rectangles [*x*, *y*] are typically the output of `meshgrid`. over a 2-D rectangular region in the x-y

plane. *z* determines the height above the plane of each vertex. If only a single *z* matrix is given, then it is plotted over the meshgrid `x = 1:columns (z), y = 1:rows (z)`. Thus, columns of *z* correspond to different *x* values and rows of *z* correspond to different *y* values.

The color of the surface is computed by linearly scaling the *z* values to fit the range of the current colormap. Use `clim` and/or change the colormap to control the appearance.

Optionally, the color of the surface can be specified independently of *z* by supplying a color matrix, *c*.

Any property/value pairs are passed directly to the underlying surface object. The full list of properties is documented at Section 15.3.3.10 [Surface Properties], page 505.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created surface object.

Note: The exact appearance of the surface can be controlled with the `shading` command or by using `set` to control surface object properties.

**See also:** [ezsurf], page 411, [surfc], page 386, [surfl], page 387, [surfnorm], page 388, [trisurf], page 908, [contour], page 348, [mesh], page 383, [surface], page 452, [meshgrid], page 398, [hidden], page 385, [shading], page 407, [colormap], page 951, [clim], page 368.

`surfc (x, y, z)`
`surfc (z)`
`surfc (..., c)`
`surfc (..., prop, val, ...)`
`surfc (hax, ...)`
`h = surfc (...)`

Plot a 3-D surface mesh with underlying contour lines.

The surface mesh is plotted using shaded rectangles. The vertices of the rectangles [*x*, *y*] are typically the output of `meshgrid`. over a 2-D rectangular region in the x-y plane. *z* determines the height above the plane of each vertex. If only a single *z* matrix is given, then it is plotted over the meshgrid `x = 1:columns (z), y = 1:rows (z)`. Thus, columns of *z* correspond to different *x* values and rows of *z* correspond to different *y* values.

The color of the surface is computed by linearly scaling the *z* values to fit the range of the current colormap. Use `clim` and/or change the colormap to control the appearance.

Optionally, the color of the surface can be specified independently of *z* by supplying a color matrix, *c*.

Any property/value pairs are passed directly to the underlying surface object. The full list of properties is documented at Section 15.3.3.10 [Surface Properties], page 505.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created surface object.

Note: The exact appearance of the surface can be controlled with the `shading` command or by using `set` to control surface object properties.

**See also:** [ezsurfc], page 412, [surf], page 385, [surfl], page 387, [surfnorm], page 388, [trisurf], page 908, [contour], page 348, [mesh], page 383, [surface], page 452, [meshgrid], page 398, [hidden], page 385, [shading], page 407, [colormap], page 951, [clim], page 368.

`surfl (`*z*`)`
`surfl (`*x*`, `*y*`, `*z*`)`
`surfl (...,` *lsrc*`)`
`surfl (`*x*`, `*y*`, `*z*`,` *lsrc*`,` *P*`)`
`surfl (...,` "*cdata*"`)`
`surfl (...,` "*light*"`)`
`surfl (`*hax*`, ...)`
*h* `= surfl (...)`

Plot a 3-D surface using shading based on various lighting models.

The surface mesh is plotted using shaded rectangles. The vertices of the rectangles [*x*, *y*] are typically the output of `meshgrid`. over a 2-D rectangular region in the x-y plane. *z* determines the height above the plane of each vertex. If only a single *z* matrix is given, then it is plotted over the meshgrid *x* `= 1:columns (`*z*`),` *y* `= 1:rows (`*z*`)`. Thus, columns of *z* correspond to different *x* values and rows of *z* correspond to different *y* values.

The default lighting mode `"cdata"`, changes the cdata property of the surface object to give the impression of a lighted surface.

The alternate mode `"light"` creates a light object to illuminate the surface.

The light source location may be specified using *lsrc* which can be a 2-element vector [azimuth, elevation] in degrees, or a 3-element vector [lx, ly, lz]. The default value is rotated 45 degrees counterclockwise to the current view.

The material properties of the surface can specified using a 4-element vector $P = [AM\ D\ SP\ exp]$ which defaults to $p = [0.55\ 0.6\ 0.4\ 10]$.

`"AM"` strength of ambient light
`"D"` strength of diffuse reflection
`"SP"` strength of specular reflection
`"EXP"` specular exponent

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created surface object.

Example:

```
colormap (bone (64));
surfl (peaks);
shading interp;
```

**See also:** [diffuse], page 395, [specular], page 395, [surf], page 385, [shading], page 407, [colormap], page 951, [clim], page 368.

`surfnorm (x, y, z)`
`surfnorm (z)`
`surfnorm (..., prop, val, ...)`
`surfnorm (hax, ...)`
`[Nx, Ny, Nz] = surfnorm (...)`

Find the vectors normal to a meshgridded surface.

If $x$ and $y$ are vectors, then a typical vertex is ($x$(j), $y$(i), $z$(i,j)). Thus, columns of $z$ correspond to different $x$ values and rows of $z$ correspond to different $y$ values. If only a single input $z$ is given then $x$ is taken to be `1:columns (z)` and $y$ is `1:rows (z)`.

If no return arguments are requested, a surface plot with the normal vectors to the surface is plotted.

Any property/value input pairs are assigned to the surface object. The full list of properties is documented at Section 15.3.3.10 [Surface Properties], page 505.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

If output arguments are requested then the components of the normal vectors are returned in *Nx*, *Ny*, and *Nz* and no plot is made. The normal vectors are unnormalized (magnitude != 1). To normalize, use

```
len = sqrt (nx.^2 + ny.^2 + nz.^2);
nx ./= len;  ny ./= len;  nz ./= len;
```

An example of the use of `surfnorm` is

```
surfnorm (peaks (25));
```

Algorithm: The normal vectors are calculated by taking the cross product of the diagonals of each of the quadrilateral faces in the meshgrid to find the normal vectors at the center of each face. Next, for each meshgrid point the four nearest normal vectors are averaged to obtain the final normal to the surface at the meshgrid point.

For surface objects, the `"VertexNormals"` property contains equivalent information, except possibly near the boundary of the surface where different interpolation schemes may yield slightly different values.

**See also:** [isonormals], page 390, [quiver3], page 357, [surf], page 385, [meshgrid], page 398.

`fv = isosurface (v, isoval)`
`fv = isosurface (v)`
`fv = isosurface (x, y, z, v, isoval)`
`fv = isosurface (x, y, z, v)`
`fvc = isosurface (..., col)`
`fv = isosurface (..., "noshare")`
`fv = isosurface (..., "verbose")`
`[f, v] = isosurface (...)`
`[f, v, c] = isosurface (...)`
`isosurface (...)`

Calculate isosurface of 3-D volume data.

An isosurface connects points with the same value and is analogous to a contour plot, but in three dimensions.

The input argument *v* is a three-dimensional array that contains data sampled over a volume.

The input *isoval* is a scalar that specifies the value for the isosurface. If *isoval* is omitted or empty, a "good" value for an isosurface is determined from *v*.

When called with a single output argument `isosurface` returns a structure array *fv* that contains the fields *faces* and *vertices* computed at the points `[x, y, z] = meshgrid (1:l, 1:m, 1:n)` where `[l, m, n] = size (v)`. The output *fv* can be used directly as input to the `patch` function.

If called with additional input arguments *x*, *y*, and *z* that are three-dimensional arrays with the same size as *v* or vectors with lengths corresponding to the dimensions of *v*, then the volume data is taken at the specified points. If *x*, *y*, or *z* are empty, the grid corresponds to the indices (`1:n`) in the respective direction (see [meshgrid], page 398).

The optional input argument *col*, which is a three-dimensional array of the same size as *v*, specifies coloring of the isosurface. The color data is interpolated, as necessary, to match *isoval*. The output structure array, in this case, has the additional field *facevertexcdata*.

If given the string input argument `"noshare"`, vertices may be returned multiple times for different faces. The default behavior is to eliminate vertices shared by adjacent faces.

The string input argument `"verbose"` is supported for MATLAB compatibility, but has no effect.

Any string arguments must be passed after the other arguments.

If called with two or three output arguments, return the information about the faces *f*, vertices *v*, and color data *c* as separate arrays instead of a single structure array.

If called with no output argument, the isosurface geometry is directly plotted with the `patch` command and a light object is added to the axes if not yet present.

For example,

```
[x, y, z] = meshgrid (1:5, 1:5, 1:5);
v = rand (5, 5, 5);
isosurface (x, y, z, v, .5);
```

will directly draw a random isosurface geometry in a graphics window.

An example of an isosurface geometry with different additional coloring:

```
N = 15;    # Increase number of vertices in each direction
iso = .4;  # Change isovalue to .1 to display a sphere
lin = linspace (0, 2, N);
[x, y, z] = meshgrid (lin, lin, lin);
v = abs ((x-.5).^2 + (y-.5).^2 + (z-.5).^2);
figure ();

subplot (2,2,1); view (-38, 20);
[f, vert] = isosurface (x, y, z, v, iso);
p = patch ("Faces", f, "Vertices", vert, "EdgeColor", "none");
pbaspect ([1 1 1]);
```

```
isonormals (x, y, z, v, p)
set (p, "FaceColor", "green", "FaceLighting", "gouraud");
light ("Position", [1 1 5]);

subplot (2,2,2); view (-38, 20);
p = patch ("Faces", f, "Vertices", vert, "EdgeColor", "blue");
pbaspect ([1 1 1]);
isonormals (x, y, z, v, p)
set (p, "FaceColor", "none", "EdgeLighting", "gouraud");
light ("Position", [1 1 5]);

subplot (2,2,3); view (-38, 20);
[f, vert, c] = isosurface (x, y, z, v, iso, y);
p = patch ("Faces", f, "Vertices", vert, "FaceVertexCData", c, ...
           "FaceColor", "interp", "EdgeColor", "none");
pbaspect ([1 1 1]);
isonormals (x, y, z, v, p)
set (p, "FaceLighting", "gouraud");
light ("Position", [1 1 5]);

subplot (2,2,4); view (-38, 20);
p = patch ("Faces", f, "Vertices", vert, "FaceVertexCData", c, ...
           "FaceColor", "interp", "EdgeColor", "blue");
pbaspect ([1 1 1]);
isonormals (x, y, z, v, p)
set (p, "FaceLighting", "gouraud");
light ("Position", [1 1 5]);
```

**See also:** [isonormals], page 390, [isocolors], page 392, [isocaps], page 391, [smooth3], page 392, [reducevolume], page 393, [reducepatch], page 394, [patch], page 451.

*vn* = isonormals (*val*, *vert*)
*vn* = isonormals (*val*, *hp*)
*vn* = isonormals (*x*, *y*, *z*, *val*, *vert*)
*vn* = isonormals (*x*, *y*, *z*, *val*, *hp*)
*vn* = isonormals (*...*, "*negate*")
isonormals (*val*, *hp*)
isonormals (*x*, *y*, *z*, *val*, *hp*)
isonormals (*...*, "*negate*")

Calculate normals to an isosurface.

The vertex normals *vn* are calculated from the gradient of the 3-dimensional array *val* (size: lxmxn) containing the data for an isosurface geometry. The normals point towards smaller values in *val*.

If called with one output argument *vn*, and the second input argument *vert* holds the vertices of an isosurface, then the normals *vn* are calculated at the vertices *vert* on a grid given by `[x, y, z] = meshgrid (1:l, 1:m, 1:n)`. The output argument *vn* has the same size as *vert* and can be used to set the `"VertexNormals"` property of the corresponding patch.

If called with additional input arguments *x*, *y*, and *z*, which are 3-dimensional arrays with the same size as *val*, then the volume data is taken at these points. Instead of the vertex data *vert*, a patch handle *hp* can be passed to the function.

If the last input argument is the string `"negate"`, compute the reverse vector normals of an isosurface geometry (i.e., pointed towards larger values in *val*).

If no output argument is given, the property `"VertexNormals"` of the patch associated with the patch handle *hp* is changed directly.

**See also:** [isosurface], page 388, [isocolors], page 392, [smooth3], page 392.

```
fvc = isocaps (v, isoval)
fvc = isocaps (v)
fvc = isocaps (x, y, z, v, isoval)
fvc = isocaps (x, y, z, v)
fvc = isocaps (..., which_caps)
fvc = isocaps (..., which_plane)
fvc = isocaps (..., "verbose")
[faces, vertices, fvcdata] = isocaps (...)
isocaps (...)
```

Create end-caps for isosurfaces of 3-D data.

This function places caps at the open ends of isosurfaces.

The input argument *v* is a three-dimensional array that contains data sampled over a volume.

The input *isoval* is a scalar that specifies the value for the isosurface. If *isoval* is omitted or empty, a "good" value for an isosurface is determined from *v*.

When called with a single output argument, `isocaps` returns a structure array *fvc* with the fields: `faces`, `vertices`, and `facevertexcdata`. The results are computed at the points `[x, y, z] = meshgrid (1:l, 1:m, 1:n)` where `[l, m, n] = size (v)`. The output *fvc* can be used directly as input to the `patch` function.

If called with additional input arguments *x*, *y*, and *z* that are three-dimensional arrays with the same size as *v* or vectors with lengths corresponding to the dimensions of *v*, then the volume data is taken at the specified points. If *x*, *y*, or *z* are empty, the grid corresponds to the indices (`1:n`) in the respective direction (see [meshgrid], page 398).

The optional parameter *which_caps* can have one of the following string values which defines how the data will be enclosed:

`"above"`, `"a"` (default)
: for end-caps that enclose the data above *isoval*.

`"below"`, `"b"`
: for end-caps that enclose the data below *isoval*.

The optional parameter *which_plane* can have one of the following string values to define which end-cap should be drawn:

`"all"` (default)
: for all of the end-caps.

`"xmin"`
: for end-caps at the lower x-plane of the data.

`"xmax"`
: for end-caps at the upper x-plane of the data.

`"ymin"` for end-caps at the lower y-plane of the data.

`"ymax"` for end-caps at the upper y-plane of the data.

`"zmin"` for end-caps at the lower z-plane of the data.

`"zmax"` for end-caps at the upper z-plane of the data.

The string input argument `"verbose"` is supported for MATLAB compatibility, but has no effect.

If called with two or three output arguments, the data for faces *faces*, vertices *vertices*, and the color data *facevertexcdata* are returned in separate arrays instead of a single structure.

If called with no output argument, the end-caps are drawn directly in the current figure with the `patch` command.

**See also:** [isosurface], page 388, [isonormals], page 390, [patch], page 451.

*cdat* = isocolors (*c*, *v*)
*cdat* = isocolors (*x*, *y*, *z*, *c*, *v*)
*cdat* = isocolors (*x*, *y*, *z*, *r*, *g*, *b*, *v*)
*cdat* = isocolors (*r*, *g*, *b*, *v*)
*cdat* = isocolors (..., *hp*)
isocolors (..., *hp*)

Compute isosurface colors.

If called with one output argument, and the first input argument *c* is a three-dimensional array that contains indexed color values, and the second input argument *v* are the vertices of an isosurface geometry, then return a matrix *cdat* with color data information for the geometry at computed points `[x, y, z] = meshgrid (1:l, 1:m, 1:n)`. The output argument *cdat* can be used to manually set the `"FaceVertexCData"` property of an isosurface patch object.

If called with additional input arguments *x*, *y* and *z* which are three-dimensional arrays of the same size as *c* then the color data is taken at those specified points.

Instead of indexed color data *c*, `isocolors` can also be called with RGB values *r*, *g*, *b*. If input arguments *x*, *y*, *z* are not given then `meshgrid` computed values are used.

Optionally, a patch handle *hp* can be given as the last input argument to all function call variations and the vertex data will be extracted from the isosurface patch object. Finally, if no output argument is given then the colors of the patch given by the patch handle *hp* are changed.

**See also:** [isosurface], page 388, [isonormals], page 390.

*smoothed_data* = smooth3 (*data*)
*smoothed_data* = smooth3 (*data*, *method*)
*smoothed_data* = smooth3 (*data*, *method*, *sz*)
*smoothed_data* = smooth3 (*data*, *method*, *sz*, *std_dev*)

Smooth values of 3-dimensional matrix *data*.

This function may be used, for example, to reduce the impact of noise in *data* before calculating isosurfaces.

*data* must be a non-singleton 3-dimensional matrix. The output *smoothed_data* is a matrix of the same size as *data*.

The option input *method* determines which convolution kernel is used for the smoothing process. Possible choices:

`"box"`, `"b"` (default)
: a convolution kernel with sharp edges.

`"gaussian"`, `"g"`
: a convolution kernel that is represented by a non-correlated trivariate normal distribution function.

*sz* is either a 3-element vector specifying the size of the convolution kernel in the x-, y- and z-directions, or a scalar. In the scalar case the same size is used for all three dimensions (`[sz, sz, sz]`). The default value is 3.

If *method* is `"gaussian"` then the optional input *std_dev* defines the standard deviation of the trivariate normal distribution function. *std_dev* is either a 3-element vector specifying the standard deviation of the Gaussian convolution kernel in x-, y- and z-directions, or a scalar. In the scalar case the same value is used for all three dimensions. The default value is 0.65.

**See also:** [isosurface], page 388, [isonormals], page 390, [patch], page 451.

```
[nx, ny, nz, nv] = reducevolume (v, r)
[nx, ny, nz, nv] = reducevolume (x, y, z, v, r)
nv = reducevolume (...)
```

Reduce the volume of the dataset in *v* according to the values in *r*.

*v* is a matrix that is non-singleton in the first 3 dimensions.

*r* can either be a vector of 3 elements representing the reduction factors in the x-, y-, and z-directions or a scalar, in which case the same reduction factor is used in all three dimensions.

`reducevolume` reduces the number of elements of *v* by taking only every *r*-th element in the respective dimension.

Optionally, *x*, *y*, and *z* can be supplied to represent the set of coordinates of *v*. They can either be matrices of the same size as *v* or vectors with sizes according to the dimensions of *v*, in which case they are expanded to matrices (see `meshgrid`, page 398).

If `reducevolume` is called with two arguments then *x*, *y*, and *z* are assumed to match the respective indices of *v*.

The reduced matrix is returned in *nv*.

Optionally, the reduced set of coordinates are returned in *nx*, *ny*, and *nz*, respectively.

Examples:

```
v = reshape (1:6*8*4, [6 8 4]);
nv = reducevolume (v, [4 3 2]);
```

```
v = reshape (1:6*8*4, [6 8 4]);
x = 1:3:24;  y = -14:5:11;  z = linspace (16, 18, 4);
[nx, ny, nz, nv] = reducevolume (x, y, z, v, [4 3 2]);
```

**See also:** [isosurface], page 388, [isonormals], page 390.

*reduced_fv* = reducepatch (*fv*)
*reduced_fv* = reducepatch (*faces*, *vertices*)
*reduced_fv* = reducepatch (*patch_handle*)
reducepatch (*patch_handle*)
*reduced_fv* = reducepatch (..., *reduction_factor*)
*reduced_fv* = reducepatch (..., "*fast*")
*reduced_fv* = reducepatch (..., "*verbose*")
[*reduced_faces*, *reduces_vertices*] = reducepatch (...)

Reduce the number of faces and vertices in a patch object while retaining the overall shape of the patch.

The input patch can be represented by a structure *fv* with the fields `faces` and `vertices`, by two matrices *faces* and *vertices* (see, e.g., the result of `isosurface`), or by a handle to a patch object *patch_handle* (see [patch], page 451).

The number of faces and vertices in the patch is reduced by iteratively collapsing the shortest edge of the patch to its midpoint (as discussed, e.g., here: https://libigl.github.io/libigl/tutorial/tutorial.html#meshdecimation).

Currently, only patches consisting of triangles are supported. The resulting patch also consists only of triangles.

If `reducepatch` is called with a handle to a valid patch *patch_handle*, and without any output arguments, then the given patch is updated immediately.

If the *reduction_factor* is omitted, the resulting structure *reduced_fv* includes approximately 50% of the faces of the original patch. If *reduction_factor* is a fraction between 0 (excluded) and 1 (excluded), a patch with approximately the corresponding fraction of faces is determined. If *reduction_factor* is an integer greater than or equal to 1, the resulting patch has approximately *reduction_factor* faces. Depending on the geometry of the patch, the resulting number of faces can differ from the given value of *reduction_factor*. This is especially true when many shared vertices are detected.

For the reduction, it is necessary that vertices of touching faces are shared. Shared vertices are detected automatically. This detection can be skipped by passing the optional string argument `"fast"`.

With the optional string arguments `"verbose"`, additional status messages are printed to the command window.

Any string input arguments must be passed after all other arguments.

If called with one output argument, the reduced faces and vertices are returned in a structure *reduced_fv* with the fields `faces` and `vertices` (see the one output option of `isosurface`).

If called with two output arguments, the reduced faces and vertices are returned in two separate matrices *reduced_faces* and *reduced_vertices*.

**See also:** [isosurface], page 388, [isonormals], page 390, [reducevolume], page 393, [patch], page 451.

shrinkfaces (*p*, *sf*)
*nfv* = shrinkfaces (*p*, *sf*)
*nfv* = shrinkfaces (*fv*, *sf*)
*nfv* = shrinkfaces (*f*, *v*, *sf*)

`[nf, nv] = shrinkfaces (...)`

Reduce the size of faces in a patch by the shrink factor *sf*.

The patch object can be specified by a graphics handle (*p*), a patch structure (*fv*) with the fields `"faces"` and `"vertices"`, or as two separate matrices (*f*, *v*) of faces and vertices.

The shrink factor *sf* is a positive number specifying the percentage of the original area the new face will occupy. If no factor is given the default is 0.3 (a reduction to 30% of the original size). A factor greater than 1.0 will result in the expansion of faces.

Given a patch handle as the first input argument and no output parameters, perform the shrinking of the patch faces in place and redraw the patch.

If called with one output argument, return a structure with fields `"faces"`, `"vertices"`, and `"facevertexcdata"` containing the data after shrinking. This structure can be used directly as an input argument to the `patch` function.

**Caution:**: Performing the shrink operation on faces which are not convex can lead to undesirable results.

Example: a triangulated 3/4 circle and the corresponding shrunken version.

```
[phi r] = meshgrid (linspace (0, 1.5*pi, 16), linspace (1, 2, 4));
tri = delaunay (phi(:), r(:));
v = [r(:).*sin(phi(:)) r(:).*cos(phi(:))];
clf ()
p = patch ("Faces", tri, "Vertices", v, "FaceColor", "none");
fv = shrinkfaces (p);
patch (fv)
axis equal
grid on
```

**See also:** [patch], page 451.

`d = diffuse (sx, sy, sz, lv)`

Calculate the diffuse reflection strength of a surface defined by the normal vector elements *sx*, *sy*, *sz*.

The light source location vector *lv* can be given as a 2-element vector [azimuth, elevation] in degrees or as a 3-element vector [x, y, z].

**See also:** [specular], page 395, [surfl], page 387.

`refl = specular (sx, sy, sz, lv, vv)`
`refl = specular (sx, sy, sz, lv, vv, se)`

Calculate the specular reflection strength of a surface defined by the normal vector elements *sx*, *sy*, *sz* using Phong's approximation.

The light source location and viewer location vectors are specified using parameters *lv* and *vv* respectively. The location vectors can given as 2-element vectors [azimuth, elevation] in degrees or as 3-element vectors [x, y, z].

An optional sixth argument specifies the specular exponent (spread) *se*. If not given, *se* defaults to 10.

**See also:** [diffuse], page 395, [surfl], page 387.

`lighting` (*type*)
`lighting` (*hax*, *type*)

Set the lighting of patch or surface graphic objects.

Valid arguments for *type* are

`"flat"` Draw objects with faceted lighting effects.

`"gouraud"`
Draw objects with linear interpolation of the lighting effects between the vertices.

`"none"` Draw objects without light and shadow effects.

If the first argument *hax* is an axes handle, then change the lighting effects of objects in this axes, rather than the current axes returned by `gca`.

The lighting effects are only visible if at least one light object is present and visible in the same axes.

**See also:** [light], page 453, [fill], page 364, [mesh], page 383, [patch], page 451, [pcolor], page 363, [surf], page 385, [surface], page 452, [shading], page 407.

`material` *shiny*
`material` *dull*
`material` *metal*
`material` *default*
`material` ([*as*, *ds*, *ss*])
`material` ([*as*, *ds*, *ss*, *se*])
`material` ([*as*, *ds*, *ss*, *se*, *scr*])
`material` (*hlist*, ...)
*mtypes* = `material` ()
*refl_props* = `material` (*mtype_string*)

Set reflectance properties for the lighting of surfaces and patches.

This function changes the ambient, diffuse, and specular strengths, as well as the specular exponent and specular color reflectance, of all `patch` and `surface` objects in the current axes. This can be used to simulate, to some extent, the reflectance properties of certain materials when used with `light`.

When called with a string, the aforementioned properties are set according to the values in the following table:

| *mtype* | **ambient-strength** | **diffuse-strength** | **specular-strength** | **specular-exponent** | **specular-color-reflectance** |
|---|---|---|---|---|---|
| `"shiny"` | 0.3 | 0.6 | 0.9 | 20 | 1.0 |
| `"dull"` | 0.3 | 0.8 | 0.0 | 10 | 1.0 |
| `"metal"` | 0.3 | 0.3 | 1.0 | 25 | 0.5 |
| `"default"` | `"default"` | `"default"` | `"default"` | `"default"` | `"default"` |

When called with a vector of three elements, the ambient, diffuse, and specular strengths of all `patch` and `surface` objects in the current axes are updated. An

optional fourth vector element updates the specular exponent, and an optional fifth vector element updates the specular color reflectance.

A list of graphic handles can also be passed as the first argument. In this case, the properties of these handles and all child `patch` and `surface` objects will be updated.

Additionally, `material` can be called with a single output argument. If called without input arguments, a column cell vector *mtypes* with the strings for all available materials is returned. If the one input argument *mtype_string* is the name of a material, a 1x5 cell vector *refl_props* with the reflectance properties of that material is returned. In both cases, no graphic properties are changed.

**See also:** [light], page 453, [fill], page 364, [mesh], page 383, [patch], page 451, [pcolor], page 363, [surf], page 385, [surface], page 452.

`camlight`
`camlight` *right*
`camlight` *left*
`camlight` *headlight*
`camlight` (*az*, *el*)
`camlight` (*...*, *style*)
`camlight` (*hl*, *...*)
`camlight` (*hax*, *...*)
*h* = `camlight` (*...*)

Add a light object to a figure using a simple interface.

When called with no arguments, a light object is added to the current plot and is placed slightly above and to the right of the camera's current position: this is equivalent to `camlight right`. The commands `camlight left` and `camlight headlight` behave similarly with the placement being either left of the camera position or centered on the camera position.

For more control, the light position can be specified by an azimuthal rotation *az* and an elevation angle *el*, both in degrees, relative to the current properties of the camera.

The optional string *style* specifies whether the light is a local point source (`"local"`, the default) or placed at infinite distance (`"infinite"`).

If the first argument *hl* is a handle to a light object, then act on this light object rather than creating a new object.

If the first argument *hax* is an axes handle, then create a new light object in this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the light object. This can be used to move or further change properties of the light object.

Examples:

Add a light object to a plot

```
sphere (36);
camlight
```

Position the light source exactly

```
camlight (45, 30);
```

Here the light is first pitched upwards (see [camup], page 404) from the camera position (see [campos], page 401) by 30 degrees. It is then yawed by 45 degrees to the right. Both rotations are centered around the camera target (see [camtarget], page 403).

Return a handle to further manipulate the light object

```
clf
sphere (36);
hl = camlight ("left");
set (hl, "color", "r");
```

**See also:** [light], page 453.

`lightangle (az, el)`
`lightangle (hax, az, el)`
`lightangle (hl, az, el)`
`hl = lightangle (...)`
`[az, el] = lightangle (hl)`

Add a light object to the current axes using spherical coordinates.

The light position is specified by an azimuthal rotation *az* and an elevation angle *el*, both in degrees.

If the first argument *hax* is an axes handle, then create a new light object in this axes, rather than the current axes returned by `gca`.

If the first argument *hl* is a handle to a light object, then act on this light object rather than creating a new object.

The optional return value *hl* is a graphics handle to the light object.

Example:

Add a light object to a plot

```
clf;
sphere (36);
lightangle (45, 30);
```

**See also:** [light], page 453, [view], page 400, [camlight], page 397.

`[xx, yy] = meshgrid (x, y)`
`[xx, yy, zz] = meshgrid (x, y, z)`
`[xx, yy] = meshgrid (x)`
`[xx, yy, zz] = meshgrid (x)`

Given vectors of *x* and *y* coordinates, return matrices *xx* and *yy* corresponding to a full 2-D grid.

The rows of *xx* are copies of *x*, and the columns of *yy* are copies of *y*. If *y* is omitted, then it is assumed to be the same as *x*.

If the optional *z* input is given, or *zz* is requested, then the output will be a full 3-D grid. If *z* is omitted and *zz* is requested, it is assumed to be the same as *y*.

`meshgrid` is most frequently used to produce input for a 2-D or 3-D function that will be plotted. The following example creates a surface plot of the "sombrero" function.

```
f = @(x,y) sin (sqrt (x.^2 + y.^2)) ./ sqrt (x.^2 + y.^2);
range = linspace (-8, 8, 41);
[X, Y] = meshgrid (range, range);
Z = f (X, Y);
surf (X, Y, Z);
```

Programming Note: `meshgrid` is restricted to 2-D or 3-D grid generation. The `ndgrid` function will generate 1-D through N-D grids. However, the functions are not completely equivalent. If $x$ is a vector of length M and $y$ is a vector of length N, then `meshgrid` will produce an output grid which is NxM. `ndgrid` will produce an output which is MxN (transpose) for the same input. Some core functions expect `meshgrid` input and others expect `ndgrid` input. Check the documentation for the function in question to determine the proper input format.

**See also:** [ndgrid], page 399, [mesh], page 383, [contour], page 348, [surf], page 385.

`[`*y1*`, `*y2*`, ..., `*yn*`] = ndgrid (`*x1*`, `*x2*`, ..., `*xn*`)`
`[`*y1*`, `*y2*`, ..., `*yn*`] = ndgrid (`*x*`)`

Given n vectors *x1*, ..., xn, `ndgrid` returns n arrays of dimension n.

The elements of the i-th output argument contains the elements of the vector *x*i repeated over all dimensions different from the i-th dimension. Calling ndgrid with only one input argument $x$ is equivalent to calling ndgrid with all n input arguments equal to $x$:

[*y1*, *y2*, ..., *y*n] = ndgrid (*x*, ..., *x*)

Programming Note: `ndgrid` is very similar to the function `meshgrid` except that the first two dimensions are transposed in comparison to `meshgrid`. Some core functions expect `meshgrid` input and others expect `ndgrid` input. Check the documentation for the function in question to determine the proper input format.

**See also:** [meshgrid], page 398.

`plot3 (`*x*, *y*, *z*`)`
`plot3 (`*x*, *y*, *z*, *prop*, *value*, ...`)`
`plot3 (`*x*, *y*, *z*, *fmt*`)`
`plot3 (`*x*, *cplx*`)`
`plot3 (`*cplx*`)`
`plot3 (`*hax*, ...`)`
*h* `= plot3 (...)`

Produce 3-D plots.

Many different combinations of arguments are possible. The simplest form is

```
plot3 (x, y, z)
```

in which the arguments are taken to be the vertices of the points to be plotted in three dimensions. If all arguments are vectors of the same length, then a single continuous line is drawn. If all arguments are matrices, then each column of is treated as a separate line. No attempt is made to transpose the arguments to make the number of rows match.

If only two arguments are given, as

```
plot3 (x, cplx)
```

the real and imaginary parts of the second argument are used as the $y$ and $z$ coordinates, respectively.

If only one argument is given, as

```
plot3 (cplx)
```

the real and imaginary parts of the argument are used as the $y$ and $z$ values, and they are plotted versus their index.

Arguments may also be given in groups of three as

```
plot3 (x1, y1, z1, x2, y2, z2, ...)
```

in which each set of three arguments is treated as a separate line or set of lines in three dimensions.

To plot multiple one- or two-argument groups, separate each group with an empty format string, as

```
plot3 (x1, c1, "", c2, "", ...)
```

Multiple property-value pairs may be specified which will affect the line objects drawn by `plot3`. If the *fmt* argument is supplied it will format the line objects in the same manner as `plot`. The full list of properties is documented at Section 15.3.3.5 [Line Properties], page 485.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created plot.

Example:

```
z = [0:0.05:5];
plot3 (cos (2*pi*z), sin (2*pi*z), z, ";helix;");
plot3 (z, exp (2i*pi*z), ";complex sinusoid;");
```

**See also:** [ezplot3], page 409, [plot], page 334.

```
view (azimuth, elevation)
view ([azimuth elevation])
view ([x y z])
view (2)
view (3)
view (hax, ...)
[azimuth, elevation] = view ()
```

Query or set the viewpoint for the current axes.

The parameters *azimuth* and *elevation* can be given as two arguments or as 2-element vector. The viewpoint can also be specified with Cartesian coordinates $x$, $y$, and $z$.

The call `view (2)` sets the viewpoint to *azimuth* = 0 and *elevation* = 90, which is the default for 2-D graphs.

The call `view (3)` sets the viewpoint to *azimuth* = -37.5 and *elevation* = 30, which is the default for 3-D graphs.

If the first argument *hax* is an axes handle, then operate on this axes rather than the current axes returned by `gca`.

If no inputs are given, return the current *azimuth* and *elevation*.

`camlookat ()`
`camlookat (`*`h`*`)`
`camlookat (`*`handle_list`*`)`
`camlookat (`*`hax`*`)`

Move the camera and adjust its properties to look at objects.

When the input is a handle *h*, the camera is set to point toward the center of the bounding box of *h*. The camera's position is adjusted so the bounding box approximately fills the field of view.

This command fixes the camera's viewing direction (`camtarget() - campos()`), camera up vector (see [camup], page 404) and viewing angle (see [camva], page 404). The camera target (see [camtarget], page 403) and camera position (see [campos], page 401) are changed.

If the argument is a list *handle_list*, then a single bounding box for all the objects is computed and the camera is then adjusted as above.

If the argument is an axis object *hax*, then the children of the axis are used as *handle_list*. When called with no inputs, it uses the current axis (see [gca], page 455).

**See also:** [camorbit], page 402, [camzoom], page 405, [camroll], page 402.

*`p`* `= campos ()`
`campos ([`*`x y z`*`])`
*`mode`* `= campos ("`*`mode`*`")`
`campos (`*`mode`*`)`
`campos (`*`hax`*`, ...)`

Get or set the camera position.

The default camera position is determined automatically based on the scene. For example, to get the camera position:

```
hf = figure();
peaks()
p = campos ()
  ⇒ p =
      -27.394  -35.701   64.079
```

We can then move the camera further up the z-axis:

```
campos (p + [0 0 10])
campos ()
  ⇒ ans =
      -27.394  -35.701   74.079
```

Having made that change, the camera position *mode* is now manual:

```
campos ("mode")
  ⇒ manual
```

We can set it back to automatic:

```
campos ("auto")
campos ()
  ⇒ ans =
      -27.394  -35.701   64.079
close (hf)
```

By default, these commands affect the current axis; alternatively, an axis can be specified by the optional argument *hax*.

**See also:** [camup], page 404, [camtarget], page 403, [camva], page 404.

`camorbit (`*`theta`*`, `*`phi`*`)`
`camorbit (`*`theta`*`, `*`phi`*`, `*`coorsys`*`)`
`camorbit (`*`theta`*`, `*`phi`*`, `*`coorsys`*`, `*`dir`*`)`
`camorbit (`*`theta`*`, `*`phi`*`, "`*`data`*`")`
`camorbit (`*`theta`*`, `*`phi`*`, "`*`data`*`", "`*`z`*`")`
`camorbit (`*`theta`*`, `*`phi`*`, "`*`data`*`", "`*`x`*`")`
`camorbit (`*`theta`*`, `*`phi`*`, "`*`data`*`", "`*`y`*`")`
`camorbit (`*`theta`*`, `*`phi`*`, "`*`data`*`", [`*`x y z`*`])`
`camorbit (`*`theta`*`, `*`phi`*`, "`*`camera`*`")`
`camorbit (`*`hax`*`, ...)`

Rotate the camera up/down and left/right around its target.

Move the camera *phi* degrees up and *theta* degrees to the right, as if it were in an orbit around its target. Example:

```
sphere ()
camorbit (30, 20)
```

These rotations are centered around the camera target (see [camtarget], page 403). First the camera position is pitched up or down by rotating it *phi* degrees around an axis orthogonal to both the viewing direction (specifically `camtarget() - campos()`) and the camera "up vector" (see [camup], page 404). Example:

```
camorbit (0, 20)
```

The second rotation depends on the coordinate system *coorsys* and direction *dir* inputs. The default for *coorsys* is `"data"`. In this case, the camera is yawed left or right by rotating it *theta* degrees around an axis specified by *dir*. The default for *dir* is `"z"`, corresponding to the vector `[0, 0, 1]`. Example:

```
camorbit (30, 0)
```

When *coorsys* is set to `"camera"`, the camera is moved left or right by rotating it around an axis parallel to the camera up vector (see [camup], page 404). The input *dir* should not be specified in this case. Example:

```
camorbit (30, 0, "camera")
```

(Note: the rotation by *phi* is unaffected by `"camera"`.)

The `camorbit` command modifies two camera properties: [campos], page 401, and [camup], page 404.

By default, this command affects the current axis; alternatively, an axis can be specified by the optional argument *hax*.

**See also:** [camzoom], page 405, [camroll], page 402, [camlookat], page 401.

`camroll (`*`theta`*`)`
`camroll (`*`hax`*`, `*`theta`*`)`

Roll the camera.

Roll the camera clockwise by *theta* degrees. For example, the following command will roll the camera by 30 degrees clockwise (to the right); this will cause the scene to appear to roll by 30 degrees to the left:

```
peaks ()
camroll (30)
```

Roll the camera back:

```
camroll (-30)
```

The following command restores the default camera roll:

```
camup ("auto")
```

By default, these commands affect the current axis; alternatively, an axis can be specified by the optional argument *hax*.

**See also:** [camzoom], page 405, [camorbit], page 402, [camlookat], page 401, [camup], page 404.

*t* = camtarget ()
camtarget ([*x* *y* *z*])
*mode* = camtarget ("*mode*")
camtarget (*mode*)
camtarget (*hax*, ...)

Get or set where the camera is pointed.

The camera target is a point in space where the camera is pointing. Usually, it is determined automatically based on the scene:

```
hf = figure();
sphere (36)
v = camtarget ()
  ⇒ v =
      0   0   0
```

We can turn the camera to point at a new target:

```
camtarget ([1 1 1])
camtarget ()
  ⇒   1   1   1
```

Having done so, the camera target *mode* is manual:

```
camtarget ("mode")
  ⇒ manual
```

This means, for example, adding new objects to the scene will not retarget the camera:

```
hold on;
peaks ()
camtarget ()
  ⇒   1   1   1
```

We can reset it to be automatic:

```
camtarget ("auto")
camtarget ()
  ⇒   0   0   0.76426
close (hf)
```

By default, these commands affect the current axis; alternatively, an axis can be specified by the optional argument *hax*.

**See also:** [campos], page 401, [camup], page 404, [camva], page 404.

*up* = camup ()
camup ([*x* *y* *z*])
*mode* = camup ("*mode*")
camup (*mode*)
camup (*hax*, ...)

Get or set the camera up vector.

By default, the camera is oriented so that "up" corresponds to the positive z-axis:

```
hf = figure ();
sphere (36)
v = camup ()
  ⇒ v =
      0   0   1
```

Specifying a new "up vector" rolls the camera and sets the mode to manual:

```
camup ([1 1 0])
camup ()
  ⇒   1   1   0
camup ("mode")
  ⇒ manual
```

Modifying the up vector does not modify the camera target (see [camtarget], page 403). Thus, the camera up vector might not be orthogonal to the direction of the camera's view:

```
camup ([1 2 3])
dot (camup (), camtarget () - campos ())
  ⇒ 6...
```

A consequence is that "pulling back" on the up vector does not pitch the camera view (as that would require changing the target). Setting the up vector is thus typically used only to roll the camera. A more intuitive command for this purpose is [camroll], page 402.

Finally, we can reset the up vector to automatic mode:

```
camup ("auto")
camup ()
  ⇒   0   0   1
close (hf)
```

By default, these commands affect the current axis; alternatively, an axis can be specified by the optional argument *hax*.

**See also:** [campos], page 401, [camtarget], page 403, [camva], page 404.

*a* = camva ()
camva (*a*)
*mode* = camva ("*mode*")
camva (*mode*)

`camva (`*hax*`, ...)`

Get or set the camera viewing angle.

The camera has a viewing angle which determines how much can be seen. By default this is:

```
hf = figure();
sphere (36)
a = camva ()
  ⇒ a =  10.340
```

To get a wider-angle view, we could double the viewing angle. This will also set the mode to manual:

```
camva (2*a)
camva ("mode")
  ⇒ manual
```

We can set it back to automatic:

```
camva ("auto")
camva ("mode")
  ⇒ auto
camva ()
  ⇒ ans =  10.340
close (hf)
```

By default, these commands affect the current axis; alternatively, an axis can be specified by the optional argument *hax*.

**See also:** [campos], page 401, [camtarget], page 403, [camup], page 404.

`camzoom (`*zf*`)`
`camzoom (`*hax*`, `*zf*`)`

Zoom the camera in or out.

A value of *zf* larger than 1 "zooms in" such that the scene appears magnified:

```
hf = figure ();
sphere (36)
camzoom (1.2)
```

A value smaller than 1 "zooms out" so the camera can see more of the scene:

```
camzoom (0.5)
```

Technically speaking, zooming affects the "viewing angle". The following command resets to the default zoom:

```
camva ("auto")
close (hf)
```

By default, these commands affect the current axis; alternatively, an axis can be specified by the optional argument *hax*.

**See also:** [camroll], page 402, [camorbit], page 402, [camlookat], page 401, [camva], page 404.

`slice (`*x*, *y*, *z*, *v*, *sx*, *sy*, *sz*`)`
`slice (`*x*, *y*, *z*, *v*, *xi*, *yi*, *zi*`)`
`slice (`*v*, *sx*, *sy*, *sz*`)`
`slice (`*v*, *xi*, *yi*, *zi*`)`
`slice (...,` *method*`)`
`slice (`*hax*, `...)`
*h* `= slice (...)`

Plot slices of 3-D data/scalar fields.

Each element of the 3-dimensional array *v* represents a scalar value at a location given by the parameters *x*, *y*, and *z*. The parameters *x*, *y*, and *z* are either 3-dimensional arrays of the same size as the array *v* in the `"meshgrid"` format or vectors. The parameters *xi*, etc. respect a similar format to *x*, etc., and they represent the points at which the array *vi* is interpolated using interp3. The vectors *sx*, *sy*, and *sz* contain points of orthogonal slices of the respective axes.

If *x*, *y*, *z* are omitted, they are assumed to be `x = 1:size (v, 2)`, `y = 1:size (v, 1)` and `z = 1:size (v, 3)`.

*method* is one of:

`"nearest"`
: Return the nearest neighbor.

`"linear"`
: Linear interpolation from nearest neighbors.

`"cubic"`
: Cubic interpolation from four nearest neighbors (not implemented yet).

`"spline"`
: Cubic spline interpolation—smooth first and second derivatives throughout the curve.

The default method is `"linear"`.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created surface object.

Examples:

```
[x, y, z] = meshgrid (linspace (-8, 8, 32));
v = sin (sqrt (x.^2 + y.^2 + z.^2)) ./ (sqrt (x.^2 + y.^2 + z.^2));
slice (x, y, z, v, [], 0, []);

[xi, yi] = meshgrid (linspace (-7, 7));
zi = xi + yi;
slice (x, y, z, v, xi, yi, zi);
```

**See also:** [interp3], page 900, [surface], page 452, [pcolor], page 363.

`ribbon (`*y*`)`
`ribbon (`*x*, *y*`)`
`ribbon (`*x*, *y*, *width*`)`
`ribbon (`*hax*, `...)`
*h* `= ribbon (...)`

Draw a ribbon plot for the columns of *y* vs. *x*.

If *x* is omitted, a vector containing the row numbers is assumed (`1:rows (Y)`). Alternatively, *x* can also be a vector with same number of elements as rows of *y* in which case the same *x* is used for each column of *y*.

The optional parameter *width* specifies the width of a single ribbon (default is 0.75).

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a vector of graphics handles to the surface objects representing each ribbon.

**See also:** [surface], page 452, [waterfall], page 408.

`shading (`*`type`*`)`
`shading (`*`hax`*`, `*`type`*`)`

Set the shading of patch or surface graphic objects.

Valid arguments for *type* are

`"flat"` Single colored patches with invisible edges.

`"faceted"`
Single colored patches with black edges.

`"interp"` Colors between patch vertices are interpolated and the patch edges are invisible.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

**See also:** [fill], page 364, [mesh], page 383, [patch], page 451, [pcolor], page 363, [surf], page 385, [surface], page 452, [hidden], page 385, [lighting], page 396.

`scatter3 (`*`x`*`, `*`y`*`, `*`z`*`)`
`scatter3 (`*`x`*`, `*`y`*`, `*`z`*`, `*`s`*`)`
`scatter3 (`*`x`*`, `*`y`*`, `*`z`*`, `*`s`*`, `*`c`*`)`
`scatter3 (...,` *`style`*`)`
`scatter3 (..., "`*`filled`*`")`
`scatter3 (...,` *`prop`*`, `*`val`*`)`
`scatter3 (`*`hax`*`, ...)`
*`h`* `= scatter3 (...)`

Draw a 3-D scatter plot.

A marker is plotted at each point defined by the coordinates in the vectors *x*, *y*, and *z*.

The size of the markers is determined by *s*, which can be a scalar or a vector of the same length as *x*, *y*, and *z*. If *s* is not given, or is an empty matrix, then a default value of 8 points is used.

The color of the markers is determined by *c*, which can be a string defining a fixed color; a 3-element vector giving the red, green, and blue components of the color; a vector of the same length as *x* that gives a scaled index into the current colormap; or an Nx3 matrix defining the RGB color of each marker individually.

The marker to use can be changed with the *style* argument, that is a string defining a marker in the same manner as the `plot` command. If no marker is specified it defaults to `"o"` or circles. If the argument `"filled"` is given then the markers are filled.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the scatter object representing the points.

```
[x, y, z] = peaks (20);
scatter3 (x(:), y(:), z(:), [], z(:));
```

Programming Note: The full list of properties is documented at Section 15.3.3.9 [Scatter Properties], page 501.

**See also:** [scatter], page 345, [patch], page 451, [plot], page 334.

`waterfall (`*x*, *y*, *z*`)`
`waterfall (`*z*`)`
`waterfall (...,` *c*`)`
`waterfall (...,` *prop*, *val*, `...)`
`waterfall (`*hax*, `...)`
*h* `= waterfall (...)`

Plot a 3-D waterfall plot.

A waterfall plot is similar to a `meshz` plot except only mesh lines for the rows of *z* (x-values) are shown.

The wireframe mesh is plotted using rectangles. The vertices of the rectangles [*x*, *y*] are typically the output of `meshgrid`. over a 2-D rectangular region in the x-y plane. *z* determines the height above the plane of each vertex. If only a single *z* matrix is given, then it is plotted over the meshgrid `x = 1:columns (z), y = 1:rows (z)`. Thus, columns of *z* correspond to different *x* values and rows of *z* correspond to different *y* values.

The color of the mesh is computed by linearly scaling the *z* values to fit the range of the current colormap. Use `clim` and/or change the colormap to control the appearance.

Optionally the color of the mesh can be specified independently of *z* by supplying a color matrix, *c*.

Any property/value pairs are passed directly to the underlying surface object. The full list of properties is documented at Section 15.3.3.10 [Surface Properties], page 505.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created surface object.

**See also:** [meshz], page 384, [mesh], page 383, [meshc], page 384, [contour], page 348, [surf], page 385, [surface], page 452, [ribbon], page 406, [meshgrid], page 398, [hidden], page 385, [shading], page 407, [colormap], page 951, [clim], page 368.

### 15.2.2.1 Aspect Ratio

For three-dimensional plots the aspect ratio can be set for data with `daspect` and for the plot box with `pbaspect`. See Section 15.2.1.1 [Axis Configuration], page 366, for controlling the x-, y-, and z-limits for plotting.

*data_aspect_ratio* = daspect ()
daspect (*data_aspect_ratio*)
daspect (*mode*)
*data_aspect_ratio_mode* = daspect ("*mode*")
daspect (*hax*, ...)

Query or set the data aspect ratio of the current axes.

The aspect ratio is a normalized 3-element vector representing the span of the x, y, and z-axis limits.

daspect (*mode*)

Set the data aspect ratio mode of the current axes. *mode* is either "auto" or "manual".

daspect ("mode")

Return the data aspect ratio mode of the current axes.

daspect (*hax*, ...)

Operate on the axes in handle *hax* instead of the current axes.

**See also:** [axis], page 366, [pbaspect], page 409, [xlim], page 369, [ylim], page 370, [zlim], page 371.

*plot_box_aspect_ratio* = pbaspect ( )
pbaspect (*plot_box_aspect_ratio*)
pbaspect (*mode*)
*plot_box_aspect_ratio_mode* = pbaspect ("*mode*")
pbaspect (*hax*, ...)

Query or set the plot box aspect ratio of the current axes.

The aspect ratio is a normalized 3-element vector representing the rendered lengths of the x, y, and z axes.

pbaspect(*mode*)

Set the plot box aspect ratio mode of the current axes. *mode* is either "auto" or "manual".

pbaspect ("mode")

Return the plot box aspect ratio mode of the current axes.

pbaspect (*hax*, ...)

Operate on the axes in handle *hax* instead of the current axes.

**See also:** [axis], page 366, [daspect], page 408, [xlim], page 369, [ylim], page 370, [zlim], page 371.

### 15.2.2.2 Three-dimensional Function Plotting

ezplot3 (*fx*, *fy*, *fz*)
ezplot3 (..., *dom*)
ezplot3 (..., *n*)
ezplot3 (..., "*animate*")
ezplot3 (*hax*, ...)

`h = ezplot3 (...)`

Plot a parametrically defined curve in three dimensions.

*fx*, *fy*, and *fz* are strings, inline functions, or function handles with one argument defining the function. By default the plot is over the domain `0 <= t <= 2*pi` with 500 points.

If *dom* is a two element vector, it represents the minimum and maximum values of *t*.

*n* is a scalar defining the number of points to use in plotting the function.

If the `"animate"` option is given then the plotting is animated in the style of `comet3`.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created plot.

```
fx = @(t) cos (t);
fy = @(t) sin (t);
fz = @(t) t;
ezplot3 (fx, fy, fz, [0, 10*pi], 100);
```

**See also:** [plot3], page 399, [comet3], page 366, [ezplot], page 379, [ezmesh], page 410, [ezsurf], page 411.

`ezmesh (f)`
`ezmesh (fx, fy, fz)`
`ezmesh (..., dom)`
`ezmesh (..., n)`
`ezmesh (..., "circ")`
`ezmesh (hax, ...)`
`h = ezmesh (...)`

Plot the mesh defined by a function.

*f* is a string, inline function, or function handle with two arguments defining the function. By default the plot is over the meshed domain `-2*pi <= x | y <= 2*pi` with 60 points in each dimension.

If three functions are passed, then plot the parametrically defined function `[fx(s, t), fy(s, t), fz(s, t)]`.

If *dom* is a two element vector, it represents the minimum and maximum values of both *x* and *y*. If *dom* is a four element vector, then the minimum and maximum values are `[xmin xmax ymin ymax]`.

*n* is a scalar defining the number of points to use in each dimension.

If the argument `"circ"` is given, then the function is plotted over a disk centered on the middle of the domain *dom*.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created surface object.

Example 1: 2-argument function

```
f = @(x,y) sqrt (abs (x .* y)) ./ (1 + x.^2 + y.^2);
ezmesh (f, [-3, 3]);
```

Example 2: parametrically defined function

```
fx = @(s,t) cos (s) .* cos (t);
fy = @(s,t) sin (s) .* cos (t);
fz = @(s,t) sin (t);
ezmesh (fx, fy, fz, [-pi, pi, -pi/2, pi/2], 20);
```

**See also:** [mesh], page 383, [ezmeshc], page 411, [ezplot], page 379, [ezsurf], page 411, [ezsurfc], page 412, [hidden], page 385.

`ezmeshc (`*f*`)`
`ezmeshc (`*fx*`, `*fy*`, `*fz*`)`
`ezmeshc (...,` *dom*`)`
`ezmeshc (...,` *n*`)`
`ezmeshc (...,` "*circ*"`)`
`ezmeshc (`*hax*`, ...)`
*h* `= ezmeshc (...)`

Plot the mesh and contour lines defined by a function.

*f* is a string, inline function, or function handle with two arguments defining the function. By default the plot is over the meshed domain `-2*pi <=` *x* `|` *y* `<= 2*pi` with 60 points in each dimension.

If three functions are passed, then plot the parametrically defined function `[`*fx*`(`*s*`, `*t*`), `*fy*`(`*s*`, `*t*`), `*fz*`(`*s*`, `*t*`)]`.

If *dom* is a two element vector, it represents the minimum and maximum values of both *x* and *y*. If *dom* is a four element vector, then the minimum and maximum values are `[xmin xmax ymin ymax]`.

*n* is a scalar defining the number of points to use in each dimension.

If the argument `"circ"` is given, then the function is plotted over a disk centered on the middle of the domain *dom*.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a 2-element vector with a graphics handle for the created mesh plot and a second handle for the created contour plot.

Example: 2-argument function

```
f = @(x,y) sqrt (abs (x .* y)) ./ (1 + x.^2 + y.^2);
ezmeshc (f, [-3, 3]);
```

**See also:** [meshc], page 384, [ezmesh], page 410, [ezplot], page 379, [ezsurf], page 411, [ezsurfc], page 412, [hidden], page 385.

`ezsurf (`*f*`)`
`ezsurf (`*fx*`, `*fy*`, `*fz*`)`
`ezsurf (...,` *dom*`)`
`ezsurf (...,` *n*`)`
`ezsurf (...,` "*circ*"`)`
`ezsurf (`*hax*`, ...)`
*h* `= ezsurf (...)`

Plot the surface defined by a function.

*f* is a string, inline function, or function handle with two arguments defining the function. By default the plot is over the meshed domain `-2*pi <= x | y <= 2*pi` with 60 points in each dimension.

If three functions are passed, then plot the parametrically defined function `[fx(s, t), fy(s, t), fz(s, t)]`.

If *dom* is a two element vector, it represents the minimum and maximum values of both *x* and *y*. If *dom* is a four element vector, then the minimum and maximum values are `[xmin xmax ymin ymax]`.

*n* is a scalar defining the number of points to use in each dimension.

If the argument `"circ"` is given, then the function is plotted over a disk centered on the middle of the domain *dom*.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created surface object.

Example 1: 2-argument function

```
f = @(x,y) sqrt (abs (x .* y)) ./ (1 + x.^2 + y.^2);
ezsurf (f, [-3, 3]);
```

Example 2: parametrically defined function

```
fx = @(s,t) cos (s) .* cos (t);
fy = @(s,t) sin (s) .* cos (t);
fz = @(s,t) sin (t);
ezsurf (fx, fy, fz, [-pi, pi, -pi/2, pi/2], 20);
```

**See also:** [surf], page 385, [ezsurfc], page 412, [ezplot], page 379, [ezmesh], page 410, [ezmeshc], page 411, [shading], page 407.

`ezsurfc (f)`
`ezsurfc (fx, fy, fz)`
`ezsurfc (..., dom)`
`ezsurfc (..., n)`
`ezsurfc (..., "circ")`
`ezsurfc (hax, ...)`
`h = ezsurfc (...)`

Plot the surface and contour lines defined by a function.

*f* is a string, inline function, or function handle with two arguments defining the function. By default the plot is over the meshed domain `-2*pi <= x | y <= 2*pi` with 60 points in each dimension.

If three functions are passed, then plot the parametrically defined function `[fx(s, t), fy(s, t), fz(s, t)]`.

If *dom* is a two element vector, it represents the minimum and maximum values of both *x* and *y*. If *dom* is a four element vector, then the minimum and maximum values are `[xmin xmax ymin ymax]`.

*n* is a scalar defining the number of points to use in each dimension.

If the argument `"circ"` is given, then the function is plotted over a disk centered on the middle of the domain *dom*.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a 2-element vector with a graphics handle for the created surface plot and a second handle for the created contour plot.

Example:

```
f = @(x,y) sqrt (abs (x .* y)) ./ (1 + x.^2 + y.^2);
ezsurfc (f, [-3, 3]);
```

**See also:** [surfc], page 386, [ezsurf], page 411, [ezplot], page 379, [ezmesh], page 410, [ezmeshc], page 411, [shading], page 407.

### 15.2.2.3 Three-dimensional Geometric Shapes

`cylinder`
`cylinder (`*r*`)`
`cylinder (`*r*`, `*n*`)`
`cylinder (`*hax*`, ...)`
`[`*x*`, `*y*`, `*z*`] = cylinder (...)`

Plot a 3-D unit cylinder.

The optional input *r* is a vector specifying the radius along the unit z-axis. The default is [1 1] indicating radius 1 at `Z == 0` and at `Z == 1`.

The optional input *n* determines the number of faces around the circumference of the cylinder. The default value is 20.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

If outputs are requested `cylinder` returns three matrices in `meshgrid` format, such that `surf (`*x*`, `*y*`, `*z*`)` generates a unit cylinder.

Example:

```
[x, y, z] = cylinder (10:-1:0, 50);
surf (x, y, z);
title ("a cone");
```

**See also:** [ellipsoid], page 414, [rectangle], page 381, [sphere], page 413.

`sphere ()`
`sphere (`*n*`)`
`sphere (`*hax*`, ...)`
`[`*x*`, `*y*`, `*z*`] = sphere (...)`

Plot a 3-D unit sphere.

The optional input *n* determines the number of faces around the circumference of the sphere. The default value is 20.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

If outputs are requested `sphere` returns three matrices in `meshgrid` format such that `surf (`*x*`, `*y*`, `*z*`)` generates a unit sphere.

Example:

```
[x, y, z] = sphere (40);
surf (3*x, 3*y, 3*z);
axis equal;
title ("sphere of radius 3");
```

**See also:** [cylinder], page 413, [ellipsoid], page 414, [rectangle], page 381.

`ellipsoid (`*xc*, *yc*, *zc*, *xr*, *yr*, *zr*`)`
`ellipsoid (`*...*, *n*`)`
`ellipsoid (`*hax*, *...*`)`
`[`*x*, *y*, *z*`] = ellipsoid (...)`

Plot a 3-D ellipsoid.

The inputs *xc*, *yc*, *zc* specify the center of the ellipsoid. The inputs *xr*, *yr*, *zr* specify the semi-major axis lengths.

The optional input *n* determines the number of faces around the circumference of the cylinder. The default value is 20.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

If outputs are requested `ellipsoid` returns three matrices in `meshgrid` format, such that `surf (`*x*, *y*, *z*`)` generates the ellipsoid.

**See also:** [cylinder], page 413, [rectangle], page 381, [sphere], page 413.

### 15.2.3 Plot Annotations

You can add titles, axis labels, legends, and arbitrary text to an existing plot. For example:

```
x = -10:0.1:10;
plot (x, sin (x));
title ("sin(x) for x = -10:0.1:10");
xlabel ("x");
ylabel ("sin (x)");
text (pi, 0.7, "arbitrary text");
legend ("sin (x)");
```

The functions `grid` and `box` may also be used to add grid and border lines to the plot. By default, the grid is off and the border lines are on.

Finally, arrows, text and rectangular or elliptic boxes can be added to highlight parts of a plot using the `annotation` function. Those objects are drawn in an invisible axes, on top of every other axes.

`title (`*string*`)`
`title (`*string*, *prop*, *val*, *...*`)`
`title (`*hax*, *...*`)`
*h* `= title (...)`

Specify the string used as a title for the current axis.

An optional list of *property*/*value* pairs can be used to change the appearance of the created title text object.

If the first argument *hax* is an axes or legend handle, then add a title to this object, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created text object.

**See also:** [xlabel], page 418, [ylabel], page 418, [zlabel], page 419, [text], page 417.

`legend ()`
`legend` *`command`*
`legend (`*`str1`*`,` *`str2`*`, ...)`
`legend (`*`charmat`*`)`
`legend ({`*`cellstr`*`})`
`legend (...,` *`property`*`,` *`value`*`, ...)`
`legend (`*`hobjs`*`, ...)`
`legend ("`*`command`*`")`
`legend (`*`hax`*`, ...)`
`legend (`*`hleg`*`, ...)`
*`hleg`* `= legend (...)`

Display a legend for the current axes using the specified strings as labels.

Legend entries may be specified as individual character string arguments, a character array, or a cell array of character strings. When label names might be confused with legend properties, or *command* arguments, the labels should be protected by specifying them as a cell array of strings.

If the first argument *hax* is an axes handle, then add a legend to this axes, rather than the current axes returned by `gca`.

If the first argument *hleg* is a legend handle, then operate on this legend rather than the legend of the current axes.

Legend labels are associated with the axes' children; The first label is assigned to the first object that was plotted in the axes, the second label to the next object plotted, etc. To label specific data objects, without labeling all objects, provide their graphic handles in the input *hobjs*.

The following customizations are available using *command*:

| | |
|---|---|
| `"show"` | Show legend on the plot |
| `"hide"` | Hide legend on the plot |
| `"toggle"` | Toggle between `"hide"` and `"show"` |
| `"boxon"` | Show a box around legend (default) |
| `"boxoff"` | Hide the box around legend |
| `"right"` | Place label text to the right of the keys (default) |
| `"left"` | Place label text to the left of the keys |
| `"off"` | Delete the legend object |

The `legend` function creates a graphics object which has various properties that can be manipulated with `get`/`set`. Alternatively, properties can be set directly when calling `legend` by including *property*/*value* pairs. If using this calling form, the

labels must be specified as a cell array of strings. Graphics object properties are documented in detail at Section 15.3.3 [Graphics Object Properties], page 460.

Following is a subset of supported legend properties:

`autoupdate`: `"off"` | {`"on"`}
: Control whether the number of legend items is updated automatically when objects are added to (or deleted from) the peer axes. For example:

```
## Create a single plot with its legend.
figure ();
plot (1:10);
legend ("Slope 1");
## Add another plot and specify its displayname so that
## the legend is correctly updated.
hold on;
plot ((1:10) * 2, "displayname", "Slope 2");
## Stop automatic updates for further plots.
legend ("autoupdate", "off");
plot ((1:10) * 3);
```

`box`: `"off"` | {`"on"`}
: Control whether the legend has a surrounding box.

`location`: `"best"` | `"bestoutside"` | `"east"` | `"eastoutside"` | `"none"` | `"north"` | {`"northeast"`} | `"northeastoutside"` | `"northoutside"` | `"northwest"` | `"northwestoutside"` | `"south"` | `"southeast"` | `"southeastoutside"` | `"southoutside"` | `"southwest"` | `"southwestoutside"` | `"west"` | `"westoutside"` Control the location of the legend.

`numcolumns`: scalar interger, def. `1`
: Control the number of columns used in the layout of the legend items. For example:

```
figure ();
plot (rand (30));
legend ("numcolumns", 3);
```

Setting `numcolumns` also forces the `numcolumnsmode` property to be set to `"manual"`.

`orientation`: `"horizontal"` | {`"vertical"`}
: Control whether the legend items are arranged vertically (column-wise) or horizontally (row-wise).

`string`: string | cell array of strings
: List of labels for the legend items. For example:

```
figure ();
plot (rand (20));
## Let legend choose names automatically
hl = legend ();
## Selectively change some names
str = get (hl, "string");
str(1:5:end) = "Garbage";
set (hl, "string", str);
```

`textcolor`: colorspec, def. `[0 0 0]`
Control the color of the text strings for legend item.

The full list of supported legend specific properties can be found at Section 15.3.3.4 [Legend Properties], page 483.

A legend is implemented as an additional axes object with the `tag` property set to `"legend"`. Properties of the legend object may be manipulated directly by using `set`.

The optional output value *hleg* is a handle to the legend object.

Implementation Note: The legend label text is either provided in the call to `legend` or is taken from the `DisplayName` property of the graphics objects. Only data objects, such as line, patch, and surface, have this property whereas axes, figures, etc. do not so they are never present in a legend. If no labels or `DisplayName` properties are available, then the label text is simply `"data1"`, `"data2"`, ..., `"dataN"`.

The legend `FontSize` property is initially set to 90% of the axes `FontSize` to which it is attached. Use `set` to override this if necessary.

`text` (*x*, *y*, *string*)
`text` (*x*, *y*, *z*, *string*)
`text` (..., *prop*, *val*, ...)
`text` (*hax*, ...)
*h* = `text` (...)

Create a text object with text *string* at position *x*, *y*, (*z*) on the current axes.

Multiple locations can be specified if *x*, *y*, (*z*) are vectors. Multiple strings can be specified with a character matrix or a cell array of strings.

Optional property/value pairs may be used to control the appearance of the text.

If the first argument *hax* is an axes handle, then add text to this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a vector of graphics handles to the created text objects.

Example 1 : multi-line text via 3 different methods

```
text (0.5, 0.8, {"Line 1", "Line 2"})
text (0.5, 0.6, ["Line 1"; "Line 2"])
text (0.5, 0.4, "Line 1\nLine 2")
```

Example 2 : text at multiple locations

```
text ([0.2, 0.2], [0.8, 0.6], "Same text at two locations")
text ([0.4, 0.4], [0.8, 0.6], {"Point 1 Text", "Point 2 text"})
text ([0.6, 0.6], [0.8, 0.6], {{"Point 1 Line 1", "Point 1 Line 2},
                               "Point 2 text"})
```

Example 2 : adjust appearance using text properties

```
ht = text (0.5, 0.5, "Hello World", "fontsize", 20);
set (ht, "color", "red");
```

Programming Notes: The full list of properties is documented at Section 15.3.3.6 [Text Properties], page 489.

Any numeric entries in a cell array will be converted to text using `sprintf ("%g")`. For more precise control of the appearance convert any numeric entries to strings using `num2str`, `sprintf`, etc., before calling `text`.

**See also:** [gtext], page 446, [title], page 414, [xlabel], page 418, [ylabel], page 418, [zlabel], page 419.

`xlabel (`*`string`*`)`
`xlabel (`*`string`*`, `*`property`*`, `*`val`*`, ...)`
`xlabel (`*`hax`*`, ...)`
*`h`* `= xlabel (...)`

Specify the string used to label the x-axis of the current axis.

An optional list of *property*/*value* pairs can be used to change the properties of the created text label.

The full list of text object properties is documented at Section 15.3.3.6 [Text Properties], page 489.

If the first argument *hax* is an axes handle, then operate on this axes rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created text object.

**See also:** [ylabel], page 418, [zlabel], page 419, [datetick], page 1037, [title], page 414, [text], page 417.

`ylabel (`*`string`*`)`
`ylabel (`*`string`*`, `*`property`*`, `*`val`*`, ...)`
`ylabel (`*`hax`*`, ...)`
*`h`* `= ylabel (...)`

Specify the string used to label the y-axis of the current axis.

If *hax* is specified then label the axis defined by *hax*.

An optional list of *property*/*value* pairs can be used to change the properties of the created text label.

The full list of text object properties is documented at Section 15.3.3.6 [Text Properties], page 489.

If the first argument *hax* is an axes handle, then operate on this axes rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created text object.

**See also:** [xlabel], page 418, [zlabel], page 419, [datetick], page 1037, [title], page 414, [text], page 417.

`zlabel (`*`string`*`)`
`zlabel (`*`string`*`, `*`property`*`, `*`val`*`, ...)`
`zlabel (`*`hax`*`, ...)`
*`h`* `= zlabel (...)`

Specify the string used to label the z-axis of the current axis.

An optional list of *property*/*value* pairs can be used to change the properties of the created text label.

The full list of text object properties is documented at Section 15.3.3.6 [Text Properties], page 489.

If the first argument *hax* is an axes handle, then operate on this axes rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created text object.

**See also:** [xlabel], page 418, [ylabel], page 418, [datetick], page 1037, [title], page 414, [text], page 417.

`clabel (`*`c`*`, `*`h`*`)`
`clabel (`*`c`*`, `*`h`*`, `*`v`*`)`
`clabel (`*`c`*`, `*`h`*`, "`*`manual`*`")`
`clabel (`*`c`*`)`
`clabel (...,` *`prop`*`, `*`val`*`, ...)`
*`hlabels`* `= clabel (...)`

Add labels to the contours of a contour plot.

The contour levels are specified by the contour matrix *c* which is returned by `contour`, `contourc`, `contourf`, and `contour3`. Contour labels are rotated to match the local line orientation and centered on the line. The position of labels along the contour line is chosen randomly.

If the argument *h* is a handle to a contour group object, then label this plot rather than the one in the current axes returned by `gca`.

By default, all contours are labeled. However, the contours to label can be specified by the vector *v*. If the `"manual"` argument is given then the contours to label can be selected with the mouse.

Additional property/value pairs that are valid properties of text objects can be given and are passed to the underlying text objects. Moreover, the contour group property `"LabelSpacing"` is available which determines the spacing between labels on a contour to be specified. The default is 144 points, or 2 inches.

The optional return value *hlabels* is a vector of graphics handles to the text objects representing each label. The `"userdata"` property of the text objects contains the numerical value of the contour label.

The full list of text object properties is documented at Section 15.3.3.6 [Text Properties], page 489.

```
[c, h] = contour (peaks (), -4 : 6);
clabel (c, h, -4:2:6, "fontsize", 12);
```

**See also:** [contour], page 348, [contourf], page 349, [contour3], page 350, [meshc], page 384, [surfc], page 386, [text], page 417.

`box`
`box` *on*
`box` *off*
`box` (*hax*, ...)

Control display of the axes border.

The argument may be either `"on"` or `"off"`. If it is omitted, the current box state is toggled.

If the first argument *hax* is an axes handle, then operate on this axes rather than the current axes returned by `gca`.

**See also:** [axis], page 366, [grid], page 420.

`grid`
`grid` *on*
`grid` *off*
`grid` *minor*
`grid` *minor on*
`grid` *minor off*
`grid` (*hax*, ...)

Control the display of plot grid lines.

The function state input may be either `"on"` or `"off"`. If it is omitted, the current grid state is toggled.

When the first argument is `"minor"` all subsequent commands modify the minor grid rather than the major grid.

If the first argument *hax* is an axes handle, then operate on this axes rather than the current axes returned by `gca`.

To control the grid lines for an individual axes use the `set` function. For example:

```
set (gca, "ygrid", "on");
```

**See also:** [axis], page 366, [box], page 420.

`colorbar`
`colorbar` (*loc*, ...)
`colorbar` (*delete_option*)
`colorbar` (*hcb*, ...)
`colorbar` (*hax*, ...)
`colorbar` (..., "*location*", *loc*, ...)
`colorbar` (..., *prop*, *val*, ...)
*h* = `colorbar` (...)

Add a colorbar to the current axes.

A colorbar displays the current colormap along with numerical rulings so that the color scale can be interpreted.

The optional input *loc* determines the location of the colorbar. If present, it must be the first argument to `colorbar`. Valid values for *loc* are

`"EastOutside"`

Place the colorbar outside the plot to the right. This is the default.

`"East"` Place the colorbar inside the plot to the right.

`"WestOutside"`
Place the colorbar outside the plot to the left.

`"West"` Place the colorbar inside the plot to the left.

`"NorthOutside"`
Place the colorbar above the plot.

`"North"` Place the colorbar at the top of the plot.

`"SouthOutside"`
Place the colorbar under the plot.

`"South"` Place the colorbar at the bottom of the plot.

To remove a colorbar from a plot use any one of the following keywords for the *delete_option*: `"off"`, `"delete"`, `"hide"`.

If the first argument *hax* is an axes handle, then the colorbar is added to this axes, rather than the current axes returned by `gca`.

If the first argument *hcb* is a handle to a colorbar object, then operate on this colorbar directly.

Additional property/value pairs are passed directly to the underlying axes object. The full list of properties is documented at Section 15.3.3.3 [Axes Properties], page 470.

The optional return value *h* is a graphics handle to the created colorbar object.

Implementation Note: A colorbar is created as an additional axes object with the `"tag"` property set to `"colorbar"`. The created object has the extra property `"location"` which controls the positioning of the colorbar.

If the argument `"peer"` is given, then the following argument is treated as the axes handle in which to add the colorbar. The `"peer"` calling syntax may be removed in the future and is not recommended.

**See also:** [colormap], page 951.

`annotation (`*`type`*`)`
`annotation ("`*`line`*`", `*`x`*`, `*`y`*`)`
`annotation ("`*`arrow`*`", `*`x`*`, `*`y`*`)`
`annotation ("`*`doublearrow`*`", `*`x`*`, `*`y`*`)`
`annotation ("`*`textarrow`*`", `*`x`*`, `*`y`*`)`
`annotation ("`*`textbox`*`", `*`pos`*`)`
`annotation ("`*`rectangle`*`", `*`pos`*`)`
`annotation ("`*`ellipse`*`", `*`pos`*`)`
`annotation (...,` *`prop`*`,` *`val`*`)`
`annotation (`*`hf`*`, ...)`
*`h`* `= annotation (...)`

Draw annotations to emphasize parts of a figure.

You may build a default annotation by specifying only the *type* of the annotation.

Otherwise you can select the type of annotation and then set its position using either *x* and *y* coordinates for line-based annotations or a position vector *pos* for others. In

either case, coordinates are interpreted using the `"units"` property of the annotation object. The default is `"normalized"`, which means the lower left hand corner of the figure has coordinates '`[0 0]`' and the upper right hand corner '`[1 1]`'.

If the first argument *hf* is a figure handle, then plot into this figure, rather than the current figure returned by `gcf`.

Further arguments can be provided in the form of *prop/val* pairs to customize the annotation appearance.

The optional return value *h* is a graphics handle to the created annotation object. This can be used with the `set` function to customize an existing annotation object.

All annotation objects share two properties:

- `"units"`: the units in which coordinates are interpreted.
  Its value may be one of `"centimeters" | "characters" | "inches" | "{normalized}" | "pixels" | "points"`.
- `"position"`: a four-element vector [x0 y0 width height].
  The vector specifies the coordinates (x0,y0) of the origin of the annotation object, its width, and its height. The width and height may be negative, depending on the orientation of the object.

Valid annotation types and their specific properties are described below:

`"line"` Constructs a line. *x* and *y* must be two-element vectors specifying the x and y coordinates of the two ends of the line.

The line can be customized using `"linewidth"`, `"linestyle"`, and `"color"` properties the same way as for `line` objects.

`"arrow"` Construct an arrow. The second point in vectors *x* and *y* specifies the arrowhead coordinates.

Besides line properties, the arrowhead can be customized using `"headlength"`, `"headwidth"`, and `"headstyle"` properties. Supported values for `"headstyle"` property are: [`"diamond" | "ellipse" | "plain" | "rectangle" | "vback1" | "{vback2}" | "vback3"`]

`"doublearrow"`
Construct a double arrow. Vectors *x* and *y* specify the arrowhead coordinates.

The line and the arrowhead can be customized as for arrow annotations, but some property names are duplicated: `"head1length"`/`"head2length"`, `"head1width"`/`"head2width"`, etc. The index 1 marks the properties of the arrowhead at the first point in *x* and *y* coordinates.

`"textarrow"`
Construct an arrow with a text label at the opposite end from the arrowhead.

Use the `"string"` property to change the text string. The line and the arrowhead can be customized as for arrow annotations, and the text can be customized using the same properties as `text` graphics objects.

Note, however, that some text property names are prefixed with "text" to distinguish them from arrow properties: "textbackgroundcolor", "textcolor", "textedgecolor", "textlinewidth", "textmargin", "textrotation".

"textbox"
: Construct a box with text inside. *pos* specifies the "position" property of the annotation.

  Use the "string" property to change the text string. You may use "backgroundcolor", "edgecolor", "linestyle", and "linewidth" properties to customize the box background color and edge appearance. A limited set of text objects properties are also available; Besides "font..." properties, you may also use "horizontalalignment" and "verticalalignment" to position the text inside the box.

  Finally, the "fitboxtotext" property controls the actual extent of the box. If "on" (the default) the box limits are fitted to the text extent.

"rectangle"
: Construct a rectangle. *pos* specifies the "position" property of the annotation.

  You may use "facecolor", "color", "linestyle", and "linewidth" properties to customize the rectangle background color and edge appearance.

"ellipse"
: Construct an ellipse. *pos* specifies the "position" property of the annotation.

  See "rectangle" annotations for customization.

**See also:** [xlabel], page 418, [ylabel], page 418, [zlabel], page 419, [title], page 414, [text], page 417, [gtext], page 446, [legend], page 415, [colorbar], page 420.

### 15.2.4 Multiple Plots on One Page

Octave can display more than one plot in a single figure. The simplest way to do this is to use the subplot function to divide the plot area into a series of subplot windows that are indexed by an integer. For example,

```
subplot (2, 1, 1)
fplot (@sin, [-10, 10]);
subplot (2, 1, 2)
fplot (@cos, [-10, 10]);
```

creates a figure with two separate axes, one displaying a sine wave and the other a cosine wave. The first call to subplot divides the figure into two plotting areas (two rows and one column) and makes the first plot area active. The grid of plot areas created by subplot is numbered in row-major order (left to right, top to bottom). After plotting a sine wave, the next call to subplot activates the second subplot area, but does not re-partition the figure.

subplot (*rows*, *cols*, *index*)
subplot (*rows*, *cols*, *index*, *hax*)

`subplot (`*rcn*`)`
`subplot (`*hax*`)`
`subplot (...,` "*align*"`)`
`subplot (...,` "*replace*"`)`
`subplot (`"*position*"`,` *pos*`)`
`subplot (...,` *prop*`,` *val*`, ...)`
*hax* `= subplot (...)`

Set up a plot grid with *rows* by *cols* subwindows and set the current axes for plotting (`gca`) to the location given by *index*.

If an axes handle *hax* is provided after the (*rows*, *cols*, *index*) arguments, the corresponding axes is turned into a subplot.

If only one numeric argument is supplied, then it must be a three digit value specifying the number of rows in digit 1, the number of columns in digit 2, and the plot index in digit 3.

The plot index runs row-wise; First, all columns in a row are numbered and then the next row is filled.

For example, a plot with 2x3 grid will have plot indices running as follows:

| | | |
|---|---|---|
| 1 | 2 | 3 |
| 4 | 5 | 6 |

*index* may also be a vector. In this case, the new axes will enclose the grid locations specified. The first demo illustrates this:

```
demo ("subplot", 1)
```

The index of the subplot to make active may also be specified by its axes handle, *hax*, returned from a previous `subplot` command.

If the option `"align"` is given then the plot boxes of the subwindows will align, but this may leave no room for axes tick marks or labels.

If the option `"replace"` is given then the subplot axes will be reset, rather than just switching the current axes for plotting to the requested subplot.

The `"position"` property can be used to exactly position the subplot axes within the current figure. The option *pos* is a 4-element vector [x, y, width, height] that determines the location and size of the axes. The values in *pos* are normalized in the range [0,1].

Any property/value pairs are passed directly to the underlying axes object. The full list of properties is documented at Section 15.3.3.3 [Axes Properties], page 470.

Any previously existing axes that would be (partly) covered by the newly created axes are deleted.

If the output *hax* is requested, subplot returns the axes handle for the subplot. This is useful for modifying the properties of a subplot using `set`.

Under some circumstances, `subplot` might not be able to identify axes that it could re-use and might replace them. If `subplot` axes should be referenced repeatedly,

consider creating and storing their axes handles beforehand instead of calling `subplot` repeatedly for the same position.

Example:

```
x = 1:10;
y = rand (16, 10);
for i_plot = 1:4
  hax(i_plot) = subplot (2, 2, i_plot);
  hold (hax(i_plot), "on");
  grid (hax(i_plot), "on");
endfor
for i_loop = 1:2
  for i_plot = 1:4
    iy = (i_loop - 1)*4 + i_plot;
    plotyy (hax(i_plot), x,y(iy,:), x,y(iy+1,:));
  endfor
 endfor
```

**See also:** [axes], page 450, [plot], page 334, [gca], page 455, [set], page 457.

### 15.2.5 Multiple Plot Windows

You can open multiple plot windows using the `figure` function. For example,

```
figure (1);
fplot (@sin, [-10, 10]);
figure (2);
fplot (@cos, [-10, 10]);
```

creates two figures, with the first displaying a sine wave and the second a cosine wave. Figure numbers must be positive integers.

`figure`
`figure` *n*
`figure` (*n*)
`figure` (..., "*property*", *value*, ...)
*h* = `figure` (...)

Create a new figure window for plotting.

If no arguments are specified, a new figure with the next available number is created.

If called with an integer *n*, and no such numbered figure exists, then a new figure with the specified number is created. If the figure already exists then it is made visible and becomes the current figure for plotting.

Multiple property-value pairs may be specified for the figure object, but they must appear in pairs.

The optional return value *h* is a graphics handle to the created figure object.

Programming Note: The full list of properties is documented at Section 15.3.3.2 [Figure Properties], page 462.

**See also:** [axes], page 450, [gcf], page 455, [shg], page 430, [clf], page 429, [close], page 431.

### 15.2.6 Manipulation of Plot Objects

`pan`
`pan` *on*
`pan` *off*
`pan` *xon*
`pan` *yon*
`pan` (*hfig*, *option*)

Control the interactive panning mode of a figure in the GUI.

Given the option `"on"` or `"off"`, set the interactive pan mode on or off.

With no arguments, toggle the current pan mode on or off.

Given the option `"xon"` or `"yon"`, enable pan mode for the x or y axis only.

If the first argument *hfig* is a figure, then operate on the given figure rather than the current figure as returned by `gcf`.

**See also:** [rotate3d], page 426, [zoom], page 426.

`rotate` (*h*, *direction*, *alpha*)
`rotate` (*...*, *origin*)

Rotate the plot object *h* through *alpha* degrees around the line with direction *direction* and origin *origin*.

The default value of *origin* is the center of the axes object that is the parent of *h*.

If *h* is a vector of handles, they must all have the same parent axes object.

Graphics objects that may be rotated are lines, surfaces, patches, and images.

`rotate3d`
`rotate3d` *on*
`rotate3d` *off*
`rotate3d` (*hfig*, *option*)

Control the interactive 3-D rotation mode of a figure in the GUI.

Given the option `"on"` or `"off"`, set the interactive rotate mode on or off.

With no arguments, toggle the current rotate mode on or off.

If the first argument *hfig* is a figure, then operate on the given figure rather than the current figure as returned by `gcf`.

**See also:** [pan], page 426, [zoom], page 426.

`zoom`
`zoom` (*factor*)
`zoom` *on*
`zoom` *off*
`zoom` *xon*
`zoom` *yon*
`zoom` *out*
`zoom` *reset*
`zoom` (*hfig*, *option*)

Zoom the current axes object or control the interactive zoom mode of a figure in the GUI.

Given a numeric argument greater than zero, zoom by the given factor. If the zoom factor is greater than one, zoom in on the plot. If the factor is less than one, zoom out. If the zoom factor is a two- or three-element vector, then the elements specify the zoom factors for the x, y, and z axes respectively.

Given the option `"on"` or `"off"`, set the interactive zoom mode on or off.

With no arguments, toggle the current zoom mode on or off.

Given the option `"xon"` or `"yon"`, enable zoom mode for the x or y-axis only.

Given the option `"out"`, zoom to the initial zoom setting.

Given the option `"reset"`, store the current zoom setting so that `zoom out` will return to this zoom level.

If the first argument *hfig* is a figure, then operate on the given figure rather than the current figure as returned by `gcf`.

**See also:** [pan], page 426, [rotate3d], page 426.

### 15.2.7 Manipulation of Plot Windows

By default, Octave refreshes the plot window when a prompt is printed, or when waiting for input. The `drawnow` function is used to cause a plot window to be updated.

`drawnow` ()
`drawnow` ("*expose*")
`drawnow` (*term*, *file*, *debug_file*)

Update figure windows and their children.

The event queue is flushed and any callbacks generated are executed.

With the optional argument `"expose"`, only graphic objects are updated and no other events or callbacks are processed.

The third calling form of `drawnow` is for debugging and is undocumented.

**See also:** [refresh], page 427.

Only figures that are modified will be updated. The `refresh` function can also be used to cause an update of the current figure, even if it is not modified.

`refresh` ()
`refresh` (*h*)

Refresh a figure, forcing it to be redrawn.

When called without an argument the current figure is redrawn. Otherwise, the figure with graphic handle *h* is redrawn.

**See also:** [drawnow], page 427.

Normally, high-level plot functions like `plot` or `mesh` call `newplot` to determine whether the state of the target axes should be initialized (the default) or if subsequent plots should be drawn on top of previous ones. To have two plots drawn over one another, use the `hold` function or manually change the axes [nextplot], page 480, property. For example,

```
hold on;
x = -10:0.1:10;
plot (x, sin (x));
plot (x, cos (x));
hold off;
```

displays sine and cosine waves on the same axes. If the hold state is off, consecutive plotting commands like this will only display the last plot.

`newplot ()`
`newplot (`*`hfig`*`)`
`newplot (`*`hax`*`)`
*`hax`* `= newplot (...)`

Prepare graphics engine to produce a new plot.

This function is called at the beginning of all high-level plotting functions. It is not normally required in user programs. `newplot` queries the `"NextPlot"` field of the current figure and axes to determine what to do.

| **Figure NextPlot** | **Action** |
|---|---|
| `"add"` (default) | Add new graphic objects to the current figure. |
| `"new"` | Create a new figure and make it the current figure. |
| `"replace"` | Delete all child objects of the figure and reset all figure properties to their defaults. However, the following four properties are not reset: Position, Units, PaperPosition, PaperUnits. In addition, the NextPlot property is set to `"add"` regardless of user-defined defaults. This is equivalent to `clf reset`. |
| `"replacechildren"` | Delete child objects whose HandleVisibility is set to `"on"`. Set NextPlot property to `"add"`. This typically clears a figure, but leaves in place hidden objects such as menubars. This is equivalent to `clf`. |

| **Axes NextPlot** | **Action** |
|---|---|
| `"add"` | Add new graphic objects to the current axes. This is equivalent to `hold on`. |
| `"replace"` (default) | Delete all child objects of the axes and reset all axes properties to their defaults. However, the following properties are not reset: Position, Units. This is equivalent to `cla reset`. |

`"replacechildren"` Delete child objects whose HandleVisibility is set to `"on"`, but leave axes properties unmodified except for ColorOrderIndex and LineStyleOrderIndex which are set to 1. This typically clears a plot, but preserves special settings such as log scaling for axes. This is equivalent to `cla`.

If the optional input *hfig* or *hax* is given then prepare the specified figure or axes rather than the current figure and axes.

The optional return value *hax* is a graphics handle to the created axes object (not figure).

`hold`
`hold` *on*
`hold` *off*
`hold (`*hax*`, ...)`

Toggle or set the `"hold"` state of the plotting engine which determines whether new graphic objects are added to the plot or replace the existing objects.

`hold on` Retain plot data and settings so that subsequent plot commands are displayed on a single graph. Line color and line style are advanced for each new plot added.

`hold all (deprecated)`
Equivalent to `hold on`.

`hold off` Restore default graphics settings which clear the graph and reset axes properties before each new plot command. (default).

`hold` Toggle the current hold state.

When given the additional argument *hax*, the hold state is modified for this axes rather than the current axes returned by `gca`.

To query the current hold state use the `ishold` function.

**See also:** [ishold], page 429, [cla], page 430, [clf], page 429, [newplot], page 428.

*tf* = `ishold`
*tf* = `ishold (`*hax*`)`
*tf* = `ishold (`*hfig*`)`

Return true if the next plot will be added to the current plot, or false if the plot device will be cleared before drawing the next plot.

If the first argument is an axes handle *hax* or figure handle *hfig* then operate on this plot rather than the current one.

**See also:** [hold], page 429, [newplot], page 428.

To clear the current figure, call the `clf` function. To clear the current axis, call the `cla` function. To bring the current figure to the top of the window stack, call the `shg` function. To delete a graphics object, call `delete` on its index. To close the figure window, call the `close` function.

`clf`
`clf` *reset*
`clf` (*hfig*)
`clf` (*hfig*, "*reset*")
*h* = `clf` (...)

Clear the current figure window.

`clf` operates by deleting child graphics objects with visible handles (HandleVisibility = `"on"`).

If the optional argument `"reset"` is specified, delete all child objects including those with hidden handles and reset all figure properties to their defaults. However, the following properties are not reset: Position, Units, PaperPosition, PaperUnits.

If the first argument *hfig* is a figure handle, then operate on this figure rather than the current figure returned by `gcf`.

The optional return value *h* is the graphics handle of the figure window that was cleared.

**See also:** [cla], page 430, [close], page 431, [delete], page 430, [reset], page 544.

`cla`
`cla` *reset*
`cla` (*hax*)
`cla` (*hax*, "*reset*")

Clear the current or specified (*hax*) axes object.

`cla` operates by deleting child graphic objects with visible handles (`HandleVisibility = "on"`). This typically clears the axes of any visual objects, but leaves in place axes limits, tick marks and labels, camera view, etc. In addition, the automatic coloring and styling of lines is reset by changing the axes properties `ColorOrderIndex`, `LinestyleOrderIndex` to 1.

If the optional argument `"reset"` is specified, delete all child objects, including those with hidden handles, and reset all axes properties to their defaults. However, the following properties are not reset: `Position`, `Units`.

If the first argument *hax* is an axes handle, then operate on this axes rather than the current axes returned by `gca`.

**See also:** [clf], page 429, [delete], page 430, [reset], page 544.

`shg`

Show the graph window.

This function makes the current figure visible, and places it on top of of all other plot windows.

Programming Note: `shg` is equivalent to `figure (gcf)` assuming that a current figure exists.

**See also:** [figure], page 425, [drawnow], page 427, [gcf], page 455.

`delete` *file*
`delete` *file1 file2 ...*
`delete` (*file*)

`delete (`*`file1`*`, `*`file2`*`, ...)`
`delete (`*`handle`*`)`

Delete the named file or graphics handle.

*file* may contain globbing patterns such as '`*`'. Multiple files to be deleted may be specified in the same function call.

*handle* may be a scalar or vector of graphic handles to delete.

Programming Note: Deleting graphics objects is the proper way to remove features from a plot without clearing the entire figure.

**See also:** [clf], page 429, [cla], page 430, [unlink], page 1039, [rmdir], page 1040.

`close`
`close (`*`h`*`)`
`close `*`figname`*
`close `*`all`*
`close `*`all hidden`*
`close `*`all force`*
*`status`* `= close (...)`

Close figure window(s).

When called with no arguments, close the current figure. This is equivalent to `close (gcf)`. If the input *h* is a graphic handle, or vector of graphics handles, then close each figure in *h*. The figure to close may also be specified by name *figname* which is matched against the `"Name"` property of all figures.

If the argument `"all"` is given then all figures with visible handles (HandleVisibility = `"on"`) are closed.

If the additional argument `"hidden"` is given then all figures, including hidden ones, are closed.

If the additional argument `"force"` is given then figures are closed even when `"closerequestfcn"` has been altered to prevent closing the window.

If the optional output *status* is requested then Octave returns 1 if the figure windows were closed successfully.

Implementation Note: `close` operates by making the handle *h* the current figure, and then calling the function specified by the `"closerequestfcn"` property of the figure. By default, the function `closereq` is used. It is possible that the function invoked will delay or abort removing the figure. To remove a figure without executing any callback functions use `delete`. When writing a callback function to close a window do not use `close` to avoid recursion.

**See also:** [closereq], page 431, [delete], page 430.

`closereq ()`

Close the current figure and delete all graphics objects associated with it.

By default, the `"closerequestfcn"` property of a new plot figure points to this function.

**See also:** [close], page 431, [delete], page 430.

### 15.2.8 Use of the "interpreter" Property

`text` (such as titles, labels, legend item) and `axes` objects feature an ["interpreter"], page 492, and a ["ticklabelinterpreter"], page 482, property respectively. It determines the manner in which special control sequences in the text are rendered.

The interpreter property can take three values: `"none"`, `"tex"`, `"latex"`.

#### 15.2.8.1 "none" interpreter

If the interpreter is set to `"none"` then no special rendering occurs—the displayed text is a verbatim copy of the specified text.

#### 15.2.8.2 "tex" interpreter

The `"tex"` interpreter implements a subset of TEX functionality when rendering text. This allows the insertion of special glyphs such as Greek characters or mathematical symbols. Special characters are inserted by using a backslash (\) character followed by a code, as shown in Table 15.1.

Besides special glyphs, the formatting of the text can be changed within the string by using the codes

| | |
|---|---|
| \bf | Bold font |
| \it | Italic font |
| \sl | Oblique Font |
| \rm | Normal font |

These codes may be used in conjunction with the { and } characters to limit the change to a part of the string. For example,

```
xlabel ('{\bf H} = a {\bf V}')
```

where the character `'a'` will not appear in bold font. Note that to avoid having Octave interpret the backslash character in the strings, the strings themselves should be in single quotes.

It is also possible to change the fontname and size within the text

| | |
|---|---|
| \fontname{*fontname*} | Specify the font to use |
| \fontsize{*size*} | Specify the size of the font to use |

The color of the text may also be changed inline using either a string (e.g., `"red"`) or numerically with a Red-Green-Blue (RGB) specification (e.g., [1 0 0], also red).

| | |
|---|---|
| \color{*color*} | Specify the color as a string |
| \color[rgb]{*R G B*} | Specify the color numerically |

Finally, superscripting and subscripting can be controlled with the `'^'` and `'_'` characters. If the `'^'` or `'_'` is followed by a { character, then all of the block surrounded by the { } pair is superscripted or subscripted. Without the { } pair, only the character immediately following the `'^'` or `'_'` is changed.

Greek Lowercase Letters

| Code | Sym | Code | Sym | Code | Sym |
|---|---|---|---|---|---|
| \alpha | $\alpha$ | \beta | $\beta$ | \gamma | $\gamma$ |
| \delta | $\delta$ | \epsilon | $\epsilon$ | \zeta | $\zeta$ |
| \eta | $\eta$ | \theta | $\theta$ | \vartheta | $\vartheta$ |
| \iota | $\iota$ | \kappa | $\kappa$ | \lambda | $\lambda$ |
| \mu | $\mu$ | \nu | $\nu$ | \xi | $\xi$ |
| \o | $o$ | \pi | $\pi$ | \varpi | $\varpi$ |
| \rho | $\rho$ | \sigma | $\sigma$ | \varsigma | $\varsigma$ |
| \tau | $\tau$ | \upsilon | $\upsilon$ | \phi | $\phi$ |
| \chi | $\chi$ | \psi | $\psi$ | \omega | $\omega$ |

Greek Uppercase Letters

| Code | Sym | Code | Sym | Code | Sym |
|---|---|---|---|---|---|
| \Gamma | $\Gamma$ | \Delta | $\Delta$ | \Theta | $\Theta$ |
| \Lambda | $\Lambda$ | \Xi | $\Xi$ | \Pi | $\Pi$ |
| \Sigma | $\Sigma$ | \Upsilon | $\Upsilon$ | \Phi | $\Phi$ |
| \Psi | $\Psi$ | \Omega | $\Omega$ | | |

Misc Symbols Type Ord

| Code | Sym | Code | Sym | Code | Sym |
|---|---|---|---|---|---|
| \aleph | $\aleph$ | \wp | $\wp$ | \Re | $\Re$ |
| \Im | $\Im$ | \partial | $\partial$ | \infty | $\infty$ |
| \prime | $\prime$ | \nabla | $\nabla$ | \surd | $\surd$ |
| \angle | $\angle$ | \forall | $\forall$ | \exists | $\exists$ |
| \neg | $\neg$ | \clubsuit | $\clubsuit$ | \diamondsuit | $\diamondsuit$ |
| \heartsuit | $\heartsuit$ | \spadesuit | $\spadesuit$ | | |

"Large" Operators

| Code | Sym | Code | Sym | Code | Sym |
|---|---|---|---|---|---|
| \int | $\int$ | | | | |

Binary operators

| Code | Sym | Code | Sym | Code | Sym |
|---|---|---|---|---|---|
| \pm | $\pm$ | \cdot | $\cdot$ | \times | $\times$ |
| \ast | $\ast$ | \circ | $\circ$ | \bullet | $\bullet$ |
| \div | $\div$ | \cap | $\cap$ | \cup | $\cup$ |
| \vee | $\vee$ | \wedge | $\wedge$ | \oplus | $\oplus$ |
| \otimes | $\otimes$ | \oslash | $\oslash$ | | |

Table 15.1: Available special characters in TEX mode

Relations

| Code | Sym | Code | Sym | Code | Sym |
|---|---|---|---|---|---|
| \leq | $\leq$ | \subset | $\subset$ | \subseteq | $\subseteq$ |
| \in | $\in$ | \geq | $\geq$ | \supset | $\supset$ |
| \supseteq | $\supseteq$ | \ni | $\ni$ | \mid | $\mid$ |
| \equiv | $\equiv$ | \sim | $\sim$ | \approx | $\approx$ |
| \cong | $\cong$ | \propto | $\propto$ | \perp | $\perp$ |

Arrows

| Code | Sym | Code | Sym | Code | Sym |
|---|---|---|---|---|---|
| \leftarrow | $\leftarrow$ | \Leftarrow | $\Leftarrow$ | \rightarrow | $\rightarrow$ |
| \Rightarrow | $\Rightarrow$ | \leftrightarrow | $\leftrightarrow$ | \uparrow | $\uparrow$ |
| \downarrow | $\downarrow$ | | | | |

Openings and Closings

| Code | Sym | Code | Sym | Code | Sym |
|---|---|---|---|---|---|
| \lfloor | $\lfloor$ | \langle | $\langle$ | \lceil | $\lceil$ |
| \rfloor | $\rfloor$ | \rangle | $\rangle$ | \rceil | $\rceil$ |

Alternate Names

| Code | Sym | Code | Sym | Code | Sym |
|---|---|---|---|---|---|
| \neq | $\neq$ | | | | |

Other (not in Appendix F Tables)

| Code | Sym | Code | Sym | Code | Sym |
|---|---|---|---|---|---|
| \ldots | $\ldots$ | \0 | $\oslash$ | \copyright | © |
| \deg | $^\circ$ | | | | |

Table 15.1: Available special characters in TEX mode (cont.)

**Caution: Degree Symbol**

Conformance to both TEX and MATLAB with respect to the `\circ` symbol is impossible. While TEX translates this symbol to Unicode 2218 (U+2218), MATLAB maps this to Unicode 00B0 (U+00B0) instead. Octave has chosen to follow the TEX specification, but has added the additional symbol `\deg` which maps to the degree symbol (U+00B0).

### 15.2.8.3 "latex" interpreter

The `"latex"` interpreter only works if an external LATEX tool chain is present. Three binaries are needed: `latex`, `dvipng`, and `dvisvgm`. If those binaries are installed but not on the path, one can still provide their respective path using the following environment variables: `OCTAVE_LATEX_BINARY`, `OCTAVE_DVIPNG_BINARY`, and `OCTAVE_DVISVG_BINARY`.

Note that Octave does not parse or validate the text strings when in `"latex"` mode—it is the responsibility of the programmer to generate valid strings which may include wrapping sections that should appear in Math mode with `'$'` characters. See, for example, https://www.latex-project.org/help/documentation/ for documentation about LaTeX typesetting.

For debugging purpose, a convenience environment variable, `OCTAVE_LATEX_DEBUG`, can be set to trigger more verbose output when Octave fails to have a given text compiled by an external LaTeX engine. For example, `"x^2"` is not a valid LaTeX string and the following example should fail

```
setenv ("OCTAVE_LATEX_DEBUG", "1")
x = 1:10;
plot (x, x.^2)
title ("x^2", "interpreter", "latex")
```

Searching the terminal output you should find some helpful info about the origin of the failure:

```
...
No file default.aux.
! Missing $ inserted.
<inserted text>
                $
l.6 x^
      2
! Missing $ inserted.
...
```

If no usable LaTeX tool chain is found at the first text rendering, using the `"latex"` interpreter is equivalent to `"none"`.

### 15.2.9 Printing and Saving Plots

The `print` command allows you to send plots to your printer and to save plots in a variety of formats. For example,

```
print -dpsc
```

prints the current figure to a color PostScript printer. And,

```
print foo.pdf
```

saves the current figure to a Portable Document encapsulated PostScript file called `foo.pdf`.

The current graphic toolkits produce very similar graphic displays, but the interpreters differ in their capability to display unusual text and in their ability to print such text. In general, the `"tex"` interpreter (default) is the best all-around performer for both on-screen display and printing. However, for the reproduction of complicated text formulas the `"latex"` interpreter is preferred. The `"none"` interpreter prints text verbatim (exactly as it appears) which is very portable, but there is no support for bold, italic, superscripts, subscripts, Greek letters, etc.

When printing with the `-painters` renderer, the default for all vector formats, two options may be considered:

- Use the `-svgconvert` option to allow for rendering LaTeX formulas. Note that the glyph are rendered as path and the original textual info are lost.

- Use one of the `-d*latex*` devices to produce a .tex file (plus supporting .eps or .pdf files) to be further processed by an external LaTeX engine. Note that the `print` function will first set the interpreter of all strings to `"latex"`, which means all strings must be valid LaTeX strings.

A complete example showing the capabilities of text printing using the `-dpdflatexstandalone` option is:

```
x = 0:0.01:3;
hf = figure ();
plot (x, erf (x));
hold on;
plot (x, x, "r");
axis ([0, 3, 0, 1]);
text (0.65, 0.6175, ...
      ['$\displaystyle\leftarrow x = {2 \over \sqrt{\pi}}' ...
       '\int_{0}^{x} e^{-t^2} dt = 0.6175$'],
      "interpreter", "latex");
xlabel ("x");
ylabel ("erf (x)");
title ("erf (x) with text annotation");
print (hf, "plot15_7", "-dpdflatexstandalone");
system ("pdflatex plot15_7");
open plot15_7.pdf
```

The result of this example can be seen in Figure 15.7

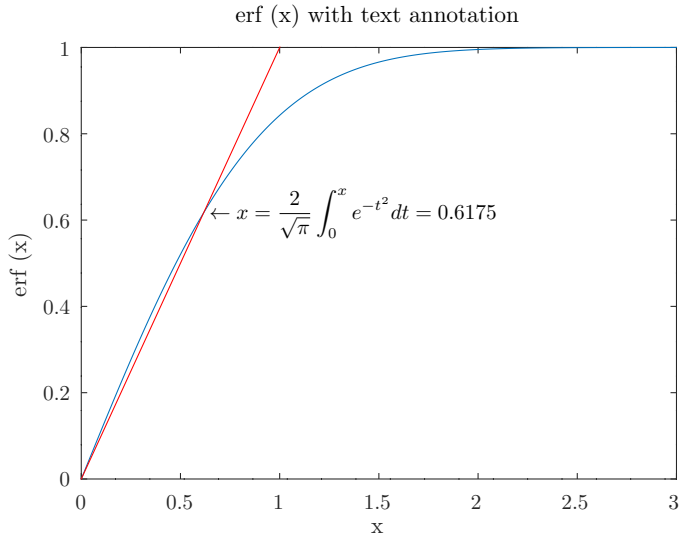



Figure 15.7: Example of inclusion of text with use of `-dpdflatexstandalone`

```
print ()
print (options)
print (filename, options)
```

`print (`*hfig*`, ...)`
*RGB* `= print ("-RGBImage", ...)`

Format a figure for printing and either save it to a file, send it to a printer, or return an RGB image.

*filename* defines the name of the output file. If the filename has no suffix then one is inferred from the specified device and appended to the filename. When neither a filename nor the `"-RGBImage"` option is present, the output is sent to the printer. The various options and filename arguments may be given in any order, except for the figure handle argument *hfig* which must be first if it is present.

Example: Print to a file using PDF and JPEG formats.

```
figure (1);
clf ();
surf (peaks);
print figure1.pdf    # The extension specifies the format
print -djpg figure1  # Will produce "figure1.jpg" file
```

If the first argument is a handle *hfig* to a figure object then it specifies the figure to print. By default, the current figure returned by `gcf` is printed.

For outputs to paged formats, for example, PostScript and PDF, the page size is specified by the figure's `papersize` property together with the `paperunits` property. The location and size of the plot on the page are specified by the figure's `paperposition` property. The orientation of the page is specified by the figure's `paperorientation` property.

For non-page formats—for example, image formats like JPEG—the width and height of the output are specified by the figure's `paperposition(3:4)` property values.

The `print` command supports many *options*:

`-fh` Specify the handle, *h*, of the figure to be printed.

Example: Print figure 1.

```
figure (1);
clf ();
surf (peaks);
figure (2);
print -f1 figure1.pdf
## Equivalent functional form:
print (1, "figure1.pdf")
```

`-P`*printer*

Set the *printer* name to which the plot is sent if no *filename* is specified.

Example: Print to printer named PS_printer using PostScript format.

```
clf ();
surf (peaks);
print -dpswrite -PPS_printer
```

`-RGBImage`

Return an M-by-N-by-3 RGB image of the figure. The size of the image depends on the formatting options. This is similar to taking a screen

capture of the plot, but formatting options may be changed such as the resolution or monochrome/color.

Example: Get the pixels of a figure image.

```
clf ();
surf (peaks);
rgb = print ("-RGBImage");
```

`-image | -opengl`
`-vector | -painters`

Specifies whether the pixel-based renderer (`-image` or `-opengl`) or vector-based renderer (`-vector` or `-painters`) is used. This is equivalent to changing the figure's `"Renderer"` property. When the figure `"RendererMode"` property is `"auto"` (the default) Octave will use the `"opengl"` renderer for raster formats (e.g., JPEG) and `"painters"` for vector formats (e.g., PDF). These options are only supported for the `"qt"` graphics toolkit.

`-svgconvert (default)`
`-nosvgconvert`

When using the `-painters` renderer, this enables or disables the SVG based backend toolchain with enhanced characteristics:

Font handling:

For interpreters `"none"` and `"tex"`, the actual font is embedded in the output file which allows for printing arbitrary characters and fonts in all vector formats.

Strings using the `"latex"` interpreter, are rendered using path objects. This looks good but note that textual info (font, characters...) are lost.

Output Simplification:

By default, the option `-painters` renders patch and surface objects using assemblies of triangles. This may lead to anti-aliasing artifacts when viewing the file. The `-svgconvert` option reconstructs polygons in order to avoid those artifacts (particularly for 2-D figures).

Transparency:

Allows for printing transparent graphics objects in PDF format. For PostScript formats the presence of any transparent object will cause the output to be rasterized.

Caution: If Octave was built against Qt version earlier than 5.13, `-svgconvert` may lead to inaccurate rendering of image objects.

`-polymerge`
`-nopolymerge`
`-polymerge-all`

When using the SVG based backend `-svgconvert`, faces are rendered as triangles. In some cases, some viewers might display fine lines where

those triangles share an edge. These options control whether all triangles that share edges are merged into polygons (`-polymerge-all` which might take some time for graphics consisting of many triangles – including line markers), only consecutive polygons are merged (`-polymerge`), or no triangles are merged at all (`-no-polymerge`). By default, only consecutive triangles sharing an edge are merged, unless the printed figure contains patch or surface graphics objects in which case all triangles that are sharing an edge are merged.

`-portrait`
`-landscape`
: Specify the orientation of the plot for printed output. For non-printed output the aspect ratio of the output corresponds to the plot area defined by the `"paperposition"` property in the orientation specified. This option is equivalent to changing the figure's `"paperorientation"` property.

`-fillpage`
`-bestfit`
: When using a page-based format (PDF, PostScript, printer) ignore the `"paperposition"` property and have the plot occupy the entire page. The option `-fillpage` will stretch the plot to occupy the page with 0.25 inch margins all around. The option `-bestfit` will expand the plot to take up as much room as possible on the page **without** distorting the original aspect ratio of the plot.

`-color`
`-mono`
: Color or monochrome output.

`-solid`
`-dashed`
: Force all lines to be solid or dashed, respectively.

`-noui`
: Don't print uicontrol objects such as pushbuttons which may overlay the plot. This is the default behavior and it is not possible to include uicontrol objects in the output without using an external screen capture tool.

`-r`*`NUM`*
: Resolution of bitmaps in dots per inch (DPI). For both metafiles and SVG the default is the screen resolution; for other formats the default is 150 DPI. To specify screen resolution, use `"-r0"`.

  Example: high resolution raster output.

```
clf ();
surf (peaks (), "facelighting", "gouraud");
light ();
print ("-r600", "lit_peaks.png");
```

`-S`*`xsize`*`,`*`ysize`*
: Plot size in pixels for raster formats including PNG, JPEG, PNG, and *unusually* SVG. For all vector formats, including PDF, PS, and EPS, the plot size is specified in points. This option is equivalent to changing the width and height of the output by setting the figure property `paperposition(3:4)`. When using the command form of the print function you must quote the *xsize*,*ysize* option to prevent the Octave in-

terpreter from recognizing the embedded comma (','). For example, by writing `"-S640,480"`.

`-tight`
`-loose` Force a tight or loose bounding box for EPS files. The default is tight.

`-preview` Add a preview to EPS files. Supported formats are:

`-interchange`
: Provide an interchange preview.

`-metafile`
: Provide a metafile preview.

`-pict` Provide a pict preview.

`-tiff` Provide a TIFF preview.

`-append` Append PostScript or PDF output to an existing file of the same type.

`-Ffontname`
`-Ffontname:size`
`-F:size` Use *fontname* and/or *fontsize* for all text. *fontname* is ignored for some devices: fig, etc.

`-ddevice` The available output format is specified by the option *device*, and is one of the following (devices marked with a `'*'` are only available with the Gnuplot toolkit):

Vector Formats

`svg` Scalable Vector Graphics.

`pdf`
`pdfcrop` Portable Document Format. The `pdf` device formats the figure for printing on paper. The size of the surrounding page and the position of the figure inside the page are defined by the [paper* figure properties], page 469.

Use `pdfcrop` if you don't want the surrounding page.

Caution: with `-nosvgconvert` option, PDF inherits the same limitations as PostScript (limited set of fonts and lack of transparency).

`eps(2)`
`epsc(2)` Encapsulated PostScript (level 1 and 2, mono and color).

The OpenGL-based graphics toolkits always generate PostScript level 3.0. They have limited support for text unless using the `-svgconvert` option (the default). Limitations include using only ASCII characters (e.g., no Greek letters) and support for just three base PostScript fonts: Helvetica (the default), Times, or Courier. Any other font will be replaced by Helvetica.

`ps(2)`
`psc(2)` Same as `eps` except that the figure is formatted for printing on paper. The size of the surrounding page and position of the figure inside the page are defined by the [paper* figure properties], page 469.

`pslatex`
`epslatex`
`pdflatex`
`pslatexstandalone`
`epslatexstandalone`
`pdflatexstandalone`
Generate a LaTeX file *`filename`*`.tex` for the text portions of a plot and a file *`filename`*`.(ps|eps|pdf)` for the remaining graphics. The graphics file suffix .ps|eps|pdf is determined by the specified device type. The LaTeX file produced by the '`standalone`' option can be processed directly by LaTeX. The file generated without the '`standalone`' option is intended to be included from another LaTeX document. In either case, the LaTeX file contains an `\includegraphics` command so that the generated graphics file is automatically included when the LaTeX file is processed. The text that is written to the LaTeX file contains the strings **exactly** as they were specified in the plot. If any special characters of the TeX mode interpreter were used, the file must be edited before LaTeX processing. Specifically, the special characters must be enclosed with dollar signs (`$ ... $`), and other characters that are recognized by LaTeX may also need editing (e.g., braces). The '`pdflatex`' device, and any of the '`standalone`' formats, are not available with the Gnuplot toolkit.

`epscairo*`
`pdfcairo*`
`epscairolatex*`
`pdfcairolatex*`
`epscairolatexstandalone*`
`pdfcairolatexstandalone*`
Generate output with Cairo renderer. The devices `epscairo` and `pdfcairo` are synonymous with the `epsc` device. The LaTeX variants generate a LaTeX file, *`filename`*`.tex`, for the text portions of a plot, and an image file, *`filename`*`.(eps|pdf)`, for the graph portion of the plot. The '`standalone`' variants behave as described for '`epslatexstandalone`' above.

`canvas*` Javascript-based drawing on an HTML5 canvas viewable in a web browser.

`emf`
`meta` Microsoft Enhanced Metafile

`fig` XFig. For the Gnuplot graphics toolkit, the additional options `-textspecial` or `-textnormal` (default) can be used to control whether the special flag should be set for the text in the figure.

`latex*`
`eepic*` LaTeX picture environment and extended picture environment.

`tikz`
`tikzstandalone*`
Generate a LaTeX file using PGF/TikZ format. The OpenGL-based toolkits create a PGF file while Gnuplot creates a TikZ file. The '`tikzstandalone`' device produces a LaTeX document which includes the TikZ file.

Raster Formats

`png` Portable Network Graphics

`jpg`
`jpeg` JPEG image

`tif`
`tiff`
`tiffn` TIFF image with LZW compression (tif, tiff) or uncompressed (tiffn).

`gif` GIF image

`pbm` PBMplus

`dumb*` ASCII art

If the device is omitted, it is inferred from the file extension, or if there is no filename then it is sent to the printer as PostScript.

`-d`*`ghostscript_device`*
Additional devices are supported by Ghostscript. Some examples are:

`ljet2p` HP LaserJet IIP

`pcx24b` 24-bit color PCX file format

`ppm` Portable Pixel Map file format

For a complete list of available formats and devices type *`system ("gs -h")`*.

When Ghostscript output is sent to a printer the size is determined by the figure's `"papersize"` property. When the output is sent to a file the size is determined by the plot box defined by the figure's `"paperposition"` property.

`-G`*`ghostscript_command`*
: Specify the command for calling Ghostscript. For Unix the default is `"gs"` and for Windows it is `"gswin32c"`.

`-TextAlphaBits=`*`n`*
`-GraphicsAlphaBits=`*`n`*
: Octave is able to produce output for various printers, bitmaps, and vector formats by using Ghostscript. For bitmap and printer output anti-aliasing is applied using Ghostscript's TextAlphaBits and GraphicsAlphaBits options. The default number of bits are 4 and 1 respectively. Allowed values for *N* are 1, 2, or 4.

`-no-append-file-extension`
: With this option, *filename* is used verbatim. That means no file extension matching the file format is appended automatically.

**See also:** [saveas], page 443, [getframe], page 950, [savefig], page 445, [hgsave], page 444, [orient], page 444, [figure], page 425.

`saveas (`*`h`*`,` *`filename`*`)`
`saveas (`*`h`*`,` *`filename`*`,` *`fmt`*`)`

Save graphics object *h* to the file *filename* in graphics format *fmt*.

If *h* is the handle to a figure object, that figure object is saved. If *h* is the handle to a different graphics object, the figure containing that graphics object is saved.

All device formats accepted by `print` may be used. Common formats are:

`ofig`
: Octave figure file format (default)

`mfig`
: Two files: Octave m-file `filename.m` containing code to open Octave figure file `filename.ofig`

`ps`
: PostScript

`eps`
: Encapsulated PostScript

`pdf`
: Portable Document Format

`jpg`
: JPEG Image

`png`
: Portable Network Graphics image

`emf`
: Enhanced MetaFile

`tif`
: TIFF Image, compressed

If *fmt* is omitted it is extracted from the extension of *filename*. The default format when there is no extension is `"ofig"`.

```
clf ();
surf (peaks);
saveas (1, "figure1.png");
```

**See also:** [print], page 436, [savefig], page 445, [hgsave], page 444, [orient], page 444.

`orient (`*`orientation`*`)`
`orient (`*`hfig`*`, `*`orientation`*`)`
*`orientation`* `= orient ()`
*`orientation`* `= orient (`*`hfig`*`)`

Query or set the print orientation for figure *hfig*.

Valid values for *orientation* are `"portrait"`, `"landscape"`, and `"tall"`.

The `"landscape"` option changes the orientation so the plot width is larger than the plot height. The `"paperposition"` is also modified so that the plot fills the page, while leaving a 0.25 inch border.

The `"tall"` option sets the orientation to `"portrait"` and fills the page with the plot, while leaving a 0.25 inch border.

The `"portrait"` option (default) changes the orientation so the plot height is larger than the plot width. It also restores the default `"paperposition"` property.

When called with no arguments, return the current print orientation.

If the argument *hfig* is omitted, then operate on the current figure returned by `gcf`.

**See also:** [print], page 436, [saveas], page 443.

`print` and `saveas` are used when work on a plot has finished and the output must be in a publication-ready format. During intermediate stages it is often better to save the graphics object and all of its associated information so that changes—to colors, axis limits, marker styles, etc.—can be made easily from within Octave. The `hgsave`/`hgload` commands can be used to save and re-create a graphics object.

`hgsave (`*`filename`*`)`
`hgsave (`*`h`*`, `*`filename`*`)`
`hgsave (`*`h`*`, `*`filename`*`, `*`fmt`*`)`

Save the graphics handle(s) *h* to the file *filename* in the format *fmt*.

If unspecified, *h* is the current figure as returned by `gcf`.

When *filename* does not have an extension the default filename extension `.ofig` will be appended.

If present, *fmt* must be one of the following:

- `-binary`, `-float-binary`
- `-hdf5`, `-float-hdf5`
- `-V7`, `-v7`, `-7`, `-mat7-binary`
- `-V6`, `-v6`, `-6`, `-mat6-binary`
- `-text`
- `-zip`, `-z`

The default format is `-binary` to minimize storage.

Programming Note: When producing graphics for final publication use `print` or `saveas`. When it is important to be able to continue to edit a figure as an Octave object, use `hgsave`/`hgload`.

**See also:** [hgload], page 445, [hdl2struct], page 459, [savefig], page 445, [saveas], page 443, [print], page 436.

`h = hgload (filename)`
`[h, old_prop] = hgload (filename, prop_struct)`

Load the graphics objects in *filename* into a vector of graphics handles *h*.

If *filename* has no extension, Octave will try to find the file with and without the default extension `.ofig`.

If provided, the elements of structure *prop_struct* will be used to override the properties of top-level objects stored in *filename*, and the saved values from *filename* will be stored in *old_prop*. *old_prop* is a cell array matching the size of *h*; each cell contains a structure of the existing property names and values before being overridden.

**See also:** [openfig], page 445, [hgsave], page 444, [struct2hdl], page 459.

`openfig`
`openfig (filename)`
`openfig (..., copies)`
`openfig (..., visibility)`
`h = openfig (...)`

Read saved figure window(s) from *filename* and return graphics handle(s) *h*.

By default, *filename* is `"Untitled.fig"`. If a full path is not specified, the file opened will be the first one encountered in the load path. If *filename* is not found and does not have an extension, a search will take place for the first file in the load path with extension `".fig"` or `".ofig"`, in that order.

*copies* is an optional input indicating whether a new figure should be created (`"new"`) or whether an existing figure may be reused (`"reuse"`). An existing figure may be reused if the `"FileName"` property matches the specified input *filename*. When a figure is reused it becomes the active figure and is shown on top of other figures. If the figure was offscreen, it is re-positioned to be onscreen. The default value for *copies* is `"new"`.

*visibility* is an optional input indicating whether to show the figure (`"visible"`) or not (`"invisible"`). When *visibility* is specified as an input to `openfig` it overrides the visibility setting stored in *filename*.

**See also:** [open], page 1060, [hgload], page 445, [savefig], page 445, [struct2hdl], page 459.

`savefig ()`
`savefig (h)`
`savefig (filename)`
`savefig (h, filename)`
`savefig (h, filename, "compact")`

Save figure windows specified by graphics handle(s) *h* to file *filename*.

If unspecified, *h* is the current figure returned by `gcf`.

If unspecified, *filename* is set to `"Untitled.fig"`. If *filename* does not have an extension then the default extension `".fig"` will be added.

If the optional third input `"compact"` is present then the data will be compressed to save more space.

**See also:** [hgsave], page 444, [hdl2struct], page 459, [openfig], page 445.

### 15.2.10 Interacting with Plots

The user can select points on a plot with the `ginput` function or select the position at which to place text on the plot with the `gtext` function using the mouse.

`[x, y, buttons] = ginput (n)`
`[x, y, buttons] = ginput ()`

Return the position and type of mouse button clicks and/or key strokes in the current figure window.

If *n* is defined, then capture *n* events before returning. When *n* is not defined `ginput` will loop until the return key `RET` is pressed.

The return values *x*, *y* are the coordinates where the mouse was clicked in the units of the current axes. The return value *button* is 1, 2, or 3 for the left, middle, or right button. If a key is pressed the ASCII value is returned in *button*.

Implementation Note: `ginput` is intenteded for 2-D plots. For 3-D plots see the *currentpoint* property of the current axes which can be transformed with knowledge of the current `view` into data units.

**See also:** [gtext], page 446, [waitforbuttonpress], page 446.

`b = waitforbuttonpress ()`

Wait for mouse click or key press over the current figure window.

The return value of *b* is 0 if a mouse button was pressed or 1 if a key was pressed.

**See also:** [waitfor], page 1020, [ginput], page 446, [kbhit], page 294.

`gtext (s)`
`gtext ({s1, s2, ...})`
`gtext ({s1; s2; ...})`
`gtext (..., prop, val, ...)`
`h = gtext (...)`

Place text on the current figure using the mouse.

The string argument *s* may be a character array or a cell array of strings. If *s* has more than one row, each row is used to create a separate text object after a mouse click. For example:

Place a single string after one mouse click

```
gtext ("I clicked here")
```

Place two strings after two mouse clicks

```
gtext ({"I clicked here"; "and there"})
```

Place two strings, each with two lines, after two mouse clicks

```
gtext ({"I clicked", "here"; "and", "there"})
```

Optional *property*/*value* pairs are passed directly to the underlying text objects.

The full list of text object properties is documented at Section 15.3.3.6 [Text Properties], page 489.

The optional return value *h* holds the graphics handle(s) to the created text object(s).

**See also:** [ginput], page 446, [text], page 417.

More sophisticated user interaction mechanisms can be obtained using the ui* family of functions, see Section 35.3 [UI Elements], page 1008.

### 15.2.11 Test Plotting Functions

The functions `sombrero` and `peaks` provide a way to check that plotting is working. Typing either `sombrero` or `peaks` at the Octave prompt should display a three-dimensional plot.

`sombrero ()`
`sombrero (n)`
`z = sombrero (...)`
`[x, y, z] = sombrero (...)`

Plot the familiar 3-D sombrero function.

The function plotted is

$$z = \frac{\sin(\sqrt{(x^2 + y^2)})}{\sqrt{(x^2 + y^2)}}$$

Called without a return argument, `sombrero` plots the surface of the above function over the meshgrid [-8,8] using `surf`.

If $n$ is a scalar the plot is made with $n$ grid lines. The default value for $n$ is 41.

When called with output arguments, return the data for the function evaluated over the meshgrid. This can subsequently be plotted with `surf (x, y, z)`.

**See also:** [peaks], page 447, [meshgrid], page 398, [mesh], page 383, [surf], page 385.

`peaks ()`
`peaks (n)`
`peaks (x, y)`
`z = peaks (...)`
`[x, y, z] = peaks (...)`

Plot a function with lots of local maxima and minima.

The function has the form

$$f(x,y) = 3(1-x)^2 e^{\left(-x^2-(y+1)^2\right)} - 10\left(\frac{x}{5} - x^3 - y^5\right) - \frac{1}{3}e^{\left(-(x+1)^2-y^2\right)}$$

Called without a return argument, `peaks` plots the surface of the above function using `surf`.

If $n$ is a scalar, `peaks` plots the value of the above function on an $n$-by-$n$ mesh over the range [-3,3]. The default value for $n$ is 49.

If $n$ is a vector, then it represents the grid values over which to calculate the function. If $x$ and $y$ are specified then the function value is calculated over the specified grid of vertices.

When called with output arguments, return the data for the function evaluated over the meshgrid. This can subsequently be plotted with `surf (x, y, z)`.

**See also:** [sombrero], page 447, [meshgrid], page 398, [mesh], page 383, [surf], page 385.

## 15.3 Graphics Data Structures

### 15.3.1 Introduction to Graphics Structures

The graphics functions use pointers, which are of class graphics_handle, in order to address the data structures which control visual display. A graphics handle may point to any one of a number of different base object types. These objects are the graphics data structures themselves. The primitive graphic object types are: `figure`, `axes`, `line`, `text`, `patch`, `scatter`, `surface`, `text`, `image`, and `light`.

Each of these objects has a function by the same name, and each of these functions returns a graphics handle pointing to an object of the corresponding type.

In addition, there are several functions which operate on properties of the graphics objects and which also return handles. This includes but is not limited to the following functions: The functions `plot` and `plot3` return a handle pointing to an object of type `line`. The function `subplot` returns a handle pointing to an object of type `axes`. The functions `fill`, `fill3`, `trimesh`, and `trisurf` return a handle pointing to an object of type patch. The function `scatter3` returns a handle to an object of type scatter. The functions `slice`, `surf`, `surfl`, `mesh`, `meshz`, `pcolor`, and `waterfall` each return a handle of type surface. The function `camlight` returns a handle to an object of type light. The functions `area`, `bar`, `barh`, `contour`, `contourf`, `contour3`, `surfc`, `meshc`, `errorbar`, `quiver`, `quiver3`, `stair`, `stem`, `stem3` each return a handle to a complex data structure as documented in Section 15.4.6.1 [Data Sources], page 553.

The graphics objects are arranged in a hierarchy:

1. The root object is returned by `groot` (historically, equivalent to the handle 0). In other words, `get (groot)` returns the properties of the root object.

2. Below the root are `figure` objects.

3. Below the `figure` objects are `axes` or `hggroup` objects.

4. Below the `axes` or `hggroup` objects are `line`, `text`, `patch`, `scatter`, `surface`, `image`, and `light` objects.

It is possible to walk this hierarchical tree by querying the `"parent"` and `"children"` properties of the graphics objects.

Graphics handles may be distinguished from function handles (see Section 11.12.1 [Function Handles], page 248) by means of the function `ishghandle`. `ishghandle` returns true if its argument is a handle of a graphics object. In addition, a figure or axes object may be tested using `isfigure` or `isaxes` respectively. To test for a specific type of graphics handle, such as a patch or line object, use `isgraphics`. The more specific test functions return true only if the argument is both a graphics handle and of the correct type (figure, axes, specified object type).

The `get` and `set` commands are used to obtain and set the values of properties of graphics objects. In addition, the `get` command may be used to obtain property names.

For example, the property `"type"` of the graphics object pointed to by the graphics handle h may be displayed by:

```
get (h, "type")
```

The properties and their current values may be obtained in the form of a structure using `s = get (h)`, where `h` is the handle of a graphics object. If only the names of the properties and the allowed values (for radio properties only) are wanted, one may use `set (h)`.

Thus, for example:

```
h = figure ();
get (h, "type")
⇒ ans = figure
set (h)
⇒
        alphamap:
        beingdeleted:  [ {off} | on ]
        busyaction:  [ cancel | {queue} ]
        buttondownfcn:
        clipping:  [ off | {on} ]
        closerequestfcn:
        color:
        colormap:
        createfcn:
        currentaxes:
        deletefcn:
        dockcontrols:  [ {off} | on ]
        filename:
        graphicssmoothing:  [ off | {on} ]
        handlevisibility:  [ callback | off | {on} ]
        ...
```

The uses of `get` and `set` are further explained in [get], page 457, [set], page 457.

*res* = `isprop (`*obj*`, "`*prop*`")`

Return true if *prop* is a property of the object *obj*.

*obj* may also be an array of objects in which case *res* will be a logical array indicating whether each handle has the property *prop*.

For plotting, *obj* is a handle to a graphics object. Otherwise, *obj* should be an instance of a class. `isprop` reports whether the class defines a property, but `Access` permissions or visibility restrictions (`Hidden = true`) may prevent use by the programmer.

**See also:** [get], page 457, [set], page 457, [properties], page 994, [ismethod], page 975, [isobject], page 974.

### 15.3.2 Graphics Objects

The hierarchy of graphics objects was explained above. See Section 15.3.1 [Introduction to Graphics Structures], page 448. Here the specific objects are described, and the properties contained in these objects are discussed. Keep in mind that graphics objects are always referenced by *handle*.

root — The top level of the hierarchy and the parent of all figure objects. Use `groot` to obtain the handle of the root graphics object.

figure — A figure window.

| | |
|---|---|
| axes | A set of axes. This object is a child of a figure object and may be a parent of line, text, image, patch, surface, or light objects. |
| line | A line in two or three dimensions. |
| text | Text annotations. |
| image | A bitmap image. |
| patch | A filled polygon, currently limited to two dimensions. |
| surface | A three-dimensional surface. |
| light | A light object used for lighting effects on patches and surfaces. |

#### 15.3.2.1 Creating Graphics Objects

You can create any graphics object primitive by calling the function of the same name as the object; In other words, `figure`, `axes`, `line`, `text`, `image`, `patch`, `surface`, and `light` functions. These fundamental graphic objects automatically become children of the current axes object as if `hold on` was in place. Separately, axes will automatically become children of the current figure object and figures will become children of the root object.

If this auto-joining feature is not desired then it is important to call `newplot` first to prepare a new figure and axes for plotting. Alternatively, the easier way is to call a high-level graphics routine which will both create the plot and then populate it with low-level graphics objects. Instead of calling `line`, use `plot`. Or use `surf` instead of `surface`. Or use `fill` or `fill3` instead of `patch`.

`axes ()`
`axes (`*`property`*`, `*`value`*`, ...)`
`axes (`*`hpar`*`, `*`property`*`, `*`value`*`, ...)`
`axes (`*`hax`*`)`
*h* `= axes (...)`

Create a Cartesian axes object and return a handle to it, or set the current axes to *hax*.

Called without any arguments, or with *property*/*value* pairs, construct a new axes. The optional argument *hpar* is a graphics handle specifying the parent for the new axes and may be a figure, uipanel, or uitab.

Called with a single axes handle argument *hax*, the function makes *hax* the current axes (as returned by `gca`). It also makes the figure which contains *hax* the current figure (as returned by `gcf`). Finally, it restacks the parent object's `children` property so that the axes *hax* appears before all other axes handles in the list. This causes *hax* to be displayed on top of any other axes objects (Z-order stacking). In addition it restacks any legend or colorbar objects associated with *hax* so that they are also visible.

Programming Note: The full list of properties is documented at Section 15.3.3.3 [Axes Properties], page 470.

**See also:** [gca], page 455, [set], page 457, [get], page 457.

`line ()`
`line (`*`x`*`, `*`y`*`)`

`line` (*x*, *y*, *z*)
`line` ("*xdata*", *x*, "*ydata*", *y*)
`line` ("*xdata*", *x*, "*ydata*", *y*, "*zdata*", *z*)
`line` (..., *property*, *value*)
`line` (*hax*, ...)
*h* = `line` (...)

Create a line object from *x* and *y* (and possibly *z*) and insert it in the current axes.

In the standard calling form the data *x*, *y*, and *z* may be scalars, vectors, or matrices. In the case of matrix inputs, `line` will attempt to orient scalars and vectors so the results can be plotted. This requires that one of the dimensions of the vector match either the number of rows or the number of columns of the matrix.

In the low-level calling form (50% higher performance) where the data is specified by name (`line ("xdata", x, ...)`) the data must be vectors. If no data is specified (`line ()`) then *x* == *y* = `[0, 1]`.

Multiple property-value pairs may be specified for the line object, but they must appear in pairs.

If called with only *property*/*value* pairs then any unspecified properties use their default values as specified on the root object.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle (or vector of handles) to the line objects created.

Programming Note: The full list of properties is documented at Section 15.3.3.5 [Line Properties], page 485.

The function `line` differs from `plot` in that line objects are inserted in to the current axes without first clearing the plot.

**See also:** [image], page 946, [patch], page 451, [rectangle], page 381, [surface], page 452, [text], page 417.

`patch` ()
`patch` (*x*, *y*, *c*)
`patch` (*x*, *y*, *z*, *c*)
`patch` ("*Faces*", *faces*, "*Vertices*", *verts*, ...)
`patch` (..., "*prop*", *val*, ...)
`patch` (..., *propstruct*, ...)
`patch` (*hax*, ...)
*h* = `patch` (...)

Create patch object in the current axes with vertices at locations (*x*, *y*) and of color *c*.

If the vertices are matrices of size MxN then each polygon patch has M vertices and a total of N polygons will be created. If some polygons do not have M vertices use NaN to represent "no vertex". If the *z* input is present then 3-D patches will be created.

The color argument *c* can take many forms. To create polygons which all share a single color use a string value (e.g., `"r"` for red), a scalar value which is scaled by

`clim` and indexed into the current colormap, or a 3-element RGB vector with the precise TrueColor.

If *c* is a vector of length N then the ith polygon will have a color determined by scaling entry *c*(i) according to `clim` and then indexing into the current colormap. More complicated coloring situations require directly manipulating patch property/value pairs.

Instead of specifying polygons by matrices *x* and *y*, it is possible to present a unique list of vertices and then a list of polygon faces created from those vertices. In this case the `"Vertices"` matrix will be an Nx2 (2-D patch) or Nx3 (3-D patch). The MxN `"Faces"` matrix describes M polygons having N vertices—each row describes a single polygon and each column entry is an index into the `"Vertices"` matrix to identify a vertex. The patch object can be created by directly passing the property/value pairs `"Vertices"`/*verts*, `"Faces"`/*faces* as inputs.

Instead of using property/value pairs, any property can be set by passing a structure *propstruct* with the respective field names.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created patch object.

Programming Notes:

1. The full list of properties is documented at Section 15.3.3.8 [Patch Properties], page 495. Useful patch properties include: `"edgecolor"`, `"facecolor"`, `"faces"`, `"vertices"`, and `"facevertexcdata"`.
2. Unexpected geometry results can occur from mixing x-y-z and face-vertex forms of defining geometry.
3. Unexpected patch color results can occur from using `"cdata"` color definitions with face-vertex defined geometry or `"facevertexcdata"` color definitions with x-y-z defined geometry.

**See also:** [fill], page 364, [get], page 457, [set], page 457.

`surface (`*x*, *y*, *z*, *c*`)`
`surface (`*x*, *y*, *z*`)`
`surface (`*z*, *c*`)`
`surface (`*z*`)`
`surface (...,` *prop*, *val*, `...)`
`surface (`*hax*, `...)`
*h* `= surface (...)`

Create a surface graphic object given matrices *x* and *y* from `meshgrid` and a matrix of values *z* corresponding to the *x* and *y* coordinates of the surface.

If *x* and *y* are vectors, then a typical vertex is (*x*(j), *y*(i), *z*(i,j)). Thus, columns of *z* correspond to different *x* values and rows of *z* correspond to different *y* values. If only a single input *z* is given then *x* is taken to be `1:columns (`*z*`)` and *y* is `1:rows (`*z*`)`.

Any property/value input pairs are assigned to the surface object.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created surface object.

Programming Note: The full list of properties is documented at Section 15.3.3.10 [Surface Properties], page 505.

**See also:** [surf], page 385, [mesh], page 383, [patch], page 451, [line], page 450.

`light ()`
`light (...,` "*prop*", *val*, `...)`
`light (`*hax*, `...)`
*h* `= light (...)`

Create a light object in the current axes or for axes *hax*.

When a light object is present in an axes object, and the properties `"EdgeLighting"` or `"FaceLighting"` of a `patch` or `surface` object are set to a value other than `"none"`, these objects are drawn with lighting effects. Supported values for Lighting properties are `"none"` (no lighting effects), `"flat"` (faceted look of the objects), and `"gouraud"` (linear interpolation of the lighting effects between the vertices). If the lighting mode is set to `"flat"`, the `"FaceNormals"` property is used for lighting. For `"gouraud"`, the `"VertexNormals"` property is used.

Up to eight light objects are supported per axes. (Implementation dependent)

Lighting is only supported for OpenGL graphic toolkits (i.e., `"fltk"` and `"qt"`).

A light object has the following properties which alter the appearance of the plot.

`"Color"`: The color of the light can be passed as an RGB-vector (e.g., `[1 0 0]` for red) or as a string (e.g., `"r"` for red). The default color is white (`[1 1 1]`).

`"Position"`: The direction from which the light emanates as a 1x3-vector. The default direction is `[1 0 1]`.

`"Style"`: This string defines whether the light emanates from a light source at infinite distance (`"infinite"`) or from a local point source (`"local"`). The default is `"infinite"`.

If the first argument *hax* is an axes handle, then add the light object to this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the created light object.

Programming Note: The full list of properties is documented at Section 15.3.3.11 [Light Properties], page 511.

**See also:** [lighting], page 396, [material], page 396, [patch], page 451, [surface], page 452.

### 15.3.2.2 Handle Functions

To determine whether a variable is a graphics object index, or an index to an axes or figure, use the functions `ishghandle`, `isgraphics`, `isaxes`, and `isfigure`.

*tf* = `ishghandle` (*h*)

Return true if *h* is a graphics handle and false otherwise.

*h* may also be a matrix of handles in which case a logical array is returned that is true where the elements of *h* are graphics handles and false where they are not.

**See also:** [isgraphics], page 454, [isaxes], page 454, [isfigure], page 454, [ishandle], page 454.

*tf* = `isgraphics` (*h*)
*tf* = `isgraphics` (*h*, *type*)

Return true if *h* is a graphics handle (of type *type*) and false otherwise.

When no *type* is specified the function is equivalent to `ishghandle`.

**See also:** [ishghandle], page 453, [ishandle], page 454, [isaxes], page 454, [isfigure], page 454.

*tf* = `ishandle` (*h*)

Return true if *h* is a handle to a graphics or Java object and false otherwise.

*h* may also be a matrix of handles in which case a logical array is returned that is true where the elements of *h* are handles to graphics or Java objects and false where they are not.

Programming Note: It is often more useful to test for a specific object type. To determine if a handle belongs to a graphics object use `ishghandle` or `isgraphics`. To determine if a handle belongs to a Java object use `isjava`.

**See also:** [ishghandle], page 453, [isgraphics], page 454, [isjava], page 1143.

*tf* = `isaxes` (*h*)

Return true if *h* is an axes graphics handle and false otherwise.

If *h* is a matrix then return a logical array which is true where the elements of *h* are axes graphics handles and false where they are not.

**See also:** [isfigure], page 454, [ishghandle], page 453, [isgraphics], page 454.

*tf* = `isfigure` (*h*)

Return true if *h* is a figure graphics handle and false otherwise.

If *h* is a matrix then return a logical array which is true where the elements of *h* are figure graphics handles and false where they are not.

**See also:** [isaxes], page 454, [ishghandle], page 453, [isgraphics], page 454.

The function `gcf` returns an index to the current figure object, or creates one if none exists. Similarly, `gca` returns the current axes object, or creates one (and its parent figure object) if none exists.

*h* = `groot` ()

Return a handle to the root graphics object.

The root graphics object is the ultimate parent of all graphics objects.

In addition, the root object contains information about the graphics system as a whole such as the `ScreenSize`. Use `get (groot)` to find out what information is available.

Defaults for the graphic system as a whole are specified by setting properties of the root graphics object that begin with `"Default"`. For example, to set the default font for all text objects to FreeSans use

```
set (groot, "DefaultTextFontName", "FreeSans")
```

Default properties can be deleted by using `set` with the special property value of `"remove"`. To undo the default font assignment above use

```
set (groot, "DefaultTextFontName", "remove")
```

Programming Note: The root graphics object is identified by the special handle value of 0. At some point this unique value may change, but code can be made resistant to future changes by using `groot` which is guaranteed to always return the root graphics object.

**See also:** [gcf], page 455, [gca], page 455, [get], page 457, [set], page 457.

*h* = `gcf ()`

Return a handle to the current figure.

The current figure is the default target for graphics output. If multiple figures exist, `gcf` returns the last created figure or the last figure that was clicked on with the mouse.

If a current figure does not exist, create one and return its handle. The handle may then be used to examine or set properties of the figure. For example,

```
fplot (@sin, [-10, 10]);
fig = gcf ();
set (fig, "numbertitle", "off", "name", "sin plot")
```

plots a sine wave, finds the handle of the current figure, and then renames the figure window to describe the contents.

Note: To find the current figure without creating a new one if it does not exist, query the `"CurrentFigure"` property of the root graphics object.

```
get (groot, "currentfigure");
```

**See also:** [gca], page 455, [gco], page 456, [gcbf], page 548, [gcbo], page 547, [get], page 457, [set], page 457.

*h* = `gca ()`

Return a handle to the current axes object.

The current axes is the default target for graphics output. In the case of a figure with multiple axes, `gca` returns the last created axes or the last axes that was clicked on with the mouse.

If no current axes object exists, create one and return its handle. The handle may then be used to examine or set properties of the axes. For example,

```
ax = gca ();
set (ax, "position", [0.5, 0.5, 0.5, 0.5]);
```

creates an empty axes object and then changes its location and size in the figure window.

Note: To find the current axes without creating a new axes object if it does not exist, query the `"CurrentAxes"` property of a figure.

```
get (gcf, "currentaxes");
```

**See also:** [gcf], page 455, [gco], page 456, [gcbf], page 548, [gcbo], page 547, [get], page 457, [set], page 457.

*h* = gco ()
*h* = gco (*hfig*)

Return a handle to the current object of the current figure, or a handle to the current object of the figure with handle *hfig*.

The current object of a figure is the object that was last clicked on. It is stored in the `"CurrentObject"` property of the target figure.

If the last mouse click did not occur on any child object of the figure, then the current object is the figure itself.

If no mouse click occurred in the target figure, this function returns an empty matrix.

Programming Note: The value returned by this function is not necessarily the same as the one returned by `gcbo` during callback execution. An executing callback can be interrupted by another callback and the current object may be changed.

**See also:** [gcbo], page 547, [gca], page 455, [gcf], page 455, [gcbf], page 548, [get], page 457, [set], page 457.

The `get` and `set` functions may be used to examine and set properties for graphics objects. For example,

```
get (groot)
    ⇒ ans =
      {
        type = root
        currentfigure = [](0x0)
        children = [](0x0)
        visible = on
        ...
      }
```

returns a structure containing all the properties of the root graphics object. As with all functions in Octave, the structure is returned by value, so modifying it will not modify the internal root object. To do that, you must use the `set` function. Also, note that in this case, the `currentfigure` property is empty, which indicates that there is no current figure window.

The `get` function may also be used to find the value of a single property. For example,

```
get (gca (), "xlim")
    ⇒ [ 0 1 ]
```

returns the range of the x-axis for the current axes object in the current figure.

To set graphics object properties, use the set function. For example,

```
set (gca (), "xlim", [-10, 10]);
```

sets the range of the x-axis for the current axes object in the current figure to '`[-10, 10]`'.

Default property values can also be queried if the `set` function is called without a value argument. When only one argument is given (a graphic handle) then a structure with defaults for all properties of the given object type is returned. For example,

```
set (gca ())
```

returns a structure containing the default property values for axes objects. If `set` is called with two arguments (a graphic handle and a property name) then only the defaults for the requested property are returned.

*val* = get (*h*)
*val* = get (*h*, *p*)

Return the value of the named property *p* from the graphics handle *h*.

If *p* is omitted, return the complete property list for *h*.

If *h* is a vector, return a cell array including the property values or lists respectively.

**See also:** [set], page 457.

set (*h*, *property*, *value*, ...)
set (*h*, {*properties*}, {*values*})
set (*h*, *pv*)
*value_list* = set (*h*, *property*)
*all_value_list* = set (*h*)

Set named property values for the graphics handle (or vector of graphics handles) *h*.

There are three ways to give the property names and values:

- as a comma-separated list of *property*, *value* pairs

  Each *property* is a string containing the property name, each *value* is a value of the appropriate type for the property. When there are multiple handles in *h*, each one is assigned the same *value*. For example:

  ```
  h = plot ([0, 1]);
  set (h, 'color', 'green');
  ```

- as a cell array of strings *properties* containing property names and a cell array *values* containing property values.

  In this case, the number of columns of *values* must match the number of elements in *properties*. The first column of *values* contains values for the first entry in *properties*, etc. The number of rows of *values* must be 1 or match the number of elements of *h*. In the first case, each handle in *h* will be assigned the same values. In the second case, the first handle in *h* will be assigned the values from the first row of *values* and so on. For example:

  ```
  h = plot ([0, 1; 1, 0]);
  set (h, {'color'}, {'green'; 'red'});
  ```

- as a structure *pv*

  This is the same as the first case where the field names of *pv* represent the property names, and the field values give the property values. As with the first case, it is only possible to set one value for a property which will be applied to all handles in *h*. For example:

```
h = plot ([0, 1]);
props.color = 'green';
set (h, props);
```

The three syntaxes for setting properties may appear in any combination.

`set` is also used to query the list of values a named property will take. *`clist`* `= set (`*`h`*`, "property")` will return the list of possible values for `"property"` in the cell list *clist*. If no output variable is used then the list is formatted and printed to the screen.

If no property is specified (*`slist`* `= set (`*`h`*`)`) then a structure *slist* is returned where the fieldnames are the properties of the object *h* and the fields are the list of possible values for each property. If no output variable is used then the list is formatted and printed to the screen.

When querying properties only a single graphics handle *h* for a single graphics object is permitted.

Example Query

```
hf = figure ();
set (hf, "paperorientation")
⇒  [ landscape | {portrait} ]
```

shows the paperorientation property can take two values with the default being `"portrait"`.

**See also:** [get], page 457.

*parent* `= ancestor (`*h*`, `*type*`)`
*parent* `= ancestor (`*h*`, `*type*`, "`*toplevel*`")`

Return the first ancestor of handle object *h* whose type matches *type*, where *type* is a character string.

If *type* is a cell array of strings, return the first parent whose type matches any of the given type strings.

If the handle object *h* itself is of type *type*, return *h*.

If `"toplevel"` is given as a third argument, return the highest parent in the object hierarchy that matches the condition, instead of the first (nearest) one.

**See also:** [findobj], page 542, [findall], page 543, [allchild], page 458.

*h* `= allchild (`*handles*`)`

Find all children, including hidden children, of a graphics object.

This function is similar to `get (h, "children")`, but also returns hidden objects (HandleVisibility = `"off"`).

If *handles* is a scalar, *h* will be a vector. Otherwise, *h* will be a cell matrix of the same size as *handles* and each cell will contain a vector of handles.

**See also:** [findall], page 543, [findobj], page 542, [get], page 457, [set], page 457.

`findfigs ()`

Find all visible figures that are currently off the screen and move them onto the screen.

**See also:** [allchild], page 458, [figure], page 425, [get], page 457, [set], page 457.

Figures can be printed or saved in many graphics formats with `print` and `saveas`. Occasionally, however, it may be useful to save the original Octave handle graphic directly so that further modifications can be made such as modifying a title or legend.

This can be accomplished with the following functions by

```
fig_struct = hdl2struct (gcf);
save myplot.fig -struct fig_struct;
...
fig_struct = load ("myplot.fig");
struct2hdl (fig_struct);
```

`s = hdl2struct (h)`

Return a structure, *s*, whose fields describe the properties of the object, and its children, associated with the handle, *h*.

The fields of the structure *s* are `"type"`, `"handle"`, `"properties"`, `"children"`, and `"special"`.

**See also:** [struct2hdl], page 459, [hgsave], page 444, [findobj], page 542.

`h = struct2hdl (s)`
`h = struct2hdl (s, p)`
`h = struct2hdl (s, p, hilev)`

Construct a graphics handle object *h* from the structure *s*.

The structure must contain the fields `"handle"`, `"type"`, `"children"`, `"properties"`, and `"special"`.

If the handle of an existing figure or axes is specified, *p*, the new object will be created as a child of that object. If no parent handle is provided then a new figure and the necessary children will be constructed using the default values from the root object.

A third boolean argument *hilev* can be passed to specify whether the function should preserve listeners/callbacks, e.g., for legends or hggroups. The default is false.

**See also:** [hdl2struct], page 459, [hgload], page 445, [findobj], page 542.

*hnew* `= copyobj (`*horig*`)`
*hnew* `= copyobj (`*horig*`, `*hparent*`)`

Construct a copy of the graphic objects associated with the handles *horig* and return new handles *hnew* to the new objects.

If a parent handle *hparent* (root, figure, axes, or hggroup) is specified, the copied object will be created as a child of *hparent*.

If *horig* is a vector of handles, and *hparent* is a scalar, then each handle in the vector *hnew* has its `"Parent"` property set to *hparent*. Conversely, if *horig* is a scalar and *hparent* a vector, then each parent object will receive a copy of *horig*. If *horig* and *hparent* are both vectors with the same number of elements then *hnew*`(i)` will have parent *hparent*`(i)`.

**See also:** [struct2hdl], page 459, [hdl2struct], page 459, [findobj], page 542.

### 15.3.3 Graphics Object Properties

In this section the graphics object properties are discussed in detail, starting with the root properties and continuing through the object hierarchy. The documentation about a specific graphics object can be displayed using `doc` function, e.g., `doc ("axes properties")` will show Section 15.3.3.3 [Axes Properties], page 470.

The allowed values for radio (string) properties can be retrieved programmatically or displayed using the one or two argument calling form of the `set` function. See [set], page 457.

Any properties marked as either unused or unimplemented in the following documentation are accepted without error by Octave. Values for those properties are stored in the object but have no effect on the object.

Default property values are enclosed in { }.

#### 15.3.3.1 Root Properties

Properties of the root graphics object:

#### Categories:

[Callback Execution], page 460 | [Command Window Display], page 460 | [Mouse Interaction], page 460 | [Object Identification], page 460 | [Parent/Children], page 461 | [Pointer Information], page 461 | [Screen Information], page 461 | [Unused], page 462

#### Callback Execution

`callbackobject` (read-only): graphics handle, def. `[](0x0)`
: Graphics handle of the current object whose callback is executing.

#### Command Window Display

`commandwindowsize` (read-only): two-element vector, def. `[0 0]`
: The number of columns and rows displayed in a newly created command window.

`fixedwidthfontname`: string, def. `"Courier"`
: Name of the fixed-width font that will be used for graphics objects when the `"fontname"` property is set to `"FixedWidth"`.

#### Mouse Interaction

`contextmenu`: graphics handle, def. `[](0x0)`
: Graphics handle of the uicontextmenu object that is currently associated to this root object.

#### Object Identification

`currentfigure`: graphics handle, def. `[](0x0)`
: Graphics handle of the current figure.

`tag`: string, def. `""`
: A user-defined string to label the graphics object.

`type` (read-only): string
: Class name of the graphics object. `type` is always `"root"`.

`userdata`: Any Octave data, def. `[](0x0)`
: User-defined data to associate with the graphics object.

## Parent/Children

`children` (read-only): vector of graphics handles, def. `[](0x0)`
: Graphics handles of the root's children.

`handlevisibility`: `"callback" | "off" | {"on"}`
: The root object handle is always visible. Changing this setting to `"callback"` or `"off"` has no effect.

`parent`: graphics handle, def. `[](0x0)`
: Root object has no parent graphics object. `parent` is always empty.

`showhiddenhandles`: `{"off"} | "on"`
: If `showhiddenhandles` is `"on"`, all graphics objects handles are visible in their parents' children list, regardless of the value of their `handlevisibility` property.

## Pointer Information

`pointerlocation`: two-element vector, def. `[0 0]`
: Global pointer location tracking is not yet implemented for root objects. `pointerlocation` is unused. `pointerlocation` for root objects will always be [0 0].

`pointerwindow` (read-only): graphics handle, def. `0`
: Pointer window tracking is not yet implemented for root objects. `pointerwindow` is unused. `pointerwindow` value for root objects will always be 0.

## Screen Information

`monitorpositions` (read-only): four-element vector
: Reports the width and height of connected monitors. Note: Octave only partially implements `monitorpositions`. Only information about the primary monitor is stored in `monitorpositions` which is the same information stored in the [`"screensize"` property], page 461.

`screendepth` (read-only): double
: Color depth in bits per pixel of the display.

`screenpixelsperinch` (read-only): double
: The screen resolution of the primary display in units of pixels per inch.

`screensize` (read-only): four-element vector
: Size of the primary display represented as the four-element vector [left, bottom, width, height].

`units`: `"centimeters"` | `"characters"` | `"inches"` | `"normalized"` | {`"pixels"`} | `"points"`
: The unit type used for the [`"monitorpositions"`], page 461, [`"pointerlocation"`], page 461, and [`"screensize"`], page 461, properties.

### Unused

`beingdeleted`: {`"off"`} | `"on"`
: `beingdeleted` is unused.

`busyaction`: `"cancel"` | {`"queue"`}
: `busyaction` is unused.

`buttondownfcn`: string | function handle, def. `[](0x0)`
: `buttondownfcn` is unused.

`clipping`: `"off"` | {`"on"`}
: `clipping` is unused.

`createfcn`: string | function handle, def. `[](0x0)`
: `createfcn` is unused.

`deletefcn`: string | function handle, def. `[](0x0)`
: `deletefcn` is unused.

`hittest`: `"off"` | {`"on"`}
: `hittest` is unused.

`interruptible`: `"off"` | {`"on"`}
: `interruptible` is unused.

`pickableparts`: `"all"` | `"none"` | {`"visible"`}
: `pickableparts` is unused.

`selected`: {`"off"`} | `"on"`
: `selected` is unused.

`selectionhighlight`: `"off"` | {`"on"`}
: `selectionhighlight` is unused.

`visible`: `"off"` | {`"on"`}
: `visible` is unused.

## 15.3.3.2 Figure Properties

Properties of `figure` objects (see [figure], page 425):

### Categories:

[Appearance], page 462 | [Callback Execution], page 463 | [Creation/Deletion], page 463 | [Display], page 464 | [Keyboard Interaction], page 464 | [Mouse Interaction], page 465 | [Object Identification], page 467 | [Object Position], page 468 | [Parent/Children], page 468 | [Printing/Saving], page 469 | [Unused], page 470

## Appearance

`alphamap`: def. 64-by-1 double
: Transparency is not yet implemented for figure objects. `alphamap` is unused.

`color`: colorspec, def. `[1 1 1]`
: Color of the figure background. See Section 15.4.1 [colorspec], page 545.

`colormap`: N-by-3 matrix, def. 256-by-3 double
: A matrix containing the RGB colormap for the current axes.

`graphicssmoothing`: `"off"` | {`"on"`}
: Use smoothing techniques to reduce the appearance of jagged lines.

`name`: string, def. `""`
: Name to be displayed in the figure title bar. The name is displayed to the right of any title determined by the `numbertitle` property.

`numbertitle`: `"off"` | {`"on"`}
: Display `"Figure"` followed by the numerical figure handle value in the figure title bar.

## Callback Execution

`busyaction`: `"cancel"` | {`"queue"`}
: Define how Octave handles the execution of this object's callback properties when it is unable to interrupt another object's executing callback. This is only relevant when the currently executing callback object has its `interruptible` property set to `"off"`. The `busyaction` property of the interrupting callback object indicates whether the interrupting callback is queued (`"queue"` (default)) or discarded (`"cancel"`). See Section 15.4.4 [Callbacks section], page 546.

`interruptible`: `"off"` | {`"on"`}
: Specify whether this object's callback functions may be interrupted by other callbacks. By default `interruptible` is `"on"` and callbacks that make use of `drawnow`, `figure`, `waitfor`, `getframe` or `pause` functions are eventually interrupted. See Section 15.4.4 [Callbacks section], page 546.

## Creation/Deletion

`beingdeleted`: {`"off"`} | `"on"`
: Property indicating that a function has initiated deletion of the object. `beingdeleted` is set to true until the object no longer exists.

`closerequestfcn`: string | function handle, def. `"closereq"`
: Function that is executed when a figure is deleted. See [closereq function], page 431.

  For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`createfcn`: string | function handle, def. `[](0x0)`
: Callback function executed immediately after figure has been created. Function is set by using default property on root object, e.g., `set (groot, "defaultfigurecreatefcn", 'disp ("figure created!")')`.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`deletefcn`: string | function handle, def. `[](0x0)`
Callback function executed immediately before figure is deleted.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Display

`visible`: `"off"` | {`"on"`}
If `visible` is `"off"`, the figure is not rendered on screen.

`windowstate`: `"fullscreen"` | `"maximized"` | `"minimized"` | {`"normal"`}
Window state adjustment is not yet implemented for figure objects. `windowstate` is unused.

`windowstyle`: `"docked"` | `"modal"` | {`"normal"`}
The window style of a figure. One of the following values:

`normal` The window may be unselected and other windows may be shown in front of the window.

`modal` The window will stay on top of all normal figures until it is dismissed.

`docked` Unimplemented.

Changing modes of a visible figure may cause the figure to close and reopen.

## Keyboard Interaction

`keypressfcn`: string | function handle, def. `[](0x0)`
Callback function executed when a key is pressed while the figure has focus. The first argument to the function is the handle of the calling figure. The second argument holds an event structure which has the following members:

`Character:`
The ASCII value of the key

`Modifier:`
A cell array containing strings representing the modifiers pressed with the key.

`Key:` Lowercase description of the key

`Source:` Graphics handle of the object executing the callback function

`EventName:`
`"KeyPress"`

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`keyreleasefcn`: string | function handle, def. `[](0x0)`
: Callback function executed when a key is release while the figure has focus. The first argument to the function is the handle of the calling figure. The second argument holds an event structure which has the following members:

  `Character:`
  : The ASCII value of the key

  `Modifier:`
  : A cell array containing strings representing the modifiers pressed with the key.

  `Key:` Lowercase description of the key

  `Source:` Graphics handle of the object executing the callback function

  `EventName:`
  : `"KeyRelease"`

  For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`windowkeypressfcn`: string | function handle, def. `[](0x0)`
: Function that is executed when a key is pressed and the figure has focus.

  For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`windowkeyreleasefcn`: string | function handle, def. `[](0x0)`
: Function that is executed when a key is released and the figure has focus.

  For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Mouse Interaction

`buttondownfcn`: string | function handle, def. `[](0x0)`
: For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`contextmenu`: graphics handle, def. `[](0x0)`
: Graphics handle of the uicontextmenu object that is currently associated to this figure object.

`currentpoint` (read-only): two-element vector, def. `[0; 0]`
: A 1-by-2 vector which holds the coordinates of the point over which the mouse pointer was when a mouse event occurred. The X and Y coordinates are in units defined by the figure's `units` property and their origin is the lower left corner of the plotting area.

  Events which set `currentpoint` are

  A mouse button was pressed
  : always

  A mouse button was released
  : only if the figure's callback `windowbuttonupfcn` is defined

The pointer was moved while pressing the mouse button (drag)
only if the figure's callback `windowbuttonmotionfcn` is defined

`hittest`: `"off"` | {`"on"`}
Specify whether figure processes mouse events or passes them to ancestors of the object. When enabled, the object may respond to mouse clicks by evaluating the `"buttondownfcn"`, showing the uicontextmenu, and eventually becoming the root `"currentobject"`. This property is only relevant when the object can accept mouse clicks which is determined by the `"pickableparts"` property. See [pickableparts property], page 470.

`menubar`: {`"figure"`} | `"none"`
Control the display of the figure menu bar at the top of the figure.

`pointer`: {`"arrow"`} | `"botl"` | `"botr"` | `"bottom"` | `"circle"` | `"cross"` | `"crosshair"` | `"custom"` | `"fleur"` | `"hand"` | `"ibeam"` | `"left"` | `"right"` | `"top"` | `"topl"` | `"topr"` | `"watch"`
Name of the mouse pointer shape associated with the canvas of the figure. When `pointer` is `"custom"`, the shape is determined by the `pointershapecdata` property.

`pointer` has no effect when the figure is in zoom, pan, or rotate mode. In this case, Octave automatically uses a pointer shape appropriate to the mode.

`pointershapecdata`: 16-by-16 or 32-by-32 Matrix, def. 16-by-16 double
m-by-m matrix defining a custom pointer. Each element defines a pixel with the element (1,1) representing the top-left pixel. A value of 1 is colored black, a value of 2 is colored white, and all other values are rendered as transparent.

`pointershapehotspot`: two-element vector, def. `[1 1]`
For custom pointers only `pointershapehotspot` defines the row and column of the pixel in `pointershapecdata` that is used as the pointer location.

`resize`: `"off"` | {`"on"`}
Control whether the figure can be resized by dragging the window borders and corners using a mouse. When `resize` is `"off"` mouse interactions are disabled but the figure can still be resized by changing its `"position"` property.

`resizefcn`: string | function handle, def. `[](0x0)`
`resizefcn` is deprecated. Use `sizechangedfcn` instead.

`selected`: {`"off"`} | `"on"`
Property indicates whether this figure is selected.

`selectionhighlight`: `"off"` | {`"on"`}
If `selectionhighlight` is `"on"`, then the figure's selection state is visually highlighted.

`selectiontype`: `"alt"` | `"extend"` | {`"normal"`} | `"open"`
Selection type of the latest mouse click.

`selectiontype` may take different values depending on the combination of mouse button and keyboard modifier that were used:

`normal`: Left-click.

`alt:` Right-click or Ctrl+Left-click.

`extend:` Shift+Left-click, Middle click, or combined Left-click and Right-click.

`open:` Double Left-click.

`sizechangedfcn`: string | function handle, def. `[](0x0)`
Callback triggered when the figure window size is changed.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`toolbar`: {`"auto"`} | `"figure"` | `"none"`
Control the display of the toolbar (along the bottom of the menubar) and the status bar. When set to `"auto"`, the display is based on the value of the [`"menubar"`], page 466, property.

`windowbuttondownfcn`: string | function handle, def. `[](0x0)`
See [windowbuttonupfcn property], page 467.

`windowbuttonmotionfcn`: string | function handle, def. `[](0x0)`
See [windowbuttonupfcn property], page 467.

`windowbuttonupfcn`: string | function handle, def. `[](0x0)`
With `windowbuttondownfcn` and `windowbuttonmotionfcn`, the mouse callback functions. These callback functions are called when a mouse button is pressed, dragged, or released respectively. When these callback functions are executed, the `currentpoint` property holds the current coordinates of the cursor.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`windowscrollwheelfcn`: string | function handle, def. `[](0x0)`
Function that is executed when a user manipulates the mouse wheel over this figure. The function is called with two input arguments. The first argument holds the handle of the calling figure. The second argument holds an event structure which has the following members:

`VerticalScrollCount:`
The number of wheel steps, typically 1 when scrolling down and -1 when scrolling up.

`VerticalScrollAmount:`
The number of lines a wheel step should scroll. This value is always 3.

`EventName:`
The event name which is `"WindowScrollWheel"`.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Object Identification

`currentaxes`: graphics handle, def. `[](0x0)`
: Handle to the graphics object of the current axes.

`currentcharacter` (read-only): def. `""`
: Tracking of the last key pressed is not yet implemented for figure objects. `currentcharacter` is unused.

`currentobject` (read-only): graphics handle, def. `[](0x0)`
: Handle to the most recently active graphics object in the figure.

`integerhandle`: `"off"` | {`"on"`}
: Assign the next lowest unused integer as the Figure number.

`nextplot`: {`"add"`} | `"new"` | `"replace"` | `"replacechildren"`
: `nextplot` is used by high level plotting functions to decide what to do with axes already present in the figure. See [newplot function], page 428.

`number` (read-only): double
: Number of the current figure.

`tag`: string, def. `""`
: A user-defined string to label the graphics object.

`type` (read-only): string
: Class name of the graphics object. `type` is always `"figure"`.

`userdata`: Any Octave data, def. `[](0x0)`
: User-defined data to associate with the graphics object.

## Object Position

`dockcontrols`: `"off"` | {`"on"`}
: Interactive figure docking is not yet implemented for figure objects. `dockcontrols` is unused.

`innerposition`: four-element vector, def. `[300 200 560 420]`
: The `"innerposition"` property is the same as the [`"position"` property], page 468.

`outerposition`: four-element vector, def. `[-1 -1 -1 -1]`
: Specify the position and size of the figure including the top menubar and the bottom status bar. The four elements of the vector are the coordinates of the lower left corner and width and height of the figure. See [units property], page 468.

`position`: four-element vector, def. `[300 200 560 420]`
: Specify the position and size of the figure canvas. The four elements of the vector are the coordinates of the lower left corner and width and height of the figure. See [units property], page 468.

`units`: `"centimeters"` | `"characters"` | `"inches"` | `"normalized"` | {`"pixels"`} | `"points"`
: The unit used to compute the [position], page 468, and [outerposition], page 468, properties.

## Parent/Children

`children` (read-only): vector of graphics handles, def. `[](0x0)`
Graphics handles of the figure's children.

`handlevisibility`: `"callback" | "off" | {"on"}`
If `handlevisibility` is `"off"`, the figure's handle is not visible in its parent's `"children"` property.

`parent`: graphics handle, def. `0`
Handle of the parent graphics object.

## Printing/Saving

`filename`: string, def. `""`
The filename used when saving the plot figure.

`inverthardcopy`: `"off" | {"on"}`
Replace the figure and axes background color with white when printing.

`paperorientation`: `"landscape" | {"portrait"}`
The value for the `papersize`, and `paperposition` properties depends upon `paperorientation`. The horizontal and vertical values for `papersize` and `paperposition` reverse order when `paperorientation` is switched between `"portrait"` and `"landscape"`.

`paperposition`: four-element vector, def. `[1.3422 3.3191 5.8156 4.3617]`
Vector `[left bottom width height]` defining the position and size of the figure (in `paperunits` units) on the printed page. The position `[left bottom]` defines the lower left corner of the figure on the page, and the size is defined by `[width height]`. For output formats not implicitly rendered on paper, `width` and `height` define the size of the image and the position information is ignored. Setting `paperposition` also forces the `paperpositionmode` property to be set to `"manual"`.

`paperpositionmode`: {`"auto"`} | `"manual"`
If `paperpositionmode` is set to `"auto"`, the `paperposition` property is automatically computed: the printed figure will have the same size as the on-screen figure and will be centered on the output page. Setting the `paperpositionmode` to `"auto"` does not modify the value of the `paperposition` property.

`papersize`: two-element vector, def. `[8.5000 11.0000]`
Vector `[width height]` defining the size of the paper for printing. Setting the `papersize` property to a value, not associated with one of the defined `papertypes` and consistent with the setting for `paperorientation`, forces the `papertype` property to the value `"<custom>"`. If `papersize` is set to a value associated with a supported `papertype` and consistent with the `paperorientation`, the `papertype` value is modified to the associated value.

`papertype`: `"<custom>"` | `"a"` | `"a0"` | `"a1"` | `"a2"` | `"a3"` | `"a4"` | `"a5"` | `"arch-a"` | `"arch-b"` | `"arch-c"` | `"arch-d"` | `"arch-e"` | `"b"` | `"b0"` | `"b1"` | `"b2"` | `"b3"` | `"b4"` | `"b5"` | `"c"` | `"d"` | `"e"` | `"tabloid"` | `"uslegal"` | {`"usletter"`}

Name of the paper used for printed output. Setting `papertype` also changes `papersize`, while maintaining consistency with the `paperorientation` property.

`paperunits`: `"centimeters"` | {`"inches"`} | `"normalized"` | `"points"`

The unit used to compute the `paperposition` property. The conversion from physical units (e.g., `"inches"`) is dependent on the `screenpixelsperinch` property of the root object.

`renderer`: {`"opengl"`} | `"painters"`

Rendering engine used for printing when `renderermode` is `"manual"`. Setting `renderer` also forces the `renderermode` property to be set to `"manual"`.

`renderermode`: {`"auto"`} | `"manual"`

Control whether the rendering engine used for printing is chosen automatically or specified by the `renderer` property. See [print function], page 436.

### Unused

`clipping`: `"off"` | {`"on"`}

`clipping` is unused.

`pickableparts` (read-only): `"all"` | `"none"` | {`"visible"`}

`pickableparts` is unused.

#### 15.3.3.3 Axes Properties

Properties of `axes` objects (see [axes], page 450):

### Categories:

[Automatic Child Properties], page 470 | [Axes Box Appearance], page 471 | [Axes Grid Appearance], page 475 | [Callback Execution], page 476 | [Camera and View Controls], page 477 | [Color and Transparency], page 477 | [Creation/Deletion], page 478 | [Display], page 478 | [Mouse Interaction], page 479 | [Object Identification], page 480 | [Object Position], page 480 | [Parent/Children], page 481 | [Text Appearance], page 481

### Automatic Child Properties

`colororder`: N-by-3 RGB matrix, def. 7-by-3 double

RGB values used by plot function for child object coloring.

`colororderindex`: integer, def. `1`

Index of the next color from the [`"colororder"` property], page 470, to be used by Axes-child objects.

`linestyleorder`: def. `"-"`

List of linestyles to be used in order by axes child objects, specified as a cell array of line specification strings. Note that the linestyle is only incremented

after cycling through the full ["colororder"], page 470, list. See Section 15.4.2 [Line Styles], page 545.

`linestyleorderindex`: whole number scalar, def. `1`
: Index of the next linestyle from the ["linestyleorder" property], page 470, to be used by Axes-child objects.

`nextseriesindex` (read-only): whole number scalar, def. `1`
: Current index value into the ["colororder"], page 470, and ["linestyleorder"], page 470, properties, indicating the item that will be used by the next child object

## Axes Box Appearance

`box`: {`"off"`} | `"on"`
: Control whether the axes has a surrounding box.

`boxstyle`: {`"back"`} | `"full"`
: For 3-D axes, control whether the `"full"` box is drawn or only the 3 `"back"` axes.

`color`: colorspec, def. `[1 1 1]`
: Color of the axes background. See Section 15.4.1 [colorspec], page 545.

`dataaspectratio`: three-element vector, def. `[1 1 1]`
: Specify the relative height and width of the data displayed in the axes. Setting `dataaspectratio` to `[1, 2]` causes the length of one unit as displayed on the x-axis to be the same as the length of 2 units on the y-axis. See [daspect function], page 408. Setting `dataaspectratio` also forces the `dataaspectratiomode` property to be set to `"manual"`.

`dataaspectratiomode`: {`"auto"`} | `"manual"`
: Current state of the data aspect ratio mode, either manually set by the ["dataaspectratio" property], page 471, or automatically set by Octave in combination with other display properties to fit the data in the current view.

`layer`: {`"bottom"`} | `"top"`
: Control whether the axes is drawn below child graphics objects (ticks, labels, etc. covered by plotted objects) or above.

`layout` (read-only): def. `[](0x0)`
: Tiled and gridded chart layout is not yet implemented for axes objects. `layout` is unused.

`linewidth`: scalar, def. `0.5000`
: Width of the main axes lines.

`tickdir`: `"both"` | {`"in"`} | `"none"` | `"out"`
: Control whether axes tick marks project `"in"` to the plot box or `"out"`. The value `"both"` will draw tick marks both in and out. The value `"none"` means no tick marks will be drawn, although tick labels will still be rendered. Setting `tickdir` also forces the `tickdirmode` property to be set to `"manual"`. Note that the default is `"in"` for 2-D and `"out"` for 3-D plots.

`tickdirmode`: {`"auto"`} | `"manual"`
: Current state of the tickdir mode, either manually set by the ["tickdir" property], page 471, or automatically set to the default for the current view.

`ticklength`: two-element vector, def. `[0.010000 0.025000]`
: Two-element vector `[2Dlen 3Dlen]` specifying the length of the tickmarks relative to the longest visible axis.

`toolbar` (read-only): def. `[](0x0)`
: AxesToolbar objects is not yet implemented for axes objects. `toolbar` is unused.

`xaxis` (read-only): def. `[](0x0)`
: Axes Ruler objects is not yet implemented for axes objects. `xaxis` is unused.

`xaxislocation`: {`"bottom"`} | `"origin"` | `"top"`
: Control the x-axis location.

`xcolor`: {colorspec} | `"none"`, def. `[0.1500 0.1500 0.1500]`
: Color of the x-axis. See Section 15.4.1 [colorspec], page 545. Setting `xcolor` also forces the `xcolormode` property to be set to `"manual"`.

`xcolormode`: {`"auto"`} | `"manual"`
: Current state of the setting determining the color that is applied to the x-axis grid lines. If set to `"auto"` and/or the ["gridcolormode" property], page 475, is set to `"manual"`, the x-axis grid color will be defined by the ["gridcolor" property], page 475. Otherwise the x-axis grid color will be defined by the ["xcolor" property], page 472.

`xdir`: {`"normal"`} | `"reverse"`
: Direction of the x-axis: `"normal"` is left to right in default 2-D and 3-D views.

`xlim`: two-element vector, def. `[0 1]`
: Two-element vector `[xmin xmax]` specifying the limits for the x-axis. Setting `xlim` also forces the `xlimmode` property to be set to `"manual"`. See [xlim function], page 369.

`xlimitmethod`: `"padded"` | {`"tickaligned"`} | `"tight"`
: Method used to determine the x-axis limits when the `xlimmode` property is `"auto"`. The default value, `"tickaligned"` makes limits align with the closest ticks. With value `"tight"` the limits are adjusted to enclose all the graphics objects in the axes, while with value `"padded"`, an additional margin of about 7% of the data extent is added around the objects. See [axis function], page 366.

`xlimmode`: {`"auto"`} | `"manual"`
: Current state of the x-axis limit selection method, either manually set with the ["xlim" property], page 472, or automatically set to span the plotted data according to the ["xlimitmethod" property], page 472.

`xminortick`: {`"off"`} | `"on"`
: Control whether minor x tick marks are displayed.

`xminortickvalues`: def. `[](0x0)`
: Position of minor tick marks. Setting `xminortickvalues` also forces the `xminortickvaluesmode` property to be set to `"manual"`.

`xminortickvaluesmode`: {`"auto"`} | `"manual"`
: Setting to determine whether the xminortick locations and spacing are set automatically by Octave or manually using the ["xminortickvalues" property], page 472.

`xtick`: vector
: Position of x tick marks. Setting `xtick` also forces the `xtickmode` property to be set to `"manual"`.

`xtickmode`: {`"auto"`} | `"manual"`
: Setting to determine whether the xtick locations and spacing are set automatically by Octave or manually using the ["xtick" property], page 473.

`yaxis` (read-only): def. `[](0x0)`
: Axes Ruler objects is not yet implemented for axes objects. `yaxis` is unused.

`yaxislocation`: {`"left"`} | `"origin"` | `"right"`
: Control the y-axis location.

`ycolor`: {colorspec} | `"none"`, def. `[0.1500 0.1500 0.1500]`
: Color of the y-axis. See Section 15.4.1 [colorspec], page 545.

`ycolormode`: {`"auto"`} | `"manual"`
: Current state of the setting determining the color that is applied to the y-axis grid lines. If set to `"auto"` and/or the ["gridcolormode" property], page 475, is set to `"manual"`, the y-axis grid color will be defined by the ["gridcolor" property], page 475. Otherwise the y-axis grid color will be defined by the ["ycolor" property], page 473.

`ydir`: {`"normal"`} | `"reverse"`
: Direction of the y-axis: `"normal"` is bottom to top in 2-D and front to back in 3-D default views.

`ylim`: two-element vector, def. `[0 1]`
: Two-element vector `[ymin ymax]` specifying the limits for the y-axis. Setting `ylim` also forces the `ylimmode` property to be set to `"manual"`. See [ylim function], page 370.

`ylimitmethod`: `"padded"` | {`"tickaligned"`} | `"tight"`
: Method used to determine the y-axis limits when the `xlimmode` property is `"auto"`. The default value, `"tickaligned"` makes limits align with the closest ticks. With value `"tight"` the limits are adjusted to enclose all the graphics objects in the axes, while with value `"padded"`, an additional margin of about 7% of the data extent is added around the objects. See [axis function], page 366.

`ylimmode`: {`"auto"`} | `"manual"`
: Current state of the y-axis limit selection method, either manually set with the ["ylim" property], page 473, or automatically set to span the plotted data according to the ["ylimitmethod" property], page 473.

`yminortickvalues`: def. `[](0x0)`
: Position of minor tick marks. Setting `yminortickvalues` also forces the `yminortickvaluesmode` property to be set to `"manual"`.

`yminortickvaluesmode`: {`"auto"`} | `"manual"`
: Setting to determine whether the yminortick locations and spacing are set automatically by Octave or manually using the ["yminortickvalues" property], page 473.

`ytick`: vector
: Position of y tick marks. Setting `ytick` also forces the `ytickmode` property to be set to `"manual"`.

`ytickmode`: {`"auto"`} | `"manual"`
: Setting to determine whether the ytick locations and spacing are set automatically by Octave or manually using the ["ytick" property], page 474.

`zaxis` (read-only): def. `[](0x0)`
: Axes Ruler objects is not yet implemented for axes objects. `zaxis` is unused.

`zcolor`: {colorspec} | `"none"`, def. `[0.1500 0.1500 0.1500]`
: Color of the z-axis. See Section 15.4.1 [colorspec], page 545.

`zcolormode`: {`"auto"`} | `"manual"`
: Current state of the setting determining the color that is applied to the z-axis grid lines. If set to `"auto"` and/or the ["gridcolormode" property], page 475, is set to `"manual"`, the z-axis grid color will be defined by the ["gridcolor" property], page 475. Otherwise the z-axis grid color will be defined by the ["zcolor" property], page 474.

`zdir`: {`"normal"`} | `"reverse"`
: Direction of the z-axis: `"normal"` is bottom to top in default 3-D views.

`zlim`: two-element vector, def. `[0 1]`
: Two-element vector `[zmin zmax]` specifying the limits for the z-axis. Setting `zlim` also forces the `zlimmode` property to be set to `"manual"`. See [zlim function], page 371.

`zlimitmethod`: `"padded"` | {`"tickaligned"`} | `"tight"`
: Method used to determine the z-axis limits when the `xlimmode` property is `"auto"`. The default value, `"tickaligned"` makes limits align with the closest ticks. With value `"tight"` the limits are adjusted to enclose all the graphics objects in the axes, while with value `"padded"`, an additional margin of about 7% of the data extent is added around the objects. See [axis function], page 366.

`zlimmode`: {`"auto"`} | `"manual"`
: Current state of the z-axis limit selection method, either manually set with the ["zlim" property], page 474, or automatically set to span the plotted data according to the ["zlimitmethod" property], page 474.

`zminortick`: {`"off"`} | `"on"`
: Control whether minor z tick marks are displayed.

`zminortickvalues`: def. `[](0x0)`
: Position of minor tick marks. Setting `zminortickvalues` also forces the `zminortickvaluesmode` property to be set to `"manual"`.

`zminortickvaluesmode`: {`"auto"`} | `"manual"`
: Setting to determine whether the zminortick locations and spacing are set automatically by Octave or manually using the [`"zminortickvalues"` property], page 474.

`ztick`: vector
: Position of z tick marks. Setting `ztick` also forces the `ztickmode` property to be set to `"manual"`.

`ztickmode`: {`"auto"`} | `"manual"`
: Setting to determine whether the ztick locations and spacing are set automatically by Octave or manually using the [`"ztick"` property], page 475.

## Axes Grid Appearance

`gridalpha`: def. `0.1500`
: Transparency is not yet implemented for axes objects. `gridalpha` is unused.

`gridalphamode`: {`"auto"`} | `"manual"`
: Transparency is not yet implemented for axes objects. `gridalphamode` is unused.

`gridcolor`: {colorspec} | `"none"`, def. `[0.1500 0.1500 0.1500]`
: Color of the major grid lines. See Section 15.4.1 [colorspec], page 545. Setting `gridcolor` also forces the `gridcolormode` property to be set to `"manual"`.

`gridcolormode`: {`"auto"`} | `"manual"`
: Current state of the gridcolor mode, either manually set by the [`"gridcolor"` property], page 475, or automatically set by Octave to the default value.

`gridlinestyle`: {`"-"`} | `"--"` | `"-."` | `":"` | `"none"`
: See Section 15.4.2 [Line Styles], page 545.

`innerposition`: four-element vector, def. `[0.1300 0.1100 0.7750 0.8150]`
: The `"innerposition"` property is the same as the [`"position"` property], page 480.

`minorgridalpha`: def. `0.2500`
: Transparency is not yet implemented for axes objects. `minorgridalpha` is unused.

`minorgridalphamode`: {`"auto"`} | `"manual"`
: Transparency is not yet implemented for axes objects. `minorgridalphamode` is unused.

`minorgridcolor`: {colorspec} | `"none"`, def. `[0.1000 0.1000 0.1000]`
: Color of the minor grid lines. See Section 15.4.1 [colorspec], page 545. Setting `minorgridcolor` also forces the `minorgridcolormode` property to be set to `"manual"`.

`minorgridcolormode`: {`"auto"`} | `"manual"`
: Current state of the minorgridcolor mode, either manually set by the [`"minorgridcolor"` property], page 475, or automatically set by Octave to the default value.

`minorgridlinestyle`: `"-"` | `"--"` | `"-."` | {`":"`} | `"none"`
: See Section 15.4.2 [Line Styles], page 545.

`xgrid`: {`"off"`} | `"on"`
: Control whether major x grid lines are displayed.

`xminorgrid`: {`"off"`} | `"on"`
: Control whether minor x grid lines are displayed.

`xscale`: {`"linear"`} | `"log"`
: Set the x-axis to a linear or logarithmic scale.

`ygrid`: {`"off"`} | `"on"`
: Control whether major y grid lines are displayed.

`yminorgrid`: {`"off"`} | `"on"`
: Control whether minor y grid lines are displayed.

`yminortick`: {`"off"`} | `"on"`
: Control whether minor y tick marks are displayed.

`yscale`: {`"linear"`} | `"log"`
: Set the y-axis to a linear or logarithmic scale.

`zgrid`: {`"off"`} | `"on"`
: Control whether major z grid lines are displayed.

`zminorgrid`: {`"off"`} | `"on"`
: Control whether minor z grid lines are displayed.

`zscale`: {`"linear"`} | `"log"`
: Set the z-axis to a linear or logarithmic scale.

## Callback Execution

`busyaction`: `"cancel"` | {`"queue"`}
: Define how Octave handles the execution of this object's callback properties when it is unable to interrupt another object's executing callback. This is only relevant when the currently executing callback object has its `interruptible` property set to `"off"`. The `busyaction` property of the interrupting callback object indicates whether the interrupting callback is queued (`"queue"` (default)) or discarded (`"cancel"`). See Section 15.4.4 [Callbacks section], page 546.

`interactions` (read-only): def. `[](0x0)`
: Interaction objects is not yet implemented for axes objects. `interactions` is unused.

`interruptible`: `"off"` | {`"on"`}
: Specify whether this object's callback functions may be interrupted by other callbacks. By default `interruptible` is `"on"` and callbacks that make use of

`drawnow`, `figure`, `waitfor`, `getframe` or `pause` functions are eventually interrupted. See Section 15.4.4 [Callbacks section], page 546.

## Camera and View Controls

`cameraposition`: three-element vector, def. `[0.5000 0.5000 9.1603]`
Coordinates of the camera position viewing the axes. Setting `cameraposition` also forces the `cameracapositionmode` property to be set to `"manual"`.

`camerapositionmode`: {`"auto"`} | `"manual"`
Current state of the camera position property, whether automatically set according to the [view function], page 400, or manually set with the [`"cameraposition"` property], page 477.

`cameratarget`: three-element vector, def. `[0.5000 0.5000 0.5000]`
Coordinates of the point at which the viewing camera is aimed. Setting `cameratarget` also forces the `cameratargetmode` property to be set to `"manual"`.

`cameratargetmode`: {`"auto"`} | `"manual"`
Current state of camera target property, either manually set with the [`"cameratarget"` property], page 477, or automatically positioned at the center of the axes plot area.

`cameraupvector`: three-element vector, def. `[0 1 0]`
A 3-element vector defining the upward direction of the current view. Note that the default is [0 1 0] for 2-D plots and [0 0 1] for 3-D plots. Setting `cameraupvector` also forces the `cameraupvectormode` property to be set to `"manual"`.

`cameraupvectormode`: {`"auto"`} | `"manual"`
Current state of camera up vector property, set to manual when the [`"cameraupvector"` property], page 477, is used to change the vector from the 2-D or 3-D default values.

`cameraviewangle`: scalar, def. `6.6086`
The camera's field of view defined as an angle between 0 and 180 degrees. Setting `cameraviewangle` also forces the `cameraviewanglemode` property to be set to `"manual"`.

`cameraviewanglemode`: {`"auto"`} | `"manual"`
Current state of the camera view angle property, either manually set with the [`"cameraviewangle"` property], page 477, or automatically set by Octave to include all visible objects.

`projection`: {`"orthographic"`} | `"perspective"`
Orthographic/perspective projection adjustment is not yet implemented for axes objects. `projection` is unused.

`view`: two-element vector, def. `[0 90]`
Two-element vector `[azimuth elevation]` specifying the viewpoint for three-dimensional plots.

## Color and Transparency

`alim`: def. `[0 1]`
: Transparency is not yet implemented for axes objects. `alim` is unused.

`alimmode`: {`"auto"`} | `"manual"`
: Transparency is not yet implemented for axes objects. `alimmode` is unused.

`alphamap`: def. `[](0x0)`
: Transparency is not yet implemented for axes objects. `alphamap` is unused.

`alphascale`: {`"linear"`} | `"log"`
: Transparency is not yet implemented for axes objects. `alphascale` is unused.

`ambientlightcolor`: def. `[1 1 1]`
: Uniform background axes lighting is not yet implemented for axes objects. `ambientlightcolor` is unused.

`clim`: two-element vector, def. `[0 1]`
: Define limits for the color axis of axes children that have the `cdata` property. Setting `clim` also forces the `climmode` property to be set to `"manual"`.

`climmode`: {`"auto"`} | `"manual"`
: Current state of the color limit mode, either manually set by the ["clim" property], page 478, or automatically set by Octave to the minimum and maximum `cdata` values of axes's children.

`colormap`: N-by-3 matrix, def. 256-by-3 double
: A matrix containing the RGB colormap for this axes object.

`colorscale`: {`"linear"`} | `"log"`
: Automatic linear/log color scaling is not yet implemented for axes objects. `colorscale` is unused.

## Creation/Deletion

`beingdeleted`: {`"off"`} | `"on"`
: Property indicating that a function has initiated deletion of the object. `beingdeleted` is set to true until the object no longer exists.

`createfcn`: string | function handle, def. `[](0x0)`
: Callback function executed immediately after axes has been created. Function is set by using default property on root object, e.g., `set (groot, "defaultaxescreatefcn", 'disp ("axes created!")')`.

  For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`deletefcn`: string | function handle, def. `[](0x0)`
: Callback function executed immediately before axes is deleted.

  For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Display

`clipping`: `"off"` | {`"on"`}
: Axes-children clipping control is not yet implemented for axes objects. `clipping` is unused.

`clippingstyle`: {`"3dbox"`} | `"rectangle"`
: Axes-children clipping control is not yet implemented for axes objects. `clippingstyle` is unused.

`visible`: `"off"` | {`"on"`}
: If `visible` is `"off"`, the axes is not rendered on screen.

## Mouse Interaction

`buttondownfcn`: string | function handle, def. `[](0x0)`
: For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`contextmenu`: graphics handle, def. `[](0x0)`
: Graphics handle of the uicontextmenu object that is currently associated to this axes object.

`currentpoint`: 2-by-3 matrix, def. 2-by-3 double
: Matrix `[xf, yf, zf; xb, yb, zb]` which holds the coordinates (in axes data units) of the point over which the mouse pointer was when the mouse button was pressed. If a mouse callback function is defined, `currentpoint` holds the pointer coordinates at the time the mouse button was pressed. For 3-D plots, the first row of the returned matrix specifies the point nearest to the current camera position and the second row the furthest point. The two points forms a line which is perpendicular to the screen.

`hittest`: `"off"` | {`"on"`}
: Specify whether axes processes mouse events or passes them to ancestors of the object. When enabled, the object may respond to mouse clicks by evaluating the `"buttondownfcn"`, showing the uicontextmenu, and eventually becoming the root `"currentobject"`. This property is only relevant when the object can accept mouse clicks which is determined by the `"pickableparts"` property. See [pickableparts property], page 479.

`mousewheelzoom`: scalar in the range (0, 1), def. `0.5000`
: Fraction of axes limits to zoom for each wheel movement.

`pickableparts`: `"all"` | `"none"` | {`"visible"`}
: Specify whether axes will accept mouse clicks. By default, `pickableparts` is `"visible"` and only visible parts of the axes or its children may react to mouse clicks. When `pickableparts` is `"all"` both visible and invisible parts (or children) may react to mouse clicks. When `pickableparts` is `"none"` mouse clicks on the object are ignored and transmitted to any objects underneath this one. When an object is configured to accept mouse clicks the `"hittest"` property will determine how they are processed. See [hittest property], page 479.

`selected`: {`"off"`} | `"on"`
: Property indicates whether this axes is selected.

`selectionhighlight`: `"off"` | {`"on"`}
: If `selectionhighlight` is `"on"`, then the axes's selection state is visually highlighted.

## Object Identification

`nextplot`: `"add"` | {`"replace"`} | `"replacechildren"`
: `nextplot` is used by high level plotting functions to decide what to do with graphics objects already present in the axes. See [newplot function], page 428. The state of `nextplot` is typically controlled using the `hold` function. See [hold function], page 429.

`tag`: string, def. `""`
: A user-defined string to label the graphics object.

`title`: graphics handle
: Graphics handle of the title text object.

`type` (read-only): string
: Class name of the graphics object. `type` is always `"axes"`.

`userdata`: Any Octave data, def. `[](0x0)`
: User-defined data to associate with the graphics object.

## Object Position

`outerposition`: four-element vector, def. `[0 0 1 1]`
: Specify the position of the plot including titles, axes, and legend. The four elements of the vector are the coordinates of the lower left corner and width and height of the plot, in units normalized to the width and height of the plot window. For example, `[0.2, 0.3, 0.4, 0.5]` sets the lower left corner of the axes at $(0.2, 0.3)$ and the width and height to be 0.4 and 0.5 respectively. See [position property], page 480.

`plotboxaspectratio`: def. `[1 1 1]`
: See [pbaspect function], page 409. Setting `plotboxaspectratio` also forces the `plotboxaspectratiomode` property to be set to `"manual"`.

`plotboxaspectratiomode`: {`"auto"`} | `"manual"`
: Current state of the plot box aspect ratio mode, either manually set by the [`"dataaspectratio"` property], page 471, or automatically set by Octave in combination with other display properties to fit the data in the current view.

`position`: four-element vector, def. `[0.1300 0.1100 0.7750 0.8150]`
: Specify the position of the plot excluding titles, axes, and legend. The four elements of the vector are the coordinates of the lower left corner and width and height of the plot, in units normalized to the width and height of the plot window. For example, `[0.2, 0.3, 0.4, 0.5]` sets the lower left corner of the axes at $(0.2, 0.3)$ and the width and height to be 0.4 and 0.5 respectively. See [outerposition property], page 480.

`positionconstraint`: `"innerposition"` | {`"outerposition"`}
: Specify which of `"innerposition"` or `"outerposition"` properties takes precedence when axes annotations extent changes. See ["innerposition" property], page 475, and ["outerposition" property], page 480.

`units`: `"centimeters"` | `"characters"` | `"inches"` | {`"normalized"`} | `"pixels"` | `"points"`
: Units used to interpret the `"position"`, `"outerposition"`, and `"tightinset"` properties.

## Parent/Children

`children` (read-only): vector of graphics handles, def. `[](0x0)`
: Graphics handles of the axes's children.

`handlevisibility`: `"callback"` | `"off"` | {`"on"`}
: If `handlevisibility` is `"off"`, the axes's handle is not visible in its parent's `"children"` property.

`parent`: graphics handle
: Handle of the parent graphics object.

`sortmethod`: `"childorder"` | {`"depth"`}
: Child display order control is not yet implemented for axes objects. `sortmethod` is unused.

## Text Appearance

`fontangle`: `"italic"` | {`"normal"`}
: Control whether the font is italic or normal.

`fontname`: string, def. `"*"`
: Name of font used for text rendering. When setting this property, the text rendering engine will search for a matching font in your system. If none is found then text is rendered using a default sans serif font (same as the default `"*"` value).

  Programming Note: On systems that don't use FontConfig natively (all but Linux), the font cache is built when Octave is installed. You will need to run `system ("fc-cache -fv")` manually after installing new fonts.

`fontsize`: scalar, def. `10`
: Size of the font used for text rendering. See [fontunits property], page 481. Setting `fontsize` also forces the `fontsizemode` property to be set to `"manual"`.

`fontsizemode`: {`"auto"`} | `"manual"`
: Current state of the fontsize mode, either manually set by the [`"fontsize"` property], page 481, or automatically set by Octave to maintain readability.

`fontsmoothing`: `"off"` | {`"on"`}
: Control whether any text associated with axes is anti-aliased.

`fontunits`: `"centimeters"` | `"inches"` | `"normalized"` | `"pixels"` | {`"points"`}
: Units used to interpret the `"fontsize"` property.

`fontweight`: `"bold"` | {`"normal"`}
: Control the variant of the base font used for text rendering.

`labelfontsizemultiplier`: def. `1.1000`
: Ratio between the x/y/zlabel fontsize and the tick label fontsize.

`legend` (read-only): def. `[](0x0)`
: Legend property control is not yet implemented for axes objects. `legend` is unused. Use the [legend function], page 415, to set legend properties.

`ticklabelinterpreter`: `"latex"` | `"none"` | {`"tex"`}
: Control the way x/y/zticklabel properties are interpreted. See Section 15.2.8 [Use of the `"interpreter"` Property], page 432.

`tightinset` (read-only): four-element vector
: Size of the `[left bottom right top]` margins around the axes that enclose labels and title annotations.

`titlefontsizemultiplier`: positive scalar, def. `1.1000`
: Ratio between the title fontsize and the tick label fontsize.

`titlefontweight`: {`"bold"`} | `"normal"`
: Control variant of base font used for the axes title.

`xlabel`: graphics handle
: Graphics handle of the x label text object.

`xticklabel`: string | cell array of strings, def. 6-by-1 cell
: Labels of x tick marks. Setting `xticklabel` also forces the `xticklabelmode` property to be set to `"manual"`.

`xticklabelmode`: {`"auto"`} | `"manual"`
: Setting to determine whether the xtick labels are set automatically by Octave or manually using the [`"xticklabel"` property], page 482.

`xticklabelrotation`: def. `0`
: Axis label rotation is not yet implemented for axes objects. `xticklabelrotation` is unused.

`ylabel`: graphics handle
: Graphics handle of the y label text object.

`yticklabel`: string | cell array of strings, def. 6-by-1 cell
: Labels of y tick marks. Setting `yticklabel` also forces the `yticklabelmode` property to be set to `"manual"`.

`yticklabelmode`: {`"auto"`} | `"manual"`
: Setting to determine whether the ytick labels are set automatically by Octave or manually using the [`"yticklabel"` property], page 482.

`yticklabelrotation`: def. `0`
: Axis label rotation is not yet implemented for axes objects. `yticklabelrotation` is unused.

`zlabel`: graphics handle
: Graphics handle of the z label text object.

`zticklabel`: string | cell array of strings, def. 6-by-1 cell
: Labels of z tick marks. Setting `zticklabel` also forces the `zticklabelmode` property to be set to `"manual"`.

`zticklabelmode`: {`"auto"`} | `"manual"`
: Setting to determine whether the ztick labels are set automatically by Octave or manually using the [`"zticklabel"` property], page 482.

`zticklabelrotation`: def. `0`
: Axis label rotation is not yet implemented for axes objects. `zticklabelrotation` is unused.

### 15.3.3.4 Legend Properties

Properties of `legend` objects (see [legend], page 415):

## Categories:

[Callback Execution], page 483 | [Layout], page 483 | [Legend Appearance], page 484 | [Object Identification], page 484 | [Object Position], page 484 | [Text Appearance], page 485

## Callback Execution

`itemhitfcn`: def. `[](0x0)`
: Callback function which is executed when a legend item is clicked. See Section 15.4.4 [Callbacks section], page 546.

  The callback function must have the following prototype `fcn (hlegend, evnt)`, where `hlegend` is the legend object handle and `evnt` is a structure with the following fields:

  `Peer` Handle of the plot object to which the clicked item is associated.

  `Region` May be `"icon"` or `"label"` depending on which part of the item is clicked.

  `SelectionType`
  : One of `"normal"`, `"extend"`, `"open"`, or `"alt"`. See [Figure `"selectiontype"`], page 466.

  `Source` Handle of the legend object.

  `EventName`
  : Name is `"ItemHit"`.

## Layout

`autoupdate`: `"off"` | {`"on"`}
: Control whether the number of legend items is updated automatically when objects are added to (or deleted from) the peer axes. For example:

```
## Create a single plot with its legend.
figure ();
plot (1:10);
legend ("Slope 1");
## Add another plot and specify its displayname so that
## the legend is correctly updated.
hold on;
plot ((1:10) * 2, "displayname", "Slope 2");
## Stop automatic updates for further plots.
legend ("autoupdate", "off");
plot ((1:10) * 3);
```

`numcolumns`: scalar integer, def. `1`
: Control the number of columns used in the layout of the legend items. For example:

```
figure ();
plot (rand (30));
legend ("numcolumns", 3);
```

Setting `numcolumns` also forces the `numcolumnsmode` property to be set to `"manual"`.

`orientation`: `"horizontal"` | {`"vertical"`}
: Control whether the legend items are arranged vertically (column-wise) or horizontally (row-wise).

## Legend Appearance

`box`: `"off"` | {`"on"`}
: Control whether the legend has a surrounding box.

`color`: colorspec, def. `[1 1 1]`
: Color of the legend background. See Section 15.4.1 [colorspec], page 545.

`edgecolor`: colorspec, def. `[0.1500 0.1500 0.1500]`
: Control the color of the legend outline.

## Object Identification

`title`: graphics handle
: Graphics handle of the title text object.

## Object Position

`location`: `"best"` | `"bestoutside"` | `"east"` | `"eastoutside"` | `"none"` | `"north"` | {`"northeast"`} | `"northeastoutside"` | `"northoutside"` | `"northwest"` | `"northwestoutside"` | `"south"` | `"southeast"` | `"southeastoutside"` | `"southoutside"` | `"southwest"` | `"southwestoutside"` | `"west"` | `"westoutside"`
: Control the location of the legend.

`position`: four-element vector
: Specify the position of the legend excluding its title. The four elements of the vector are the coordinates of the lower left corner and width and height of the legend. Changing this property also switches the `"location"` to `"none"`.

`units`: `"centimeters"` | `"characters"` | `"inches"` | {`"normalized"`} | `"pixels"` | `"points"`
: Units used to interpret the `"position"` property.

### Text Appearance

`fontangle`: `"italic"` | {`"normal"`}
: Control whether the font is italic or normal.

`fontname`: string, def. `"*"`
: Name of font used for text rendering. When setting this property, the text rendering engine will search for a matching font in your system. If none is found then text is rendered using a default sans serif font (same as the default `"*"` value).

  Programming Note: On systems that don't use FontConfig natively (all but Linux), the font cache is built when Octave is installed. You will need to run `system ("fc-cache -fv")` manually after installing new fonts.

`fontsize`: scalar, def. `9`
: Size of the font used for text rendering. See [fontunits property], page 485. Setting `fontsize` also forces the `fontsizemode` property to be set to `"manual"`.

`fontunits`: `"centimeters"` | `"inches"` | `"normalized"` | `"pixels"` | {`"points"`}
: Units used to interpret the `"fontsize"` property.

`fontweight`: `"bold"` | {`"normal"`}
: Control the variant of the base font used for text rendering.

`string`: string | cell array of strings
: List of labels for the legend items. For example:

```
figure ();
plot (rand (20));
## Let legend choose names automatically
hl = legend ();
## Selectively change some names
str = get (hl, "string");
str(1:5:end) = "Garbage";
set (hl, "string", str);
```

`textcolor`: colorspec, def. `[0 0 0]`
: Control the color of the text strings for legend items.

#### 15.3.3.5 Line Properties

Properties of `line` objects (see [line], page 450):

### Categories:

[Callback Execution], page 486 | [Coordinate Data], page 486 | [Creation/Deletion], page 486 | [Display], page 487 | [Legend Options], page 487 | [Line Appearance], page 487 | [Marker Appearance], page 487 | [Mouse Interaction], page 487 | [Object Identification], page 488 | [Parent/Children], page 488

### Callback Execution

`busyaction`: `"cancel"` | {`"queue"`}

Define how Octave handles the execution of this object's callback properties when it is unable to interrupt another object's executing callback. This is only relevant when the currently executing callback object has its `interruptible` property set to `"off"`. The `busyaction` property of the interrupting callback object indicates whether the interrupting callback is queued (`"queue"` (default)) or discarded (`"cancel"`). See Section 15.4.4 [Callbacks section], page 546.

`interruptible`: `"off"` | {`"on"`}

Specify whether this object's callback functions may be interrupted by other callbacks. By default `interruptible` is `"on"` and callbacks that make use of `drawnow`, `figure`, `waitfor`, `getframe` or `pause` functions are eventually interrupted. See Section 15.4.4 [Callbacks section], page 546.

### Coordinate Data

`xdata`: vector, def. `[0 1]`

Vector of x data to be plotted.

`xdatasource`: string, def. `""`

Name of a vector in the current base workspace to use as x data.

`ydata`: vector, def. `[0 1]`

Vector of y data to be plotted.

`ydatasource`: string, def. `""`

Name of a vector in the current base workspace to use as y data.

`zdata`: vector, def. `[](0x0)`

Vector of z data to be plotted.

`zdatasource`: string, def. `""`

Name of a vector in the current base workspace to use as z data.

### Creation/Deletion

`beingdeleted`: {`"off"`} | `"on"`

Property indicating that a function has initiated deletion of the object. `beingdeleted` is set to true until the object no longer exists.

`createfcn`: string | function handle, def. `[](0x0)`

Callback function executed immediately after line has been created. Function is set by using default property on root object, e.g., `set (groot, "defaultlinecreatefcn", 'disp ("line created!")')`.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`deletefcn`: string | function handle, def. `[](0x0)`
Callback function executed immediately before line is deleted.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Display

`clipping`: `"off"` | {`"on"`}
If `clipping` is `"on"`, the line is clipped in its parent axes limits.

`visible`: `"off"` | {`"on"`}
If `visible` is `"off"`, the line is not rendered on screen.

## Legend Options

`displayname`: string | cell array of strings, def. `""`
Text for the legend entry corresponding to this line.

## Line Appearance

`color`: colorspec, def. `[0 0 0]`
Color of the line object. See Section 15.4.1 [colorspec], page 545.

`linejoin`: `"chamfer"` | `"miter"` | {`"round"`}
Control the shape of the junction of line segments. This property currently only affects the printed output.

`linestyle`: {`"-"`} | `"--"` | `"-."` | `":"` | `"none"`
See Section 15.4.2 [Line Styles], page 545.

`linewidth`: scalar, def. `0.5000`
Width of the line object measured in points.

## Marker Appearance

`marker`: `"*"` | `"+"` | `"."` | `"<"` | `">"` | `"^"` | `"_"` | `"d"` | `"diamond"` | `"h"` | `"hexagram"` | {`"none"`} | `"o"` | `"p"` | `"pentagram"` | `"s"` | `"square"` | `"v"` | `"x"` | `"|"`
Shape of the marker for each data point. See Section 15.4.3 [Marker Styles], page 546.

`markeredgecolor`: {`"auto"`} | `"none"`
Color of the edge of the markers. When set to `"auto"`, the marker edges have the same color as the line. If set to `"none"`, no marker edges are displayed. This property can also be set to any color. See Section 15.4.1 [colorspec], page 545.

`markerfacecolor`: `"auto"` | {`"none"`}
Color of the face of the markers. When set to `"auto"`, the marker faces have the same color as the line. If set to `"none"`, the marker faces are not displayed. This property can also be set to any color. See Section 15.4.1 [colorspec], page 545.

`markersize`: scalar, def. `6`
Size of the markers measured in points.

## Mouse Interaction

`buttondownfcn`: string | function handle, def. `[](0x0)`
: For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`contextmenu`: graphics handle, def. `[](0x0)`
: Graphics handle of the uicontextmenu object that is currently associated to this line object.

`hittest`: `"off"` | {`"on"`}
: Specify whether line processes mouse events or passes them to ancestors of the object. When enabled, the object may respond to mouse clicks by evaluating the `"buttondownfcn"`, showing the uicontextmenu, and eventually becoming the root `"currentobject"`. This property is only relevant when the object can accept mouse clicks which is determined by the `"pickableparts"` property. See [pickableparts property], page 488.

`pickableparts`: `"all"` | `"none"` | {`"visible"`}
: Specify whether line will accept mouse clicks. By default, `pickableparts` is `"visible"` and only visible parts of the line or its children may react to mouse clicks. When `pickableparts` is `"all"` both visible and invisible parts (or children) may react to mouse clicks. When `pickableparts` is `"none"` mouse clicks on the object are ignored and transmitted to any objects underneath this one. When an object is configured to accept mouse clicks the `"hittest"` property will determine how they are processed. See [hittest property], page 488.

`selected`: {`"off"`} | `"on"`
: Property indicates whether this line is selected.

`selectionhighlight`: `"off"` | {`"on"`}
: If `selectionhighlight` is `"on"`, then the line's selection state is visually highlighted.

## Object Identification

`tag`: string, def. `""`
: A user-defined string to label the graphics object.

`type` (read-only): string
: Class name of the graphics object. `type` is always `"line"`.

`userdata`: Any Octave data, def. `[](0x0)`
: User-defined data to associate with the graphics object.

## Parent/Children

`children` (read-only): vector of graphics handles, def. `[](0x0)`
: Child objects for lines is not yet implemented for line objects. `children` is unused.

`handlevisibility`: `"callback"` | `"off"` | {`"on"`}
: If `handlevisibility` is `"off"`, the line's handle is not visible in its parent's `"children"` property.

`parent`: graphics handle
: Handle of the parent graphics object.

#### 15.3.3.6 Text Properties

Properties of `text` objects (see [text], page 417):

### Categories:

[Callback Execution], page 489 | [Creation/Deletion], page 489 | [Display], page 489 | [Mouse Interaction], page 490 | [Object Identification], page 490 | [Object Position], page 490 | [Parent/Children], page 491 | [Text Appearance], page 491 | [Text Box Appearance], page 492

### Callback Execution

`busyaction`: `"cancel"` | {`"queue"`}
: Define how Octave handles the execution of this object's callback properties when it is unable to interrupt another object's executing callback. This is only relevant when the currently executing callback object has its `interruptible` property set to `"off"`. The `busyaction` property of the interrupting callback object indicates whether the interrupting callback is queued (`"queue"` (default)) or discarded (`"cancel"`). See Section 15.4.4 [Callbacks section], page 546.

`interruptible`: `"off"` | {`"on"`}
: Specify whether this object's callback functions may be interrupted by other callbacks. By default `interruptible` is `"on"` and callbacks that make use of `drawnow`, `figure`, `waitfor`, `getframe` or `pause` functions are eventually interrupted. See Section 15.4.4 [Callbacks section], page 546.

### Creation/Deletion

`beingdeleted`: {`"off"`} | `"on"`
: Property indicating that a function has initiated deletion of the object. `beingdeleted` is set to true until the object no longer exists.

`createfcn`: string | function handle, def. `[](0x0)`
: Callback function executed immediately after text has been created. Function is set by using default property on root object, e.g., `set (groot, "defaulttextcreatefcn", 'disp ("text created!")')`.

  For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`deletefcn`: string | function handle, def. `[](0x0)`
: Callback function executed immediately before text is deleted.

  For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

### Display

`clipping`: `"off"` | {`"on"`}
: If `clipping` is `"on"`, the text is clipped in its parent axes limits.

`visible`: `"off"` | {`"on"`}
: If `visible` is `"off"`, the text is not rendered on screen.

## Mouse Interaction

`buttondownfcn`: string | function handle, def. `[](0x0)`
: For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`contextmenu`: graphics handle, def. `[](0x0)`
: Graphics handle of the uicontextmenu object that is currently associated to this text object.

`hittest`: `"off"` | {`"on"`}
: Specify whether text processes mouse events or passes them to ancestors of the object. When enabled, the object may respond to mouse clicks by evaluating the `"buttondownfcn"`, showing the uicontextmenu, and eventually becoming the root `"currentobject"`. This property is only relevant when the object can accept mouse clicks which is determined by the `"pickableparts"` property. See [pickableparts property], page 490.

`pickableparts`: `"all"` | `"none"` | {`"visible"`}
: Specify whether text will accept mouse clicks. By default, `pickableparts` is `"visible"` and only visible parts of the text or its children may react to mouse clicks. When `pickableparts` is `"all"` both visible and invisible parts (or children) may react to mouse clicks. When `pickableparts` is `"none"` mouse clicks on the object are ignored and transmitted to any objects underneath this one. When an object is configured to accept mouse clicks the `"hittest"` property will determine how they are processed. See [hittest property], page 490.

`selected`: {`"off"`} | `"on"`
: Property indicates whether this text is selected.

`selectionhighlight`: `"off"` | {`"on"`}
: If `selectionhighlight` is `"on"`, then the text's selection state is visually highlighted.

## Object Identification

`tag`: string, def. `""`
: A user-defined string to label the graphics object.

`type` (read-only): string
: Class name of the graphics object. `type` is always `"text"`.

`userdata`: Any Octave data, def. `[](0x0)`
: User-defined data to associate with the graphics object.

## Object Position

`extent` (read-only): four-element vector
: Vector `[x0 y0 width height]` indicating the size and location of the text string.

`horizontalalignment`: `"center"` | {`"left"`} | `"right"`
: Specifies the horizontal location of the point set by the [`"position"` property], page 491, relative to the text.

`position`: three-element vector, def. `[0 0 0]`
: Vector `[X0 Y0 Z0]` where X0, Y0, and Z0 indicate the position of the text anchor as defined by `verticalalignment` and `horizontalalignment`.

`rotation`: scalar, def. `0`
: The angle of rotation for the displayed text, measured in degrees.

`units`: `"centimeters"` | {`"data"`} | `"inches"` | `"normalized"` | `"pixels"` | `"points"`
: Sets the measurement unit or method applied to the [`"position"`], page 491, and [`"extent"`], page 490, properties. The default option `"data"` uses the same units and limits as the data plotted in the figure. The `"normalized"` option applies a unitless 0 to 1 scale to the limits along each axis of the displayed data.

`verticalalignment`: `"baseline"` | `"bottom"` | `"cap"` | {`"middle"`} | `"top"`
: Specifies the vertical location of the point set by the [`"position"` property], page 491, relative to the text. Note that `"top"` and `"bottom"` align to the edge of the text box while `"cap"` and `"baseline"` refer to the edges of the text itself.

## Parent/Children

`children` (read-only): vector of graphics handles, def. `[](0x0)`
: text objects have no child objects. `children` is unused.

`handlevisibility`: `"callback"` | `"off"` | {`"on"`}
: If `handlevisibility` is `"off"`, the text's handle is not visible in its parent's `"children"` property.

`parent`: graphics handle
: Handle of the parent graphics object.

## Text Appearance

`color`: colorspec, def. `[0 0 0]`
: Color of the text. See Section 15.4.1 [colorspec], page 545.

`editing`: {`"off"`} | `"on"`
: Interactive text editing is not yet implemented for text objects. `editing` is unused.

`fontangle`: `"italic"` | {`"normal"`}
: Control whether the font is italic or normal.

`fontname`: string, def. `"*"`
: Name of font used for text rendering. When setting this property, the text rendering engine will search for a matching font in your system. If none is found then text is rendered using a default sans serif font (same as the default `"*"` value).

  Programming Note: On systems that don't use FontConfig natively (all but Linux), the font cache is built when Octave is installed. You will need to run `system ("fc-cache -fv")` manually after installing new fonts.

`fontsize`: scalar, def. `10`
: Size of the font used for text rendering. See [fontunits property], page 492.

`fontsmoothing`: `"off"` | {`"on"`}
: Control whether anti-aliasing is used when rendering text.

`fontunits`: `"centimeters"` | `"inches"` | `"normalized"` | `"pixels"` | {`"points"`}
: Units used to interpret the `"fontsize"` property.

`fontweight`: `"bold"` | {`"normal"`}
: Control the variant of the base font used for text rendering.

`interpreter`: `"latex"` | `"none"` | {`"tex"`}
: Control the way the `"string"` property is interpreted. See Section 15.2.8 [Use of the `"interpreter"` Property], page 432.

`string`: string, def. `""`
: The text object string content.

### Text Box Appearance

`backgroundcolor`: colorspec, def. `"none"`
: Color of the background area. See Section 15.4.1 [colorspec], page 545.

`edgecolor`: colorspec, def. `"none"`
: Color of the outline of the background area. See Section 15.4.1 [colorspec], page 545.

`linestyle`: {`"-"`} | `"--"` | `"-."` | `":"` | `"none"`
: Style of the text box outline. See Section 15.4.2 [Line Styles], page 545.

`linewidth`: scalar, def. `0.5000`
: Width of the text box outline.

`margin`: scalar, def. `3`
: Margins between the borders of the background area and the texts. The value is currently interpreted as pixels, regardless of the `"fontunits"` property.

## 15.3.3.7 Image Properties

Properties of `image` objects (see [image], page 946):

### Categories:

[Callback Execution], page 492 | [Creation/Deletion], page 493 | [Display], page 493 | [Image Data], page 493 | [Mouse Interaction], page 494 | [Object Identification], page 494 | [Parent/Children], page 495

### Callback Execution

`busyaction`: `"cancel"` | {`"queue"`}
: Define how Octave handles the execution of this object's callback properties when it is unable to interrupt another object's executing callback. This is only relevant when the currently executing callback object has its `interruptible` property set to `"off"`. The `busyaction` property of the interrupting callback

object indicates whether the interrupting callback is queued (`"queue"` (default)) or discarded (`"cancel"`). See Section 15.4.4 [Callbacks section], page 546.

`interruptible`: `"off"` | {`"on"`}
Specify whether this object's callback functions may be interrupted by other callbacks. By default `interruptible` is `"on"` and callbacks that make use of `drawnow`, `figure`, `waitfor`, `getframe` or `pause` functions are eventually interrupted. See Section 15.4.4 [Callbacks section], page 546.

## Creation/Deletion

`beingdeleted`: {`"off"`} | `"on"`
Property indicating that a function has initiated deletion of the object. `beingdeleted` is set to true until the object no longer exists.

`createfcn`: string | function handle, def. `[](0x0)`
Callback function executed immediately after image has been created. Function is set by using default property on root object, e.g., `set (groot, "defaultimagecreatefcn", 'disp ("image created!")')`.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`deletefcn`: string | function handle, def. `[](0x0)`
Callback function executed immediately before image is deleted.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Display

`clipping`: `"off"` | {`"on"`}
If `clipping` is `"on"`, the image is clipped in its parent axes limits.

`visible`: `"off"` | {`"on"`}
If `visible` is `"off"`, the image is not rendered on screen.

## Image Data

`alphadata`: scalar | matrix, def. `1`
Transparency is not yet implemented for image objects. `alphadata` is unused.

`alphadatamapping`: `"direct"` | {`"none"`} | `"scaled"`
Transparency is not yet implemented for image objects. `alphadatamapping` is unused.

`cdata`: array, def. 64-by-64 double
Color data for the image object. Data is either stored as a 2-D matrix where each element's value determines that pixel's color according to the current colormap, or as a 3-D array where the third dimension contains separate red, blue, and green components for each pixel. For RGB arrays, the values that map to minimum and maximum color value depend on the class of `"cdata"`. Floating point and logical values range from 0 to 1 while integer value range from `intmin` to `intmax` for that integer class.

`cdatamapping`: {`"direct"`} | `"scaled"`

Sets the method for mapping data from the ["cdata" property], page 493, to the current colormap. `"Direct"` mapping selects the color using the `"cdata"` value as an index to the current colormap. `"Scaled"` mapping scales the `"cdata"` values to the range specified in the ["clim" axes property], page 478.

`xdata`: two-element vector, def. `[1 64]`

Two-element vector `[xfirst xlast]` specifying the x coordinates of the centers of the first and last columns of the image.

Setting `xdata` to the empty matrix ([]) will restore the default value of `[1 columns(image)]`.

`ydata`: two-element vector, def. `[1 64]`

Two-element vector `[yfirst ylast]` specifying the y coordinates of the centers of the first and last rows of the image.

Setting `ydata` to the empty matrix ([]) will restore the default value of `[1 rows(image)]`.

## Mouse Interaction

`buttondownfcn`: string | function handle, def. `[](0x0)`

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`contextmenu`: graphics handle, def. `[](0x0)`

Graphics handle of the uicontextmenu object that is currently associated to this image object.

`hittest`: `"off"` | {`"on"`}

Specify whether image processes mouse events or passes them to ancestors of the object. When enabled, the object may respond to mouse clicks by evaluating the `"buttondownfcn"`, showing the uicontextmenu, and eventually becoming the root `"currentobject"`. This property is only relevant when the object can accept mouse clicks which is determined by the `"pickableparts"` property. See [pickableparts property], page 494.

`pickableparts`: `"all"` | `"none"` | {`"visible"`}

Specify whether image will accept mouse clicks. By default, `pickableparts` is `"visible"` and only visible parts of the image or its children may react to mouse clicks. When `pickableparts` is `"all"` both visible and invisible parts (or children) may react to mouse clicks. When `pickableparts` is `"none"` mouse clicks on the object are ignored and transmitted to any objects underneath this one. When an object is configured to accept mouse clicks the `"hittest"` property will determine how they are processed. See [hittest property], page 494.

`selected`: {`"off"`} | `"on"`

Property indicates whether this image is selected.

`selectionhighlight`: `"off"` | {`"on"`}

If `selectionhighlight` is `"on"`, then the image's selection state is visually highlighted.

## Object Identification

`tag`: string, def. `""`
A user-defined string to label the graphics object.

`type` (read-only): string
Class name of the graphics object. `type` is always `"image"`.

`userdata`: Any Octave data, def. `[](0x0)`
User-defined data to associate with the graphics object.

## Parent/Children

`children` (read-only): vector of graphics handles, def. `[](0x0)`
Child objects of Images is not yet implemented for image objects. `children` is unused.

`handlevisibility`: `"callback"` | `"off"` | {`"on"`}
If `handlevisibility` is `"off"`, the image's handle is not visible in its parent's `"children"` property.

`parent`: graphics handle
Handle of the parent graphics object.

### 15.3.3.8 Patch Properties

Properties of `patch` objects (see [patch], page 451):

## Categories:

[Callback Execution], page 495 | [Color and Transparency], page 495 | [Coordinate Data], page 497 | [Creation/Deletion], page 497 | [Display], page 498 | [Legend Options], page 498 | [Lighting], page 498 | [Marker Appearance], page 499 | [Mouse Interaction], page 499 | [Object Identification], page 500 | [Outline Appearance], page 500 | [Parent/Children], page 500

## Callback Execution

`busyaction`: `"cancel"` | {`"queue"`}
Define how Octave handles the execution of this object's callback properties when it is unable to interrupt another object's executing callback. This is only relevant when the currently executing callback object has its `interruptible` property set to `"off"`. The `busyaction` property of the interrupting callback object indicates whether the interrupting callback is queued (`"queue"` (default)) or discarded (`"cancel"`). See Section 15.4.4 [Callbacks section], page 546.

`interruptible`: `"off"` | {`"on"`}
Specify whether this object's callback functions may be interrupted by other callbacks. By default `interruptible` is `"on"` and callbacks that make use of `drawnow`, `figure`, `waitfor`, `getframe` or `pause` functions are eventually interrupted. See Section 15.4.4 [Callbacks section], page 546.

## Color and Transparency

`alphadatamapping`: `"direct"` | `"none"` | {`"scaled"`}

Transparency is not yet implemented for patch objects. `alphadatamapping` is unused.

`cdata`: scalar | matrix | array, def. `[](0x0)`

Data defining the patch object color relative to its x/y/z-coordinate data. Patch color can be defined using indices into the current colormap or as RGB triplets, where the RGB colors are defined along the third dimension. These colors can be separately defined for the entire patch object, for individual faces, or for individual vertices, and is determined by the shape of `"cdata"` as follows:

If `"cdata"` is a scalar index into the current colormap or a 1-by-1-by-3 RGB triplet, it defines the color of all faces and edges.

If the patch object has N faces, and `"cdata"` is a 1-by-N vector of colormap indices or a 1-by-N-by-3 RGB array, it defines the color of each face.

If the patch object has N faces and M vertices per face, and `cdata` is a M-by-N matrix of colormap indices or a M-by-N-by-3 RGB array, it defines the color at each vertex. (The shape of `"cdata"` should match that of [`"xdata"`], page 497, [`"ydata"`], page 497, and [`"zdata"`], page 497.)

`cdatamapping`: `"direct"` | {`"scaled"`}

Sets the method for mapping data from the [`"cdata"`], page 496, or [`"cdata"`], page 496, property to the current colormap. `"Direct"` mapping selects the color using the `"cdata"` or `"facevertexcdata"` value as an index to the current colormap. `"Scaled"` mapping scales the `"cdata"` or `"facevertexcdata"` values to the range specified in the [`"clim"` axes property], page 478.

`facealpha`: scalar | `"flat"` | `"interp"`, def. `1`

Transparency level of the faces of the patch object. Only double values are supported at present where a value of 0 means complete transparency and a value of 1 means solid faces without transparency. Setting the property to `"flat"` or `"interp"` causes the faces to not being rendered. Additionally, the faces are not sorted from back to front which might lead to unexpected results when rendering layered transparent faces.

`facecolor`: {colorspec} | `"none"` | `"flat"` | `"interp"`, def. `[0 0 0]`

Color of the faces of the patch object, specified as either a valid color specification or one of `"none"`, `"flat"`, or `"interp"`. `"flat"` and `"interp"` will set either a single color for each face or a color interpolated across the face's vertices using the color value data stored in either the [`"cdata"`], page 496, or [`"facevertexcdata"`], page 497, properties. See Section 15.4.1 [colorspec], page 545.

`facelighting`: {`"flat"`} | `"gouraud"` | `"none"` | `"phong"`

When set to a value other than `"none"`, the faces of the object are drawn with light and shadow effects. Supported values are `"none"` (no lighting effects), `"flat"` (faceted look), and `"gouraud"` (linear interpolation of the lighting effects between the vertices). `"phong"` is deprecated and has the same effect as `"gouraud"`.

`facevertexalphadata`: def. `[](0x0)`
: Face-Vertex transparency control is not yet implemented for patch objects. `facevertexalphadata` is unused.

`facevertexcdata`: scalar | matrix, def. `[](0x0)`
: Data defining the patch object color relative to its face-vertex data. Patch color can be defined using indices into the current colormap or as RGB triplets, where the RGB colors are defined in the rows of `"facevertexcdata"`. These colors can be separately defined for the entire patch object, for individual faces, or for individual vertices, and is determined by the shape of `"facevertexcdata"` as follows:

  If `facevertexcdata` is a scalar index into the current colormap or a 1-by-3 RGB triplet, it defines the color of all faces and edges.

  If the patch object has N faces, and `facevertexcdata` is a N-by-1 column vector of indices or a N-by-3 RGB matrix, it defines the color of each one of the N faces.

  If the patch object has M vertices, and `facevertexcdata` is a M-by-1 column vector of indices or a M-by-3 RGB matrix, it defines the color at each vertex.

## Coordinate Data

`faces`: vector | matrix, def. `[1 2 3]`
: patch faces connectivity list stored as an M x N matrix, with each of the M faces defined by a row of up to N vertices, and each element contains the row index of a vertex stored in the [vertices property], page 497. Faces with fewer than N vertices use NaN values to fill empty row elements.

`vertices`: vector | matrix, def. 3-by-2 double
: patch vertex list stored as an N x 3 matrix, with each row containing the x, y, and z coordinates of the vector, and used with the [faces property], page 497, to define patch structure.

`xdata`: vector | matrix, def. `[0; 1; 0]`
: patch vertex x-coordinates.

`ydata`: vector | matrix, def. `[1; 1; 0]`
: patch vertex y-coordinates.

`zdata`: vector | matrix, def. `[](0x0)`
: patch vertex z-coordinates.

## Creation/Deletion

`beingdeleted`: {`"off"`} | `"on"`
: Property indicating that a function has initiated deletion of the object. `beingdeleted` is set to true until the object no longer exists.

`createfcn`: string | function handle, def. `[](0x0)`
: Callback function executed immediately after patch has been created. Function is set by using default property on root object, e.g., `set (groot, "defaultpatchcreatefcn", 'disp ("patch created!")')`.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`deletefcn`: string | function handle, def. `[](0x0)`
Callback function executed immediately before patch is deleted.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Display

`clipping`: `"off"` | {`"on"`}
If `clipping` is `"on"`, the patch is clipped in its parent axes limits.

`visible`: `"off"` | {`"on"`}
If `visible` is `"off"`, the patch is not rendered on screen.

## Legend Options

`displayname`: def. `""`
Text of the legend entry corresponding to this patch.

## Lighting

`ambientstrength`: scalar, def. `0.3000`
Strength of the ambient light. Value between 0.0 and 1.0.

`backfacelighting`: `"lit"` | {`"reverselit"`} | `"unlit"`
`"lit"`: The normals are used as is for lighting. `"reverselit"`: The normals are always oriented towards the point of view. `"unlit"`: Faces with normals pointing away from the point of view are unlit.

`diffusestrength`: scalar, def. `0.6000`
Strength of the diffuse reflection. Value between 0.0 (no diffuse reflection) and 1.0 (full diffuse reflection).

`facenormals`: def. `[](0x0)`
Face normals are used for lighting the edges or faces if the `edgelighting` or `facelighting` properties are set to `"flat"`. Setting `facenormals` also forces the `facenormalsmode` property to be set to `"manual"`.

`facenormalsmode`: {`"auto"`} | `"manual"`
If this property is set to `"auto"`, `facenormals` are automatically calculated if the `edgelighting` or `facelighting` property are set to `"flat"` and at least one `light` object is present and visible in the same axes.

`specularcolorreflectance`: scalar, def. `1`
Reflectance for specular color. Value between 0.0 (color of underlying face) and 1.0 (color of light source).

`specularexponent`: scalar, def. `10`
Exponent for the specular reflection. The lower the value, the more the reflection is spread out.

`specularstrength`: scalar, def. `0.9000`
: Strength of the specular reflection. Value between 0.0 (no specular reflection) and 1.0 (full specular reflection).

`vertexnormals`: def. `[](0x0)`
: Vertex normals are used for lighting the edges or faces if the `edgelighting` or `facelighting` properties are set to `"gouraud"`. Setting `vertexnormals` also forces the `vertexnormalsmode` property to be set to `"manual"`.

`vertexnormalsmode`: {`"auto"`} | `"manual"`
: If this property is set to `"auto"`, `vertexnormals` are automatically calculated if the `edgelighting` or `facelighting` property are set to `"gouraud"` and at least one `light` object is present and visible in the same axes.

## Marker Appearance

`marker`: `"*"` | `"+"` | `"."` | `"<"` | `">"` | `"^"` | `"_"` | `"d"` | `"diamond"` | `"h"` | `"hexagram"` | {`"none"`} | `"o"` | `"p"` | `"pentagram"` | `"s"` | `"square"` | `"v"` | `"x"` | `"|"`
: See [line marker property], page 487.

`markeredgecolor`: {`"auto"`} | `"flat"` | `"none"`
: See [line markeredgecolor property], page 487.

`markerfacecolor`: `"auto"` | `"flat"` | {`"none"`}
: See [line markerfacecolor property], page 487.

`markersize`: scalar, def. `6`
: See [line markersize property], page 487.

## Mouse Interaction

`buttondownfcn`: string | function handle, def. `[](0x0)`
: For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`contextmenu`: graphics handle, def. `[](0x0)`
: Graphics handle of the uicontextmenu object that is currently associated to this patch object.

`hittest`: `"off"` | {`"on"`}
: Specify whether patch processes mouse events or passes them to ancestors of the object. When enabled, the object may respond to mouse clicks by evaluating the `"buttondownfcn"`, showing the uicontextmenu, and eventually becoming the root `"currentobject"`. This property is only relevant when the object can accept mouse clicks which is determined by the `"pickableparts"` property. See [pickableparts property], page 499.

`pickableparts`: `"all"` | `"none"` | {`"visible"`}
: Specify whether patch will accept mouse clicks. By default, `pickableparts` is `"visible"` and only visible parts of the patch or its children may react to mouse clicks. When `pickableparts` is `"all"` both visible and invisible parts (or children) may react to mouse clicks. When `pickableparts` is `"none"` mouse clicks on the object are ignored and transmitted to any objects underneath this one.

When an object is configured to accept mouse clicks the `"hittest"` property will determine how they are processed. See [hittest property], page 499.

`selected`: {`"off"`} | `"on"`
: Property indicates whether this patch is selected.

`selectionhighlight`: `"off"` | {`"on"`}
: If `selectionhighlight` is `"on"`, then the patch's selection state is visually highlighted.

## Object Identification

`tag`: string, def. `""`
: A user-defined string to label the graphics object.

`type` (read-only): string
: Class name of the graphics object. `type` is always `"patch"`.

`userdata`: Any Octave data, def. `[](0x0)`
: User-defined data to associate with the graphics object.

## Outline Appearance

`edgealpha`: scalar | matrix, def. `1`
: Transparency is not yet implemented for patch objects. `edgealpha` is unused.

`edgecolor`: colorspec | `"none"` | `"flat"` | `"interp"`, def. `[0 0 0]`
: Color of the edges of the patch object, specified as either a valid color specification or one of `"none"`, `"flat"`, or `"interp"`. `"flat"` and `"interp"` will set either a single color for each edge or a color interpolated between edge's vertices using the color value data stored in [`"cdata"`], page 496. See Section 15.4.1 [colorspec], page 545.

`edgelighting`: `"flat"` | `"gouraud"` | {`"none"`} | `"phong"`
: When set to a value other than `"none"`, the edges of the object are drawn with light and shadow effects. Supported values are `"none"` (no lighting effects), `"flat"` (faceted look), and `"gouraud"` (linear interpolation of the lighting effects between the vertices). `"phong"` is deprecated and has the same effect as `"gouraud"`.

`linestyle`: {`"-"`} | `"--"` | `"-."` | `":"` | `"none"`
: See Section 15.4.2 [Line Styles], page 545.

`linewidth`: scalar, def. `0.5000`
: See [line linewidth property], page 487.

## Parent/Children

`children` (read-only): vector of graphics handles, def. `[](0x0)`
: Child objects for Patch objects is not yet implemented for patch objects. `children` is unused.

`handlevisibility`: `"callback"` | `"off"` | {`"on"`}
: If `handlevisibility` is `"off"`, the patch's handle is not visible in its parent's `"children"` property.

`parent`: graphics handle
: Handle of the parent graphics object.

### 15.3.3.9 Scatter Properties

Properties of `scatter` objects (see [scatter], page 345):

## Categories:

[Callback Execution], page 501 | [Color Data], page 501 | [Coordinate Data], page 501 | [Creation/Deletion], page 503 | [Display], page 503 | [Legend Options], page 503 | [Marker Appearance], page 503 | [Mouse Interaction], page 504 | [Object Identification], page 505 | [Parent/Children], page 505

## Callback Execution

`busyaction`: `"cancel"` | {`"queue"`}
: Define how Octave handles the execution of this object's callback properties when it is unable to interrupt another object's executing callback. This is only relevant when the currently executing callback object has its `interruptible` property set to `"off"`. The `busyaction` property of the interrupting callback object indicates whether the interrupting callback is queued (`"queue"` (default)) or discarded (`"cancel"`). See Section 15.4.4 [Callbacks section], page 546.

`interruptible`: `"off"` | {`"on"`}
: Specify whether this object's callback functions may be interrupted by other callbacks. By default `interruptible` is `"on"` and callbacks that make use of `drawnow`, `figure`, `waitfor`, `getframe` or `pause` functions are eventually interrupted. See Section 15.4.4 [Callbacks section], page 546.

## Color Data

`cdata`: scalar | matrix, def. `[0 0.4470 0.7410]`
: Data defining the scatter object color.

  If `cdata` is a scalar index into the current colormap or a RGB triplet, it defines the color of all scatter markers.

  If `cdata` is an N-by-1 vector of indices or an N-by-3 (RGB) matrix, it defines the color of each one of the N scatter markers.

`cdatamode`: {`"auto"`} | `"manual"`
: If `cdatamode` is `"auto"`, `cdata` is set to the color from the `colororder` of the ancestor axes corresponding to the `seriesindex`.

`cdatasource`: string, def. `""`
: The name of a workspace variable that contains data that will be used for the [`"cdata"` property], page 496. Data is transferred into `"cdata"` using the [refreshdata function], page 553.

`seriesindex`: def. `1`
: Each scatter object in the same axes is assigned an incrementing integer. This corresponds to the index into the `colororder` of the ancestor axes that is used if `cdatamode` is set to `"auto"`.

## Coordinate Data

`latitudedata`: def. `[](0x0)`
: Geographic coordinate scatter plotting is not yet implemented for scatter objects. `latitudedata` is unused.

`latitudedatasource`: def. `""`
: Geographic coordinate scatter plotting is not yet implemented for scatter objects. `latitudedatasource` is unused.

`longitudedata`: def. `[](0x0)`
: Geographic coordinate scatter plotting is not yet implemented for scatter objects. `longitudedata` is unused.

`longitudedatasource`: def. `""`
: Geographic coordinate scatter plotting is not yet implemented for scatter objects. `longitudedatasource` is unused.

`rdata`: def. `[](0x0)`
: Polar coordinates for scatter plotting is not yet implemented for scatter objects. `rdata` is unused.

`rdatasource`: def. `""`
: Polar coordinates for scatter plotting is not yet implemented for scatter objects. `rdatasource` is unused.

`thetadata`: def. `[](0x0)`
: Polar coordinates for scatter plotting is not yet implemented for scatter objects. `thetadata` is unused.

`thetadatasource`: def. `""`
: Polar coordinates for scatter plotting is not yet implemented for scatter objects. `thetadatasource` is unused.

`xdata`: vector, def. `[](0x0)`
: Vector with the x coordinates of the scatter object.

`xdatasource`: string, def. `""`
: The name of a workspace variable that contains data that will be used for the [`"xdata"` property], page 502. Data is transferred into `"xdata"` using the [refreshdata function], page 553.

`ydata`: vector, def. `[](0x0)`
: Vector with the y coordinates of the scatter object.

`ydatasource`: string, def. `""`
: The name of a workspace variable that contains data that will be used for the [`"ydata"` property], page 502. Data is transferred into `"ydata"` using the [refreshdata function], page 553.

`zdata`: [] | vector, def. `[](0x0)`
: For 3-D data, vector with the y coordinates of the scatter object.

`zdatasource`: string, def. `""`
: The name of a workspace variable that contains data that will be used for the ["zdata" property], page 502. Data is transferred into `"zdata"` using the [refreshdata function], page 553.

## Creation/Deletion

`beingdeleted`: {`"off"`} | `"on"`
: Property indicating that a function has initiated deletion of the object. `beingdeleted` is set to true until the object no longer exists.

`createfcn`: string | function handle, def. `[](0x0)`
: Callback function executed immediately after scatter has been created. Function is set by using default property on root object, e.g., `set (groot, "defaultscattercreatefcn", 'disp ("scatter created!")')`.

  For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`deletefcn`: string | function handle, def. `[](0x0)`
: Callback function executed immediately before scatter is deleted.

  For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Display

`clipping`: `"off"` | {`"on"`}
: If `clipping` is `"on"`, the scatter is clipped in its parent axes limits.

`visible`: `"off"` | {`"on"`}
: If `visible` is `"off"`, the scatter is not rendered on screen.

## Legend Options

`annotation`: def. `[](0x0)`
: Legend appearance toggling from within the scatter object is not yet implemented for scatter objects. `annotation` is unused.

`displayname`: def. `""`
: Text of the legend entry corresponding to this scatter object.

## Marker Appearance

`linewidth`: scalar, def. `0.5000`
: Line width of the edge of the markers.

`marker`: `"*"` | `"+"` | `"."` | `"<"` | `">"` | `"^"` | `"_"` | `"d"` | `"diamond"` | `"h"` | `"hexagram"` | `"none"` | {`"o"`} | `"p"` | `"pentagram"` | `"s"` | `"square"` | `"v"` | `"x"` | `"|"`
: See [line marker property], page 487.

`markeredgealpha`: scalar, def. `1`
: Transparency level of the faces of the markers where a value of 0 means complete transparency and a value of 1 means solid faces without transparency. Note that

the markers are not sorted from back to front which might lead to unexpected results when rendering layered transparent markers or in combination with other transparent objects.

`markeredgecolor`: {`"none"`} | {`"flat"`} | colorspec
: Color of the edge of the markers. `"none"` means that the edges are transparent and `"flat"` means that the value from `cdata` is used. See [line markeredgecolor property], page 487.

`markerfacealpha`: scalar, def. `1`
: Transparency level of the faces of the markers where a value of 0 means complete transparency and a value of 1 means solid faces without transparency. Note that the markers are not sorted from back to front which might lead to unexpected results when rendering layered transparent markers or in combination with other transparent objects.

`markerfacecolor`: {{`"none"`}} | `"flat"` | `"auto"` | colorspec
: Color of the face of the markers. `"none"` means that the faces are transparent, `"flat"` means that the value from `cdata` is used, and `"auto"` uses the `color` property of the ancestor axes. See [line markerfacecolor property], page 487.

`sizedata`: [] | scalar | vector, def. `[](0x0)`
: Size of the area of the marker. A scalar value applies to all markers. If `cdata` is an N-by-1 vector, it defines the color of each one of the N scatter markers.

`sizedatasource`: def. `""`
: Data from workspace variables is not yet implemented for scatter objects. `sizedatasource` is unused.

## Mouse Interaction

`buttondownfcn`: string | function handle, def. `[](0x0)`
: For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`contextmenu`: graphics handle, def. `[](0x0)`
: Graphics handle of the uicontextmenu object that is currently associated to this scatter object.

`datatiptemplate`: def. `[](0x0)`
: Data tip objects is not yet implemented for scatter objects. `datatiptemplate` is unused.

`hittest`: `"off"` | {`"on"`}
: Specify whether scatter processes mouse events or passes them to ancestors of the object. When enabled, the object may respond to mouse clicks by evaluating the `"buttondownfcn"`, showing the uicontextmenu, and eventually becoming the root `"currentobject"`. This property is only relevant when the object can accept mouse clicks which is determined by the `"pickableparts"` property. See [pickableparts property], page 504.

`pickableparts`: `"all"` | `"none"` | {`"visible"`}
: Specify whether scatter will accept mouse clicks. By default, `pickableparts` is `"visible"` and only visible parts of the scatter or its children may react to

mouse clicks. When `pickableparts` is `"all"` both visible and invisible parts (or children) may react to mouse clicks. When `pickableparts` is `"none"` mouse clicks on the object are ignored and transmitted to any objects underneath this one. When an object is configured to accept mouse clicks the `"hittest"` property will determine how they are processed. See [hittest property], page 504.

`selected`: {`"off"`} | `"on"`
: Property indicates whether this scatter is selected.

`selectionhighlight`: `"off"` | {`"on"`}
: If `selectionhighlight` is `"on"`, then the scatter's selection state is visually highlighted.

## Object Identification

`tag`: string, def. `""`
: A user-defined string to label the graphics object.

`type` (read-only): string
: Class name of the graphics object. `type` is always `"scatter"`.

`userdata`: Any Octave data, def. `[](0x0)`
: User-defined data to associate with the graphics object.

## Parent/Children

`children` (read-only): vector of graphics handles, def. `[](0x0)`
: Child objects for Scatter plots is not yet implemented for scatter objects. `children` is unused.

`handlevisibility`: `"callback"` | `"off"` | {`"on"`}
: If `handlevisibility` is `"off"`, the scatter's handle is not visible in its parent's `"children"` property.

`parent`: graphics handle
: Handle of the parent graphics object.

### 15.3.3.10 Surface Properties

Properties of `surface` objects (see [surface], page 452):

## Categories:

[Callback Execution], page 505 | [Color and Transparency], page 506 | [Coordinate Data], page 506 | [Creation/Deletion], page 507 | [Display], page 507 | [Faces Appearance], page 507 | [Legend Options], page 508 | [Lighting], page 508 | [Marker Appearance], page 509 | [Mouse Interaction], page 509 | [Object Identification], page 510 | [Outline Appearance], page 510 | [Parent/Children], page 510

## Callback Execution

`busyaction`: `"cancel"` | {`"queue"`}
: Define how Octave handles the execution of this object's callback properties when it is unable to interrupt another object's executing callback. This is only

relevant when the currently executing callback object has its `interruptible` property set to `"off"`. The `busyaction` property of the interrupting callback object indicates whether the interrupting callback is queued (`"queue"` (default)) or discarded (`"cancel"`). See Section 15.4.4 [Callbacks section], page 546.

`interruptible`: `"off"` | {`"on"`}
: Specify whether this object's callback functions may be interrupted by other callbacks. By default `interruptible` is `"on"` and callbacks that make use of `drawnow`, `figure`, `waitfor`, `getframe` or `pause` functions are eventually interrupted. See Section 15.4.4 [Callbacks section], page 546.

## Color and Transparency

`alphadata`: scalar | matrix, def. `1`
: Transparency is not yet implemented for surface objects. `alphadata` is unused.

`alphadatamapping`: `"direct"` | `"none"` | {`"scaled"`}
: Transparency is not yet implemented for surface objects. `alphadatamapping` is unused.

`cdata`: array, def. 3-by-3 double
: Color data values for surface vertices. Data is stored either as a 2-D matrix the same size as ["zdata"], page 507, where each element's value determines that vertex's color according to the current colormap, or as a 3-D array where the third dimension contains separate red, blue, and green components for each vertex.

`cdatamapping`: `"direct"` | {`"scaled"`}
: Sets the method for mapping data from the ["cdata" property], page 506, to the current colormap. `"Direct"` mapping selects the color using the `"cdata"` value as an index to the current colormap. `"Scale"` mapping scales the `"cdata"` values to the range specified in the ["clim" axes property], page 478.

`cdatasource`: string, def. `""`
: The name of a workspace variable that contains data that will be used for the ["cdata" property], page 506. Data is transferred into `"cdata"` using the [refreshdata function], page 553.

## Coordinate Data

`xdata`: matrix, def. `[1 2 3]`
: Data for the x-coordinate.

`xdatasource`: string, def. `""`
: The name of a workspace variable that contains data that will be used for the ["xdata" property], page 506. Data is transferred into `"xdata"` using the [refreshdata function], page 553.

`ydata`: matrix, def. `[1; 2; 3]`
: Data for the y-coordinate.

`ydatasource`: string, def. `""`
The name of a workspace variable that contains data that will be used for the ["ydata" property], page 506. Data is transferred into `"ydata"` using the [refreshdata function], page 553.

`zdata`: matrix, def. 3-by-3 double
Data for the z-coordinate.

`zdatasource`: string, def. `""`
The name of a workspace variable that contains data that will be used for the ["zdata" property], page 507. Data is transferred into `"zdata"` using the [refreshdata function], page 553.

## Creation/Deletion

`beingdeleted`: {`"off"`} | `"on"`
Property indicating that a function has initiated deletion of the object. `beingdeleted` is set to true until the object no longer exists.

`createfcn`: string | function handle, def. `[](0x0)`
Callback function executed immediately after surface has been created. Function is set by using default property on root object, e.g., `set (groot, "defaultsurfacecreatefcn", 'disp ("surface created!")')`.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`deletefcn`: string | function handle, def. `[](0x0)`
Callback function executed immediately before surface is deleted.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Display

`clipping`: `"off"` | {`"on"`}
If `clipping` is `"on"`, the surface is clipped in its parent axes limits.

`visible`: `"off"` | {`"on"`}
If `visible` is `"off"`, the surface is not rendered on screen.

## Faces Appearance

`facealpha`: scalar | `"flat"` | `"interp"` | `"texturemap"`, def. `1`
Transparency level of the faces of the surface object. Only double values are supported at present where a value of 0 means complete transparency and a value of 1 means solid faces without transparency. Setting the property to `"flat"`, `"interp"` or `"texturemap"` causes the faces to not being rendered. Additionally, the faces are not sorted from back to front which might lead to unexpected results when rendering layered transparent faces.

`facecolor`: colorspec | `"none"` | {`"flat"`} | `"interp"`
Color of the faces of the surface object, specified as either a valid color specification or one of `"none"`, `"flat"`, or `"interp"`. `"flat"` and `"interp"` will set

either a single color for each face or a color interpolated across the face's vertices using the color value data stored in ["cdata"], page 506. See Section 15.4.1 [colorspec], page 545.

`facelighting`: {`"flat"`} | `"gouraud"` | `"none"` | `"phong"`
When set to a value other than `"none"`, the faces of the object are drawn with light and shadow effects. Supported values are `"none"` (no lighting effects), `"flat"` (faceted look), and `"gouraud"` (linear interpolation of the lighting effects between the vertices). `"phong"` is deprecated and has the same effect as `"gouraud"`.

`facenormals`: def. `[](0x0)`
Face normals are used for lighting the edges or faces if the `edgelighting` or `facelighting` properties are set to `"flat"`. Setting `facenormals` also forces the `facenormalsmode` property to be set to `"manual"`.

`facenormalsmode`: {`"auto"`} | `"manual"`
If this property is set to `"auto"`, `facenormals` are automatically calculated if the `edgelighting` or `facelighting` property are set to `"flat"` and at least one `light` object is present and visible in the same axes.

## Legend Options

`displayname`: def. `""`
Text for the legend entry corresponding to this surface.

## Lighting

`ambientstrength`: scalar, def. `0.3000`
Strength of the ambient light. Value between 0.0 and 1.0.

`backfacelighting`: `"lit"` | {`"reverselit"`} | `"unlit"`
`"lit"`: The normals are used as is for lighting. `"reverselit"`: The normals are always oriented towards the point of view. `"unlit"`: Faces with normals pointing away from the point of view are unlit.

`diffusestrength`: scalar, def. `0.6000`
Strength of the diffuse reflection. Value between 0.0 (no diffuse reflection) and 1.0 (full diffuse reflection).

`specularcolorreflectance`: scalar, def. `1`
Reflectance for specular color. Value between 0.0 (color of underlying face) and 1.0 (color of light source).

`specularexponent`: scalar, def. `10`
Exponent for the specular reflection. The lower the value, the more the reflection is spread out.

`specularstrength`: scalar, def. `0.9000`
Strength of the specular reflection. Value between 0.0 (no specular reflection) and 1.0 (full specular reflection).

`vertexnormals`: def. `[](0x0)`
: Vertex normals are used for lighting the edges or faces if the `edgelighting` or `facelighting` properties are set to `"gouraud"`. Setting `vertexnormals` also forces the `vertexnormalsmode` property to be set to `"manual"`.

`vertexnormalsmode`: {`"auto"`} | `"manual"`
: If this property is set to `"auto"`, `vertexnormals` are automatically calculated if the `edgelighting` or `facelighting` property are set to `"gouraud"` and at least one `light` object is present and visible in the same axes.

## Marker Appearance

`marker`: `"*"` | `"+"` | `"."` | `"<"` | `">"` | `"^"` | `"_"` | `"d"` | `"diamond"` | `"h"` | `"hexagram"` | {`"none"`} | `"o"` | `"p"` | `"pentagram"` | `"s"` | `"square"` | `"v"` | `"x"` | `"|"`
: See Section 15.4.3 [Marker Styles], page 546.

`markeredgecolor`: {`"auto"`} | `"flat"` | `"none"`
: See [line markeredgecolor property], page 487.

`markerfacecolor`: `"auto"` | `"flat"` | {`"none"`}
: See [line markerfacecolor property], page 487.

`markersize`: scalar, def. `6`
: See [line markersize property], page 487.

## Mouse Interaction

`buttondownfcn`: string | function handle, def. `[](0x0)`
: For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`contextmenu`: graphics handle, def. `[](0x0)`
: Graphics handle of the uicontextmenu object that is currently associated to this surface object.

`hittest`: `"off"` | {`"on"`}
: Specify whether surface processes mouse events or passes them to ancestors of the object. When enabled, the object may respond to mouse clicks by evaluating the `"buttondownfcn"`, showing the uicontextmenu, and eventually becoming the root `"currentobject"`. This property is only relevant when the object can accept mouse clicks which is determined by the `"pickableparts"` property. See [pickableparts property], page 509.

`pickableparts`: `"all"` | `"none"` | {`"visible"`}
: Specify whether surface will accept mouse clicks. By default, `pickableparts` is `"visible"` and only visible parts of the surface or its children may react to mouse clicks. When `pickableparts` is `"all"` both visible and invisible parts (or children) may react to mouse clicks. When `pickableparts` is `"none"` mouse clicks on the object are ignored and transmitted to any objects underneath this one. When an object is configured to accept mouse clicks the `"hittest"` property will determine how they are processed. See [hittest property], page 509.

`selected`: {`"off"`} | `"on"`
: Property indicates whether this surface is selected.

`selectionhighlight`: `"off"` | {`"on"`}
: If `selectionhighlight` is `"on"`, then the surface's selection state is visually highlighted.

## Object Identification

`tag`: string, def. `""`
: A user-defined string to label the graphics object.

`type` (read-only): string
: Class name of the graphics object. `type` is always `"surface"`.

`userdata`: Any Octave data, def. `[](0x0)`
: User-defined data to associate with the graphics object.

## Outline Appearance

`edgealpha`: scalar, def. `1`
: Transparency is not yet implemented for surface objects. `edgealpha` is unused.

`edgecolor`: colorspec | `"none"` | `"flat"` | `"interp"`, def. `[0 0 0]`
: Color of the edges of the surface object, specified as either a valid color specification or one of `"none"`, `"flat"`, or `"interp"`. `"flat"` and `"interp"` will set either a single color for each edge or a color interpolated between two adjacent vertices using the color value data stored in [`"cdata"`], page 506. See Section 15.4.1 [colorspec], page 545.

`edgelighting`: `"flat"` | `"gouraud"` | {`"none"`} | `"phong"`
: When set to a value other than `"none"`, the edges of the object are drawn with light and shadow effects. Supported values are `"none"` (no lighting effects), `"flat"` (faceted look), and `"gouraud"` (linear interpolation of the lighting effects between the vertices). `"phong"` is deprecated and has the same effect as `"gouraud"`.

`linestyle`: {`"-"`} | `"--"` | `"-."` | `":"` | `"none"`
: See Section 15.4.2 [Line Styles], page 545.

`linewidth`: scalar, def. `0.5000`
: See [line linewidth property], page 487.

`meshstyle`: {`"both"`} | `"column"` | `"row"`
: Specifies whether to display the edges associated with the surface data's rows, columns, or both.

## Parent/Children

`children` (read-only): vector of graphics handles, def. `[](0x0)`
: Child objects for Surfaces is not yet implemented for surface objects. `children` is unused.

`handlevisibility`: `"callback"` | `"off"` | {`"on"`}
: If `handlevisibility` is `"off"`, the surface's handle is not visible in its parent's `"children"` property.

`parent`: graphics handle
: Handle of the parent graphics object.

### 15.3.3.11 Light Properties

Properties of `light` objects (see [light], page 453):

## Categories:

[Callback Execution], page 511 | [Creation/Deletion], page 511 | [Display], page 511 | [Lighting], page 512 | [Mouse Interaction], page 512 | [Object Identification], page 512 | [Parent/Children], page 513

## Callback Execution

`busyaction`: `"cancel"` | {`"queue"`}
: Define how Octave handles the execution of this object's callback properties when it is unable to interrupt another object's executing callback. This is only relevant when the currently executing callback object has its `interruptible` property set to `"off"`. The `busyaction` property of the interrupting callback object indicates whether the interrupting callback is queued (`"queue"` (default)) or discarded (`"cancel"`). See Section 15.4.4 [Callbacks section], page 546.

`interruptible`: `"off"` | {`"on"`}
: Specify whether this object's callback functions may be interrupted by other callbacks. By default `interruptible` is `"on"` and callbacks that make use of `drawnow`, `figure`, `waitfor`, `getframe` or `pause` functions are eventually interrupted. See Section 15.4.4 [Callbacks section], page 546.

## Creation/Deletion

`beingdeleted`: {`"off"`} | `"on"`
: Property indicating that a function has initiated deletion of the object. `beingdeleted` is set to true until the object no longer exists.

`createfcn`: string | function handle, def. `[](0x0)`
: Callback function executed immediately after light has been created. Function is set by using default property on root object, e.g., `set (groot, "defaultlightcreatefcn", 'disp ("light created!")')`.

  For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`deletefcn`: string | function handle, def. `[](0x0)`
: Callback function executed immediately before light is deleted.

  For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Display

`clipping`: `"off"` | {`"on"`}
: If `clipping` is `"on"`, the light is clipped in its parent axes limits.

`visible`: `"off"` | {`"on"`}
: If `visible` is `"off"`, the light is not rendered on screen.

## Lighting

`color`: colorspec, def. `[1 1 1]`
: Color of the light source. See Section 15.4.1 [colorspec], page 545.

`position`: def. `[1 0 1]`
: Position of the light source.

`style`: {`"infinite"`} | `"local"`
: This string defines whether the light emanates from a light source at infinite distance (`"infinite"`) or from a local point source (`"local"`).

## Mouse Interaction

`buttondownfcn`: string | function handle, def. `[](0x0)`
: For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`contextmenu`: graphics handle, def. `[](0x0)`
: Graphics handle of the uicontextmenu object that is currently associated to this light object.

`hittest`: `"off"` | {`"on"`}
: Specify whether light processes mouse events or passes them to ancestors of the object. When enabled, the object may respond to mouse clicks by evaluating the `"buttondownfcn"`, showing the uicontextmenu, and eventually becoming the root `"currentobject"`. This property is only relevant when the object can accept mouse clicks which is determined by the `"pickableparts"` property. See [pickableparts property], page 512.

`pickableparts`: `"all"` | `"none"` | {`"visible"`}
: Specify whether light will accept mouse clicks. By default, `pickableparts` is `"visible"` and only visible parts of the light or its children may react to mouse clicks. When `pickableparts` is `"all"` both visible and invisible parts (or children) may react to mouse clicks. When `pickableparts` is `"none"` mouse clicks on the object are ignored and transmitted to any objects underneath this one. When an object is configured to accept mouse clicks the `"hittest"` property will determine how they are processed. See [hittest property], page 512.

`selected`: {`"off"`} | `"on"`
: Property indicates whether this light is selected.

`selectionhighlight`: `"off"` | {`"on"`}
: If `selectionhighlight` is `"on"`, then the light's selection state is visually highlighted.

## Object Identification

`tag`: string, def. `""`
A user-defined string to label the graphics object.

`type` (read-only): string
Class name of the graphics object. `type` is always `"light"`.

`userdata`: Any Octave data, def. `[](0x0)`
User-defined data to associate with the graphics object.

## Parent/Children

`children` (read-only): vector of graphics handles, def. `[](0x0)`
light objects have no child objects. `children` is unused.

`handlevisibility`: `"callback"` | `"off"` | {`"on"`}
If `handlevisibility` is `"off"`, the light's handle is not visible in its parent's `"children"` property.

`parent`: graphics handle
Handle of the parent graphics object.

### 15.3.3.12 Uimenu Properties

Properties of `uimenu` objects (see [uimenu], page 1014):

## Categories:

[Appearance], page 513 | [Callback Execution], page 513 | [Creation/Deletion], page 514 | [Display], page 514 | [Keyboard Interaction], page 514 | [Menu Options], page 514 | [Mouse Interaction], page 514 | [Object Identification], page 515 | [Object Position], page 515 | [Parent/Children], page 515

## Appearance

`foregroundcolor`: colorspec, def. `[0 0 0]`
The color value of the text for this menu entry.

`separator`: {`"off"`} | `"on"`
State indicating whether a separator line will be drawn above the current menu position.

## Callback Execution

`busyaction`: `"cancel"` | {`"queue"`}
Define how Octave handles the execution of this object's callback properties when it is unable to interrupt another object's executing callback. This is only relevant when the currently executing callback object has its `interruptible` property set to `"off"`. The `busyaction` property of the interrupting callback object indicates whether the interrupting callback is queued (`"queue"` (default)) or discarded (`"cancel"`). See Section 15.4.4 [Callbacks section], page 546.

`interruptible`: `"off"` | {`"on"`}
: Specify whether this object's callback functions may be interrupted by other callbacks. By default `interruptible` is `"on"` and callbacks that make use of `drawnow`, `figure`, `waitfor`, `getframe` or `pause` functions are eventually interrupted. See Section 15.4.4 [Callbacks section], page 546.

`menuselectedfcn`: string | function handle, def. `[](0x0)`
: Function that is called when this menu item is executed. For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Creation/Deletion

`beingdeleted`: {`"off"`} | `"on"`
: Property indicating that a function has initiated deletion of the object. `beingdeleted` is set to true until the object no longer exists.

`createfcn`: string | function handle, def. `[](0x0)`
: Callback function executed immediately after uimenu has been created. Function is set by using default property on root object, e.g., `set (groot, "defaultuimenucreatefcn", 'disp ("uimenu created!")')`.

  For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`deletefcn`: string | function handle, def. `[](0x0)`
: Callback function executed immediately before uimenu is deleted.

  For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Display

`clipping`: `"off"` | {`"on"`}
: If `clipping` is `"on"`, the uimenu is clipped in its parent axes limits.

`visible`: `"off"` | {`"on"`}
: If `visible` is `"off"`, the uimenu is not rendered on screen.

## Keyboard Interaction

`accelerator`: character, def. `""`
: A character that when pressed together with CTRL will execute this menu entry (e.g., `"x"` for `CTRL+x`).

## Menu Options

`checked`: {`"off"`} | `"on"`
: Sets whether or not a mark appears at this menu entry.

`enable`: `"off"` | {`"on"`}
: Sets whether this menu entry is active or is grayed out.

`text`: string, def. `""`
: The text for this menu entry. A `"&"` character can be used to mark the [`"accelerator"` key], page 514,

## Mouse Interaction

`buttondownfcn`: string | function handle, def. `[](0x0)`
: `buttondownfcn` is unused.

`contextmenu`: graphics handle, def. `[](0x0)`
: Graphics handle of the uicontextmenu object that is currently associated to this uimenu object.

`hittest`: `"off"` | {`"on"`}
: Specify whether uimenu processes mouse events or passes them to ancestors of the object. When enabled, the object may respond to mouse clicks by evaluating the `"buttondownfcn"`, showing the uicontextmenu, and eventually becoming the root `"currentobject"`. This property is only relevant when the object can accept mouse clicks which is determined by the `"pickableparts"` property. See [pickableparts property], page 515.

`pickableparts`: `"all"` | `"none"` | {`"visible"`}
: Specify whether uimenu will accept mouse clicks. By default, `pickableparts` is `"visible"` and only visible parts of the uimenu or its children may react to mouse clicks. When `pickableparts` is `"all"` both visible and invisible parts (or children) may react to mouse clicks. When `pickableparts` is `"none"` mouse clicks on the object are ignored and transmitted to any objects underneath this one. When an object is configured to accept mouse clicks the `"hittest"` property will determine how they are processed. See [hittest property], page 515.

`selected`: {`"off"`} | `"on"`
: Property indicates whether this uimenu is selected.

`selectionhighlight`: `"off"` | {`"on"`}
: If `selectionhighlight` is `"on"`, then the uimenu's selection state is visually highlighted.

## Object Identification

`tag`: string, def. `""`
: A user-defined string to label the graphics object.

`type` (read-only): string
: Class name of the graphics object. `type` is always `"uimenu"`.

`userdata`: Any Octave data, def. `[](0x0)`
: User-defined data to associate with the graphics object.

## Object Position

`position`: scalar, def. `4`
: A scalar value containing the relative menu position from the left or top depending on the orientation of the menu.

## Parent/Children

`children` (read-only): vector of graphics handles, def. `[](0x0)`
: Graphics handles of the uimenu's children.

`handlevisibility`: `"callback"` | `"off"` | {`"on"`}
: If `handlevisibility` is `"off"`, the uimenu's handle is not visible in its parent's `"children"` property.

`parent`: graphics handle
: Handle of the parent graphics object.

### 15.3.3.13 Uibuttongroup Properties

Properties of `uibuttongroup` objects (see [uibuttongroup], page 1009):

## Categories:

[Appearance], page 516 | [Button Group Operation], page 516 | [Callback Execution], page 516 | [Creation/Deletion], page 517 | [Display], page 517 | [Mouse Interaction], page 517 | [Object Identification], page 518 | [Object Position], page 518 | [Parent/Children], page 518 | [Text Appearance], page 519

## Appearance

`backgroundcolor`: colorspec, def. `[0.9400 0.9400 0.9400]`
: The color value of the background of this buttongroup.

`bordertype`: `"beveledin"` | `"beveledout"` | {`"etchedin"`} | `"etchedout"` | `"line"` | `"none"`
: Sets whether or not a line border will surround the buttongroup.

`borderwidth`: whole number scalar, def. `1`
: The width of the line border in pixels.

`foregroundcolor`: colorspec, def. `[0 0 0]`
: The color value of the title text for this buttongroup.

`highlightcolor`: colorspec, def. `[1 1 1]`
: The color value of the line bordering this buttongroup.

`shadowcolor`: colorspec, def. `[0.7000 0.7000 0.7000]`
: The color value of the line surrounding the border line around this buttongroup.

`sizechangedfcn`: string | function handle, def. `[](0x0)`
: Callback triggered when the buttongroup size is changed.

  For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Button Group Operation

`selectedobject`: def. `[](0x0)`
: Graphic handle of the currently selected item in the buttongroup.

## Callback Execution

`busyaction`: `"cancel"` | {`"queue"`}
: Define how Octave handles the execution of this object's callback properties when it is unable to interrupt another object's executing callback. This is only

relevant when the currently executing callback object has its `interruptible` property set to `"off"`. The `busyaction` property of the interrupting callback object indicates whether the interrupting callback is queued (`"queue"` (default)) or discarded (`"cancel"`). See Section 15.4.4 [Callbacks section], page 546.

`interruptible`: `"off"` | {`"on"`}
Specify whether this object's callback functions may be interrupted by other callbacks. By default `interruptible` is `"on"` and callbacks that make use of `drawnow`, `figure`, `waitfor`, `getframe` or `pause` functions are eventually interrupted. See Section 15.4.4 [Callbacks section], page 546.

`resizefcn`: string | function handle, def. `[](0x0)`
`resizefcn` is deprecated. Use `sizechangedfcn` instead.

`selectionchangedfcn`: string | function handle, def. `[](0x0)`
Callback triggered when the selected item within the buttongroup is changed.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Creation/Deletion

`beingdeleted`: {`"off"`} | `"on"`
Property indicating that a function has initiated deletion of the object. `beingdeleted` is set to true until the object no longer exists.

`createfcn`: string | function handle, def. `[](0x0)`
Callback function executed immediately after uibuttongroup has been created. Function is set by using default property on root object, e.g., `set (groot, "defaultuibuttongroupcreatefcn", 'disp ("uibuttongroup created!")')`.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`deletefcn`: string | function handle, def. `[](0x0)`
Callback function executed immediately before uibuttongroup is deleted.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Display

`clipping`: `"off"` | {`"on"`}
If `clipping` is `"on"`, the uibuttongroup is clipped in its parent axes limits.

`visible`: `"off"` | {`"on"`}
If `visible` is `"off"`, the uibuttongroup is not rendered on screen.

## Mouse Interaction

`buttondownfcn`: string | function handle, def. `[](0x0)`
For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`contextmenu`: graphics handle, def. `[](0x0)`
: Graphics handle of the uicontextmenu object that is currently associated to this uibuttongroup object.

`hittest`: `"off"` | {`"on"`}
: Specify whether uibuttongroup processes mouse events or passes them to ancestors of the object. When enabled, the object may respond to mouse clicks by evaluating the `"buttondownfcn"`, showing the uicontextmenu, and eventually becoming the root `"currentobject"`. This property is only relevant when the object can accept mouse clicks which is determined by the `"pickableparts"` property. See [pickableparts property], page 518.

`pickableparts`: `"all"` | `"none"` | {`"visible"`}
: Specify whether uibuttongroup will accept mouse clicks. By default, `pickableparts` is `"visible"` and only visible parts of the uibuttongroup or its children may react to mouse clicks. When `pickableparts` is `"all"` both visible and invisible parts (or children) may react to mouse clicks. When `pickableparts` is `"none"` mouse clicks on the object are ignored and transmitted to any objects underneath this one. When an object is configured to accept mouse clicks the `"hittest"` property will determine how they are processed. See [hittest property], page 518.

`selected`: {`"off"`} | `"on"`
: Property indicates whether this uibuttongroup is selected.

`selectionhighlight`: `"off"` | {`"on"`}
: If `selectionhighlight` is `"on"`, then the uibuttongroup's selection state is visually highlighted.

## Object Identification

`tag`: string, def. `""`
: A user-defined string to label the graphics object.

`type` (read-only): string
: Class name of the graphics object. `type` is always `"uibuttongroup"`.

`userdata`: Any Octave data, def. `[](0x0)`
: User-defined data to associate with the graphics object.

## Object Position

`position`: four-element vector, def. `[0 0 1 1]`
: Size of the buttongroup represented as the four-element vector [left, bottom, width, height].

`units`: `"centimeters"` | `"characters"` | `"inches"` | {`"normalized"`} | `"pixels"` | `"points"`
: Unit of measurement used to interpret the `"position"` property.

## Parent/Children

`children` (read-only): vector of graphics handles, def. `[](0x0)`
: Graphics handles of the uibuttongroup's children.

`handlevisibility`: `"callback"` | `"off"` | {`"on"`}
: If `handlevisibility` is `"off"`, the uibuttongroup's handle is not visible in its parent's `"children"` property.

`parent`: graphics handle
: Handle of the parent graphics object.

## Text Appearance

`fontangle`: `"italic"` | {`"normal"`}
: Control whether the font is italic or normal.

`fontname`: string, def. `"*"`
: Name of font used for text rendering. When setting this property, the text rendering engine will search for a matching font in your system. If none is found then text is rendered using a default sans serif font (same as the default `"*"` value).

  Programming Note: On systems that don't use FontConfig natively (all but Linux), the font cache is built when Octave is installed. You will need to run `system ("fc-cache -fv")` manually after installing new fonts.

`fontsize`: scalar, def. `10`
: Size of the font used for text rendering. See [fontunits property], page 519.

`fontunits`: `"centimeters"` | `"inches"` | `"normalized"` | `"pixels"` | {`"points"`}
: Units used to interpret the `"fontsize"` property.

`fontweight`: `"bold"` | {`"normal"`}
: Control the variant of the base font used for text rendering.

`title`: string, def. `""`
: The text for the buttongroup title.

`titleposition`: `"centerbottom"` | `"centertop"` | `"leftbottom"` | {`"lefttop"`} | `"rightbottom"` | `"righttop"`
: Relative position of the title within the buttongroup.

### 15.3.3.14 Uicontextmenu Properties

Properties of `uicontextmenu` objects (see [uicontextmenu], page 1015):

## Categories:

[Callback Execution], page 519 | [Creation/Deletion], page 520 | [Display], page 520 | [Mouse Interaction], page 520 | [Object Identification], page 521 | [Object Position], page 521 | [Parent/Children], page 521

## Callback Execution

`busyaction`: `"cancel"` | {`"queue"`}
: Define how Octave handles the execution of this object's callback properties when it is unable to interrupt another object's executing callback. This is only relevant when the currently executing callback object has its `interruptible` property set to `"off"`. The `busyaction` property of the interrupting callback

object indicates whether the interrupting callback is queued (`"queue"` (default)) or discarded (`"cancel"`). See Section 15.4.4 [Callbacks section], page 546.

`callback`: string, def. `[](0x0)`
A string consisting of a valid Octave expression that will be executed whenever this item is selected.

`interruptible`: `"off"` | {`"on"`}
Specify whether this object's callback functions may be interrupted by other callbacks. By default `interruptible` is `"on"` and callbacks that make use of `drawnow`, `figure`, `waitfor`, `getframe` or `pause` functions are eventually interrupted. See Section 15.4.4 [Callbacks section], page 546.

## Creation/Deletion

`beingdeleted`: {`"off"`} | `"on"`
Property indicating that a function has initiated deletion of the object. `beingdeleted` is set to true until the object no longer exists.

`createfcn`: string | function handle, def. `[](0x0)`
Callback function executed immediately after uicontextmenu has been created. Function is set by using default property on root object, e.g., `set (groot, "defaultuicontextmenucreatefcn", 'disp ("uicontextmenu created!")')`.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`deletefcn`: string | function handle, def. `[](0x0)`
Callback function executed immediately before uicontextmenu is deleted.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Display

`clipping`: `"off"` | {`"on"`}
If `clipping` is `"on"`, the uicontextmenu is clipped in its parent axes limits.

`visible`: `"off"` | {`"on"`}
If `visible` is `"off"`, the uicontextmenu is not rendered on screen.

## Mouse Interaction

`buttondownfcn`: string | function handle, def. `[](0x0)`
`buttondownfcn` is unused.

`contextmenu`: graphics handle, def. `[](0x0)`
Graphics handle of the uicontextmenu object that is currently associated to this uicontextmenu object.

`hittest`: `"off"` | {`"on"`}
Specify whether uicontextmenu processes mouse events or passes them to ancestors of the object. When enabled, the object may respond to mouse clicks by

evaluating the `"buttondownfcn"`, showing the uicontextmenu, and eventually becoming the root `"currentobject"`. This property is only relevant when the object can accept mouse clicks which is determined by the `"pickableparts"` property. See [pickableparts property], page 521.

`pickableparts`: `"all"` | `"none"` | {`"visible"`}
: Specify whether uicontextmenu will accept mouse clicks. By default, `pickableparts` is `"visible"` and only visible parts of the uicontextmenu or its children may react to mouse clicks. When `pickableparts` is `"all"` both visible and invisible parts (or children) may react to mouse clicks. When `pickableparts` is `"none"` mouse clicks on the object are ignored and transmitted to any objects underneath this one. When an object is configured to accept mouse clicks the `"hittest"` property will determine how they are processed. See [hittest property], page 520.

`selected`: {`"off"`} | `"on"`
: Property indicates whether this uicontextmenu is selected.

`selectionhighlight`: `"off"` | {`"on"`}
: If `selectionhighlight` is `"on"`, then the uicontextmenu's selection state is visually highlighted.

## Object Identification

`tag`: string, def. `""`
: A user-defined string to label the graphics object.

`type` (read-only): string
: Class name of the graphics object. `type` is always `"uicontextmenu"`.

`userdata`: Any Octave data, def. `[](0x0)`
: User-defined data to associate with the graphics object.

## Object Position

`position`: def. `[0 0]`
: Manually setting location for uicontextmenu to appear is not yet implemented for uicontextmenu objects. `position` is unused.

## Parent/Children

`children` (read-only): vector of graphics handles, def. `[](0x0)`
: Graphics handles of the uicontextmenu's children.

`handlevisibility`: `"callback"` | `"off"` | {`"on"`}
: If `handlevisibility` is `"off"`, the uicontextmenu's handle is not visible in its parent's `"children"` property.

`parent`: graphics handle
: Handle of the parent graphics object.

### 15.3.3.15 Uipanel Properties

Properties of `uipanel` objects (see [uipanel], page 1008):

### Categories:

[Appearance], page 522 | [Callback Execution], page 522 | [Creation/Deletion], page 522 | [Display], page 523 | [Mouse Interaction], page 523 | [Object Identification], page 524 | [Object Position], page 524 | [Parent/Children], page 524 | [Text Appearance], page 524

### Appearance

`backgroundcolor`: colorspec, def. `[0.9400 0.9400 0.9400]`
The color value of the background of this panel.

`bordertype`: `"beveledin"` | `"beveledout"` | {`"etchedin"`} | `"etchedout"` | `"line"` | `"none"`
Sets whether or not a line border will surround the panel.

`borderwidth`: whole number scalar, def. `1`
The width of the line border in pixels.

`foregroundcolor`: colorspec, def. `[0 0 0]`
The color value of the title text for this panel.

`highlightcolor`: colorspec, def. `[1 1 1]`
The color value of the line bordering this panel.

`shadowcolor`: colorspec, def. `[0.7000 0.7000 0.7000]`
The color value of the line surrounding the border line around this panel. See Section 15.4.1 [colorspec], page 545.

### Callback Execution

`busyaction`: `"cancel"` | {`"queue"`}
Define how Octave handles the execution of this object's callback properties when it is unable to interrupt another object's executing callback. This is only relevant when the currently executing callback object has its `interruptible` property set to `"off"`. The `busyaction` property of the interrupting callback object indicates whether the interrupting callback is queued (`"queue"` (default)) or discarded (`"cancel"`). See Section 15.4.4 [Callbacks section], page 546.

`interruptible`: `"off"` | {`"on"`}
Specify whether this object's callback functions may be interrupted by other callbacks. By default `interruptible` is `"on"` and callbacks that make use of `drawnow`, `figure`, `waitfor`, `getframe` or `pause` functions are eventually interrupted. See Section 15.4.4 [Callbacks section], page 546.

`resizefcn`: string | function handle, def. `[](0x0)`
`resizefcn` is deprecated. Use `sizechangedfcn` instead.

`sizechangedfcn`: string | function handle, def. `[](0x0)`
Callback triggered when the panel size is changed.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Creation/Deletion

`beingdeleted`: {`"off"`} | `"on"`
: Property indicating that a function has initiated deletion of the object. `beingdeleted` is set to true until the object no longer exists.

`createfcn`: string | function handle, def. `[](0x0)`
: Callback function executed immediately after uipanel has been created. Function is set by using default property on root object, e.g., `set (groot, "defaultuipanelcreatefcn", 'disp ("uipanel created!")')`.

  For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`deletefcn`: string | function handle, def. `[](0x0)`
: Callback function executed immediately before uipanel is deleted.

  For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Display

`clipping`: `"off"` | {`"on"`}
: If `clipping` is `"on"`, the uipanel is clipped in its parent axes limits.

`visible`: `"off"` | {`"on"`}
: If `visible` is `"off"`, the uipanel is not rendered on screen.

## Mouse Interaction

`buttondownfcn`: string | function handle, def. `[](0x0)`
: For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`contextmenu`: graphics handle, def. `[](0x0)`
: Graphics handle of the uicontextmenu object that is currently associated to this uipanel object.

`hittest`: `"off"` | {`"on"`}
: Specify whether uipanel processes mouse events or passes them to ancestors of the object. When enabled, the object may respond to mouse clicks by evaluating the `"buttondownfcn"`, showing the uicontextmenu, and eventually becoming the root `"currentobject"`. This property is only relevant when the object can accept mouse clicks which is determined by the `"pickableparts"` property. See [pickableparts property], page 523.

`pickableparts`: `"all"` | `"none"` | {`"visible"`}
: Specify whether uipanel will accept mouse clicks. By default, `pickableparts` is `"visible"` and only visible parts of the uipanel or its children may react to mouse clicks. When `pickableparts` is `"all"` both visible and invisible parts (or children) may react to mouse clicks. When `pickableparts` is `"none"` mouse clicks on the object are ignored and transmitted to any objects underneath this one. When an object is configured to accept mouse clicks the `"hittest"` property will determine how they are processed. See [hittest property], page 523.

`selected`: {`"off"`} | `"on"`
: Property indicates whether this uipanel is selected.

`selectionhighlight`: `"off"` | {`"on"`}
: If `selectionhighlight` is `"on"`, then the uipanel's selection state is visually highlighted.

## Object Identification

`tag`: string, def. `""`
: A user-defined string to label the graphics object.

`type` (read-only): string
: Class name of the graphics object. `type` is always `"uipanel"`.

`userdata`: Any Octave data, def. `[](0x0)`
: User-defined data to associate with the graphics object.

## Object Position

`position`: four-element vector, def. `[0 0 1 1]`
: Size of the panel represented as the four-element vector [left, bottom, width, height].

`units`: `"centimeters"` | `"characters"` | `"inches"` | {`"normalized"`} | `"pixels"` | `"points"`
: Unit of measurement used to interpret the `"position"` property.

## Parent/Children

`children` (read-only): vector of graphics handles, def. `[](0x0)`
: Graphics handles of the uipanel's children.

`handlevisibility`: `"callback"` | `"off"` | {`"on"`}
: If `handlevisibility` is `"off"`, the uipanel's handle is not visible in its parent's `"children"` property.

`parent`: graphics handle
: Handle of the parent graphics object.

## Text Appearance

`fontangle`: `"italic"` | {`"normal"`}
: Control whether the font is italic or normal.

`fontname`: string, def. `"*"`
: Name of font used for text rendering. When setting this property, the text rendering engine will search for a matching font in your system. If none is found then text is rendered using a default sans serif font (same as the default `"*"` value).

  Programming Note: On systems that don't use FontConfig natively (all but Linux), the font cache is built when Octave is installed. You will need to run `system ("fc-cache -fv")` manually after installing new fonts.

`fontsize`: scalar, def. `10`
: Size of the font used for text rendering. See [fontunits property], page 525.

`fontunits`: `"centimeters"` | `"inches"` | `"normalized"` | `"pixels"` | {`"points"`}
: Units used to interpret the `"fontsize"` property.

`fontweight`: `"bold"` | {`"normal"`}
: Control the variant of the base font used for text rendering.

`title`: string, def. `""`
: The text for the panel title.

`titleposition`: `"centerbottom"` | `"centertop"` | `"leftbottom"` | {`"lefttop"`} | `"rightbottom"` | `"righttop"`
: Relative position of the title within the panel.

### 15.3.3.16 Uicontrol Properties

Properties of `uicontrol` objects (see [uicontrol], page 1009):

## Categories:

[Appearance], page 525 | [Callback Execution], page 525 | [Control Options], page 526 | [Creation/Deletion], page 527 | [Display], page 527 | [Mouse Interaction], page 527 | [Object Identification], page 528 | [Object Position], page 528 | [Parent/Children], page 529 | [Text Appearance], page 529

## Appearance

`backgroundcolor`: colorspec, def. `[0.9400 0.9400 0.9400]`
: The color value of the background of this control object.

`cdata`: array, def. `[](0x0)`
: Image data used to represent the control object, stored as a M x N x 3 RGB array.

`extent` (read-only): four-element vector
: Size of the text string associated to the uicontrol returned in the form `[0 0 width height]` (the two first elements are always zero).

  For multi-line strings the returned `width` and `height` indicate the size of the rectangle enclosing all lines.

`foregroundcolor`: colorspec, def. `[0 0 0]`
: The color value of the text for this control object. See Section 15.4.1 [colorspec], page 545.

`style`: `"checkbox"` | `"edit"` | `"frame"` | `"listbox"` | `"popupmenu"` | {`"pushbutton"`} | `"radiobutton"` | `"slider"` | `"text"` | `"togglebutton"`
: The type of control object created. For a complete description of available control styles, see the [`"uicontrol"` function], page 1009,

## Callback Execution

`busyaction`: `"cancel"` | {`"queue"`}
: Define how Octave handles the execution of this object's callback properties when it is unable to interrupt another object's executing callback. This is only relevant when the currently executing callback object has its `interruptible` property set to `"off"`. The `busyaction` property of the interrupting callback object indicates whether the interrupting callback is queued (`"queue"` (default)) or discarded (`"cancel"`). See Section 15.4.4 [Callbacks section], page 546.

`callback`: string, def. `[](0x0)`
: A string consisting of a valid Octave expression that will be executed whenever this control is activated.

`interruptible`: `"off"` | {`"on"`}
: Specify whether this object's callback functions may be interrupted by other callbacks. By default `interruptible` is `"on"` and callbacks that make use of `drawnow`, `figure`, `waitfor`, `getframe` or `pause` functions are eventually interrupted. See Section 15.4.4 [Callbacks section], page 546.

`keypressfcn`: string | function handle, def. `[](0x0)`
: Callback function executed when a key is pressed while the uicontrol object has focus. The first argument to the function is the handle of the calling uicontrol. The second argument holds an event structure which has the following members:

  `Character:`
  : The ASCII value of the key

  `Modifier:`
  : A cell array containing strings representing the modifiers pressed with the key.

  `Key:` Lowercase description of the key

  `Source:` Graphics handle of the object executing the callback function

  `EventName:`
  : `"KeyPress"`

  For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Control Options

`enable`: `"inactive"` | `"off"` | {`"on"`}
: Sets whether this control object is active or is grayed out.

`listboxtop`: scalar, def. `1`
: The index of the string option that will appear at the top of `"listbox"` controls

`max`: scalar, def. `1`
: The maximum control value, whose effect on the control is dependent on the control type. For `"checkbox"`, `"togglebutton"`, and `"radiobutton"` controls, the `"max"` value is assigned to the `"value"` property when the control object

is selected. For "slider" controls, "max" defines the maximum value of the slider. For "edit" and "listbox" controls, if Max - Min > 1, then the control will permit multiple line entries or list item selections, respectively.

min: scalar, def. 0
The minimum control value, whose effect on the control is dependent on the control type. For "checkbox", "togglebutton", and "radiobutton" controls, the "min" value is assigned to the "value" property when the control object is not selected. For "slider" controls, "min" defines the minimum value of the slider. For "edit" and "listbox" controls, if Max - Min > 1, then the control will permit multiple line entries or list item selections, respectively.

sliderstep: two-element vector, def. [0.010000 0.100000]
The fractional step size, measured relative to the Min - Max span of the slider, that the slider moves when the user clicks on the object. "sliderstep" is specified as a two-element vector consisting of [minor major], where "minor" is the step size for clicking on the slider arrows, and "major" is the step size for clicking within the slider bar.

value: scalar, def. 0
A numerical value associated with the current state of the control object, with the meaning of the value dependent on the "style" of the control object.

## Creation/Deletion

beingdeleted: {"off"} | "on"
Property indicating that a function has initiated deletion of the object. beingdeleted is set to true until the object no longer exists.

createfcn: string | function handle, def. [](0x0)
Callback function executed immediately after uicontrol has been created. Function is set by using default property on root object, e.g., set (groot, "defaultuicontrolcreatefcn", 'disp ("uicontrol created!")').

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

deletefcn: string | function handle, def. [](0x0)
Callback function executed immediately before uicontrol is deleted.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Display

clipping: "off" | {"on"}
If clipping is "on", the uicontrol is clipped in its parent axes limits.

visible: "off" | {"on"}
If visible is "off", the uicontrol is not rendered on screen.

## Mouse Interaction

`buttondownfcn`: string | function handle, def. `[](0x0)`
: For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`contextmenu`: graphics handle, def. `[](0x0)`
: Graphics handle of the uicontextmenu object that is currently associated to this uicontrol object.

`hittest`: `"off"` | {`"on"`}
: Specify whether uicontrol processes mouse events or passes them to ancestors of the object. When enabled, the object may respond to mouse clicks by evaluating the `"buttondownfcn"`, showing the uicontextmenu, and eventually becoming the root `"currentobject"`. This property is only relevant when the object can accept mouse clicks which is determined by the `"pickableparts"` property. See [pickableparts property], page 528.

`pickableparts`: `"all"` | `"none"` | {`"visible"`}
: Specify whether uicontrol will accept mouse clicks. By default, `pickableparts` is `"visible"` and only visible parts of the uicontrol or its children may react to mouse clicks. When `pickableparts` is `"all"` both visible and invisible parts (or children) may react to mouse clicks. When `pickableparts` is `"none"` mouse clicks on the object are ignored and transmitted to any objects underneath this one. When an object is configured to accept mouse clicks the `"hittest"` property will determine how they are processed. See [hittest property], page 528.

`selected`: {`"off"`} | `"on"`
: Property indicates whether this uicontrol is selected.

`selectionhighlight`: `"off"` | {`"on"`}
: If `selectionhighlight` is `"on"`, then the uicontrol's selection state is visually highlighted.

`tooltipstring`: string, def. `""`
: A text string that appears in a tooltip when the mouse pointer hovers over the control object.

## Object Identification

`tag`: string, def. `""`
: A user-defined string to label the graphics object.

`type` (read-only): string
: Class name of the graphics object. `type` is always `"uicontrol"`.

`userdata`: Any Octave data, def. `[](0x0)`
: User-defined data to associate with the graphics object.

## Object Position

`position`: four-element vector, def. `[0 0 80 30]`
: Size of the control object represented as the four-element vector [left, bottom, width, height].

`units`: `"centimeters"` | `"characters"` | `"inches"` | `"normalized"` | {`"pixels"`} | `"points"`
: Unit of measurement used to interpret the `"position"` property.

### Parent/Children

`children` (read-only): vector of graphics handles, def. `[](0x0)`
: Graphics handles of the uicontrol's children.

`handlevisibility`: `"callback"` | `"off"` | {`"on"`}
: If `handlevisibility` is `"off"`, the uicontrol's handle is not visible in its parent's `"children"` property.

`parent`: graphics handle
: Handle of the parent graphics object.

### Text Appearance

`fontangle`: `"italic"` | {`"normal"`}
: Control whether the font is italic or normal.

`fontname`: string, def. `"*"`
: Name of font used for text rendering. When setting this property, the text rendering engine will search for a matching font in your system. If none is found then text is rendered using a default sans serif font (same as the default `"*"` value).

  Programming Note: On systems that don't use FontConfig natively (all but Linux), the font cache is built when Octave is installed. You will need to run `system ("fc-cache -fv")` manually after installing new fonts.

`fontsize`: scalar, def. `10`
: Size of the font used for text rendering. See [fontunits property], page 529.

`fontunits`: `"centimeters"` | `"inches"` | `"normalized"` | `"pixels"` | {`"points"`}
: Units used to interpret the `"fontsize"` property.

`fontweight`: `"bold"` | {`"normal"`}
: Control the variant of the base font used for text rendering.

`horizontalalignment`: {`"center"`} | `"left"` | `"right"`
: Specifies the horizontal justification of the text within the uicontrol object.

`string`: string, def. `""`
: The text appearing with the control object.

`verticalalignment`: `"bottom"` | {`"middle"`} | `"top"`
: Specifies the vertical position of the text in the uicontrol object.

## 15.3.3.17 Uitable Properties

Properties of `uitable` objects (see [uitable], page 1011):

## Categories:

[Appearance], page 530 | [Callback Execution], page 530 | [Creation/Deletion], page 531 | [Display], page 531 | [Mouse Interaction], page 532 | [Object Identification], page 532 | [Object Position], page 533 | [Parent/Children], page 533 | [Table Data], page 533 | [Table Operation], page 534 | [Text Appearance], page 534

## Appearance

`backgroundcolor`: colorspec, def. 2-by-3 double

Color of the background of the table specified as a 3-element RBG vector. If `backgroundcolor` has multiple rows, the colors cycle repeatedly if the [`"rowstriping"` property], page 530, is on.

`foregroundcolor`: colorspec, def. `[0 0 0]`

Color of the data text in this table. See Section 15.4.1 [colorspec], page 545.

`rowstriping`: `"off"` | {`"on"`}

Setting to indicate whether the table background color will use different colors for each row. Colors are drawn from the [`"backgroundcolor"` property], page 530, in a repeating pattern.

## Callback Execution

`busyaction`: `"cancel"` | {`"queue"`}

Define how Octave handles the execution of this object's callback properties when it is unable to interrupt another object's executing callback. This is only relevant when the currently executing callback object has its `interruptible` property set to `"off"`. The `busyaction` property of the interrupting callback object indicates whether the interrupting callback is queued (`"queue"` (default)) or discarded (`"cancel"`). See Section 15.4.4 [Callbacks section], page 546.

`celleditcallback`: string, def. `[](0x0)`

A string consisting of a valid Octave expression that will be executed whenever a table cell is edited.

`cellselectioncallback`: string, def. `[](0x0)`

A string consisting of a valid Octave expression that will be executed whenever a table cell is selected.

`interruptible`: `"off"` | {`"on"`}

Specify whether this object's callback functions may be interrupted by other callbacks. By default `interruptible` is `"on"` and callbacks that make use of `drawnow`, `figure`, `waitfor`, `getframe` or `pause` functions are eventually interrupted. See Section 15.4.4 [Callbacks section], page 546.

`keypressfcn`: string | function handle, def. `[](0x0)`

Callback function executed when a key is pressed while the table object has focus. The first argument to the function is the handle of the calling table. The second argument holds an event structure which has the following members:

`Character`:

The ASCII value of the key

`Modifier:`
A cell array containing strings representing the modifiers pressed with the key.

`Key:` Lowercase description of the key

`Source:` Graphics handle of the object executing the callback function

`EventName:`
`"KeyPress"`

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`keyreleasefcn`: string | function handle, def. `[](0x0)`
Callback function executed when a key is released while the table object has focus. The first argument to the function is the handle of the calling table. The second argument holds an event structure which has the following members:

`Character:`
The ASCII value of the key

`Modifier:`
A cell array containing strings representing the modifiers pressed with the key.

`Key:` Lowercase description of the key

`Source:` Graphics handle of the object executing the callback function

`EventName:`
`"KeyRelease"`

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Creation/Deletion

`beingdeleted`: {`"off"`} | `"on"`
Property indicating that a function has initiated deletion of the object. `beingdeleted` is set to true until the object no longer exists.

`createfcn`: string | function handle, def. `[](0x0)`
Callback function executed immediately after uitable has been created. Function is set by using default property on root object, e.g., `set (groot, "defaultuitablecreatefcn", 'disp ("uitable created!")')`.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`deletefcn`: string | function handle, def. `[](0x0)`
Callback function executed immediately before uitable is deleted.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Display

`clipping`: `"off"` | {`"on"`}
: If `clipping` is `"on"`, the uitable is clipped in its parent axes limits.

`visible`: `"off"` | {`"on"`}
: If `visible` is `"off"`, the uitable is not rendered on screen.

## Mouse Interaction

`buttondownfcn`: string | function handle, def. `[](0x0)`
: For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`contextmenu`: graphics handle, def. `[](0x0)`
: Graphics handle of the uicontextmenu object that is currently associated to this uitable object.

`hittest`: `"off"` | {`"on"`}
: Specify whether uitable processes mouse events or passes them to ancestors of the object. When enabled, the object may respond to mouse clicks by evaluating the `"buttondownfcn"`, showing the uicontextmenu, and eventually becoming the root `"currentobject"`. This property is only relevant when the object can accept mouse clicks which is determined by the `"pickableparts"` property. See [pickableparts property], page 532.

`pickableparts`: `"all"` | `"none"` | {`"visible"`}
: Specify whether uitable will accept mouse clicks. By default, `pickableparts` is `"visible"` and only visible parts of the uitable or its children may react to mouse clicks. When `pickableparts` is `"all"` both visible and invisible parts (or children) may react to mouse clicks. When `pickableparts` is `"none"` mouse clicks on the object are ignored and transmitted to any objects underneath this one. When an object is configured to accept mouse clicks the `"hittest"` property will determine how they are processed. See [hittest property], page 532.

`selected`: {`"off"`} | `"on"`
: Property indicates whether this uitable is selected.

`selectionhighlight`: `"off"` | {`"on"`}
: If `selectionhighlight` is `"on"`, then the uitable's selection state is visually highlighted.

`tooltipstring`: string, def. `""`
: A text string that appears in a tooltip when the mouse pointer hovers over the table object.

## Object Identification

`tag`: string, def. `""`
: A user-defined string to label the graphics object.

`type` (read-only): string
: Class name of the graphics object. `type` is always `"uitable"`.

`userdata`: Any Octave data, def. `[](0x0)`
: User-defined data to associate with the graphics object.

## Object Position

`extent` (read-only): four-element vector
: A four-element vector indicating the size of the table. The first two elements of the array are always zero, while the third and fourth elements contain the height and width of the table.

`position`: four-element vector, def. `[20 20 300 300]`
: The position and size of the table. The four elements of the vector are the coordinates of the lower left corner and width and height of the figure. See [units property], page 533.

`units`: `"centimeters"` | `"characters"` | `"inches"` | `"normalized"` | {`"pixels"`} | `"points"`
: Unit of measurement used to interpret the `"position"` property.

## Parent/Children

`children` (read-only): vector of graphics handles, def. `[](0x0)`
: Graphics handles of the uitable's children.

`handlevisibility`: `"callback"` | `"off"` | {`"on"`}
: If `handlevisibility` is `"off"`, the uitable's handle is not visible in its parent's `"children"` property.

`parent`: graphics handle
: Handle of the parent graphics object.

## Table Data

`columnformat`: def. `{}(0x0)`
: The display format for numeric data in each column. Valid formats include `"char"`, `"logical"`, `"numeric"`, or a valid format setting from the [format function], page 288,

`columnname`: def. `"numbered"`
: Column names specified as either `"numbered"` or a 1 x N cell string vector containing the names to be used for each column heading.

`columnwidth`: def. `"auto"`
: Setting for determining width of each column, valid options include: `"auto"`, `"fit"`, evenly divided multiples specified as `"1x"`, `"2x"`, etc., or a 1 x N cell vector where each element corresponds to one of N table columns, and contains any of the above string options or a fixed width specified in pixels. The `"1x"` property is not yet implemented.

`data`: matrix, def. `[](0x0)`
: The data contained in the table specified as either a 2-D numeric, logical, or cell array.

`rowname`: def. `"numbered"`
: Row names specified as either `"numbered"` or a N x 1 cell string vector containing the names to be used for each row heading.

### Table Operation

`columneditable`: logical row vector, def. `[](0x0)`
: A logical indicator of whether the columns are editable. It consists of either a 1 x N vector of logical values where true or false indicate the corresponding column is or is not editable, respectively, or an empty logical array indicating that no column is editable.

`enable`: `"off"` | {`"on"`}
: Sets whether this table object is active or is grayed out.

`rearrangeablecolumns`: {`"off"`} | `"on"`
: Indicates whether or not the ability to move columns by clicking and dragging the column headers.

### Text Appearance

`fontangle`: `"italic"` | {`"normal"`}
: Control whether the font is italic or normal.

`fontname`: string, def. `"*"`
: Name of font used for text rendering. When setting this property, the text rendering engine will search for a matching font in your system. If none is found then text is rendered using a default sans serif font (same as the default `"*"` value).

  Programming Note: On systems that don't use FontConfig natively (all but Linux), the font cache is built when Octave is installed. You will need to run `system ("fc-cache -fv")` manually after installing new fonts.

`fontsize`: scalar, def. `10`
: Size of the font used for text rendering. See [fontunits property], page 534.

`fontunits`: `"centimeters"` | `"inches"` | `"normalized"` | `"pixels"` | {`"points"`}
: Units used to interpret the `"fontsize"` property.

`fontweight`: `"bold"` | {`"normal"`}
: Control the variant of the base font used for text rendering.

## 15.3.3.18 Uitoolbar Properties

Properties of `uitoolbar` objects (see [uitoolbar], page 1015):

### Categories:

[Callback Execution], page 534 | [Creation/Deletion], page 535 | [Display], page 535 | [Mouse Interaction], page 535 | [Object Identification], page 536 | [Parent/Children], page 536

## Callback Execution

`busyaction`: `"cancel"` | {`"queue"`}
: Define how Octave handles the execution of this object's callback properties when it is unable to interrupt another object's executing callback. This is only relevant when the currently executing callback object has its `interruptible` property set to `"off"`. The `busyaction` property of the interrupting callback object indicates whether the interrupting callback is queued (`"queue"` (default)) or discarded (`"cancel"`). See Section 15.4.4 [Callbacks section], page 546.

`interruptible`: `"off"` | {`"on"`}
: Specify whether this object's callback functions may be interrupted by other callbacks. By default `interruptible` is `"on"` and callbacks that make use of `drawnow`, `figure`, `waitfor`, `getframe` or `pause` functions are eventually interrupted. See Section 15.4.4 [Callbacks section], page 546.

## Creation/Deletion

`beingdeleted`: {`"off"`} | `"on"`
: Property indicating that a function has initiated deletion of the object. `beingdeleted` is set to true until the object no longer exists.

`createfcn`: string | function handle, def. `[](0x0)`
: Callback function executed immediately after uitoolbar has been created. Function is set by using default property on root object, e.g., `set (groot, "defaultuitoolbarcreatefcn", 'disp ("uitoolbar created!")')`.

  For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`deletefcn`: string | function handle, def. `[](0x0)`
: Callback function executed immediately before uitoolbar is deleted.

  For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Display

`clipping`: `"off"` | {`"on"`}
: If `clipping` is `"on"`, the uitoolbar is clipped in its parent axes limits.

`visible`: `"off"` | {`"on"`}
: If `visible` is `"off"`, the uitoolbar is not rendered on screen.

## Mouse Interaction

`buttondownfcn`: string | function handle, def. `[](0x0)`
: `buttondownfcn` is unused.

`contextmenu`: graphics handle, def. `[](0x0)`
: Graphics handle of the uicontextmenu object that is currently associated to this uitoolbar object.

`hittest`: `"off"` | {`"on"`}
: Specify whether uitoolbar processes mouse events or passes them to ancestors of the object. When enabled, the object may respond to mouse clicks by evaluating the `"buttondownfcn"`, showing the uicontextmenu, and eventually becoming the root `"currentobject"`. This property is only relevant when the object can accept mouse clicks which is determined by the `"pickableparts"` property. See [pickableparts property], page 536.

`pickableparts`: `"all"` | `"none"` | {`"visible"`}
: Specify whether uitoolbar will accept mouse clicks. By default, `pickableparts` is `"visible"` and only visible parts of the uitoolbar or its children may react to mouse clicks. When `pickableparts` is `"all"` both visible and invisible parts (or children) may react to mouse clicks. When `pickableparts` is `"none"` mouse clicks on the object are ignored and transmitted to any objects underneath this one. When an object is configured to accept mouse clicks the `"hittest"` property will determine how they are processed. See [hittest property], page 535.

`selected`: {`"off"`} | `"on"`
: Property indicates whether this uitoolbar is selected.

`selectionhighlight`: `"off"` | {`"on"`}
: If `selectionhighlight` is `"on"`, then the uitoolbar's selection state is visually highlighted.

## Object Identification

`tag`: string, def. `""`
: A user-defined string to label the graphics object.

`type` (read-only): string
: Class name of the graphics object. `type` is always `"uitoolbar"`.

`userdata`: Any Octave data, def. `[](0x0)`
: User-defined data to associate with the graphics object.

## Parent/Children

`children` (read-only): vector of graphics handles, def. `[](0x0)`
: Graphics handles of the uitoolbar's children.

`handlevisibility`: `"callback"` | `"off"` | {`"on"`}
: If `handlevisibility` is `"off"`, the uitoolbar's handle is not visible in its parent's `"children"` property.

`parent`: graphics handle
: Handle of the parent graphics object.

### 15.3.3.19 Uipushtool Properties

Properties of `uipushtool` objects (see [uipushtool], page 1016):

### Categories:

[Appearance], page 537 | [Callback Execution], page 537 | [Creation/Deletion], page 537 | [Display], page 538 | [Mouse Interaction], page 538 | [Object Identification], page 538 | [Parent/Children], page 539 | [Pushtool Operation], page 539

### Appearance

`__named_icon__`: string, def. `""`

The name of an bundled icon file to use as the image for the pushtool object.

`cdata`: array, def. `[](0x0)`

Image data used to represent the pushtool object, stored as a M x N x 3 RGB array.

`separator`: {`"off"`} | `"on"`

State indicating whether a separator line will be drawn next to the current pushtool position.

### Callback Execution

`busyaction`: `"cancel"` | {`"queue"`}

Define how Octave handles the execution of this object's callback properties when it is unable to interrupt another object's executing callback. This is only relevant when the currently executing callback object has its `interruptible` property set to `"off"`. The `busyaction` property of the interrupting callback object indicates whether the interrupting callback is queued (`"queue"` (default)) or discarded (`"cancel"`). See Section 15.4.4 [Callbacks section], page 546.

`clickedcallback`: string, def. `[](0x0)`

A string consisting of a valid Octave expression that will be executed whenever this control object is clicked.

`interruptible`: `"off"` | {`"on"`}

Specify whether this object's callback functions may be interrupted by other callbacks. By default `interruptible` is `"on"` and callbacks that make use of `drawnow`, `figure`, `waitfor`, `getframe` or `pause` functions are eventually interrupted. See Section 15.4.4 [Callbacks section], page 546.

### Creation/Deletion

`beingdeleted`: {`"off"`} | `"on"`

Property indicating that a function has initiated deletion of the object. `beingdeleted` is set to true until the object no longer exists.

`createfcn`: string | function handle, def. `[](0x0)`

Callback function executed immediately after uipushtool has been created. Function is set by using default property on root object, e.g., `set (groot, "defaultuipushtoolcreatefcn", 'disp ("uipushtool created!")')`.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`deletefcn`: string | function handle, def. `[](0x0)`
: Callback function executed immediately before uipushtool is deleted.

  For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Display

`clipping`: `"off"` | {`"on"`}
: If `clipping` is `"on"`, the uipushtool is clipped in its parent axes limits.

`visible`: `"off"` | {`"on"`}
: If `visible` is `"off"`, the uipushtool is not rendered on screen.

## Mouse Interaction

`buttondownfcn`: string | function handle, def. `[](0x0)`
: `buttondownfcn` is unused.

`contextmenu`: graphics handle, def. `[](0x0)`
: Graphics handle of the uicontextmenu object that is currently associated to this uipushtool object.

`hittest`: `"off"` | {`"on"`}
: Specify whether uipushtool processes mouse events or passes them to ancestors of the object. When enabled, the object may respond to mouse clicks by evaluating the `"buttondownfcn"`, showing the uicontextmenu, and eventually becoming the root `"currentobject"`. This property is only relevant when the object can accept mouse clicks which is determined by the `"pickableparts"` property. See [pickableparts property], page 538.

`pickableparts`: `"all"` | `"none"` | {`"visible"`}
: Specify whether uipushtool will accept mouse clicks. By default, `pickableparts` is `"visible"` and only visible parts of the uipushtool or its children may react to mouse clicks. When `pickableparts` is `"all"` both visible and invisible parts (or children) may react to mouse clicks. When `pickableparts` is `"none"` mouse clicks on the object are ignored and transmitted to any objects underneath this one. When an object is configured to accept mouse clicks the `"hittest"` property will determine how they are processed. See [hittest property], page 538.

`selected`: {`"off"`} | `"on"`
: Property indicates whether this uipushtool is selected.

`selectionhighlight`: `"off"` | {`"on"`}
: If `selectionhighlight` is `"on"`, then the uipushtool's selection state is visually highlighted.

`tooltipstring`: string, def. `""`
: A text string that appears in a tooltip when the mouse pointer hovers over the pushtool object.

### Object Identification

`tag`: string, def. `""`
A user-defined string to label the graphics object.

`type` (read-only): string
Class name of the graphics object. `type` is always `"uipushtool"`.

`userdata`: Any Octave data, def. `[](0x0)`
User-defined data to associate with the graphics object.

### Parent/Children

`children` (read-only): vector of graphics handles, def. `[](0x0)`
Graphics handles of the uipushtool's children.

`handlevisibility`: `"callback"` | `"off"` | {`"on"`}
If `handlevisibility` is `"off"`, the uipushtool's handle is not visible in its parent's `"children"` property.

`parent`: graphics handle
Handle of the parent graphics object.

### Pushtool Operation

`enable`: `"off"` | {`"on"`}
Sets whether this pushtool object is active or is grayed out.

## 15.3.3.20 Uitoggletool Properties

Properties of `uitoggletool` objects (see [uitoggletool], page 1017):

### Categories:

[Appearance], page 539 | [Callback Execution], page 539 | [Creation/Deletion], page 540 | [Display], page 540 | [Mouse Interaction], page 540 | [Object Identification], page 541 | [Parent/Children], page 541 | [Toggle Operation], page 542

### Appearance

`__named_icon__`: string, def. `""`
The name of an bundled icon file to use as the image for the toggletool object.

`cdata`: array, def. `[](0x0)`
Image data used to represent the toggletool object, stored as a M x N x 3 RGB array.

`separator`: {`"off"`} | `"on"`
Setting to draw a vertical line to the left of the toggletool.

### Callback Execution

`busyaction`: `"cancel"` | {`"queue"`}
Define how Octave handles the execution of this object's callback properties when it is unable to interrupt another object's executing callback. This is only

relevant when the currently executing callback object has its `interruptible` property set to `"off"`. The `busyaction` property of the interrupting callback object indicates whether the interrupting callback is queued (`"queue"` (default)) or discarded (`"cancel"`). See Section 15.4.4 [Callbacks section], page 546.

`clickedcallback`: string, def. `[](0x0)`
A string consisting of a valid Octave expression that will be executed whenever this control object is clicked.

`interruptible`: `"off" | {"on"}`
Specify whether this object's callback functions may be interrupted by other callbacks. By default `interruptible` is `"on"` and callbacks that make use of `drawnow`, `figure`, `waitfor`, `getframe` or `pause` functions are eventually interrupted. See Section 15.4.4 [Callbacks section], page 546.

`offcallback`: string, def. `[](0x0)`
A string consisting of a valid Octave expression that will be executed whenever this control object is toggled off.

`oncallback`: string, def. `[](0x0)`
A string consisting of a valid Octave expression that will be executed whenever this control object is toggled on.

## Creation/Deletion

`beingdeleted`: `{"off"} | "on"`
Property indicating that a function has initiated deletion of the object. `beingdeleted` is set to true until the object no longer exists.

`createfcn`: string | function handle, def. `[](0x0)`
Callback function executed immediately after uitoggletool has been created. Function is set by using default property on root object, e.g., `set (groot, "defaultuitoggletoolcreatefcn", 'disp ("uitoggletool created!")')`.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

`deletefcn`: string | function handle, def. `[](0x0)`
Callback function executed immediately before uitoggletool is deleted.

For information on how to write graphics listener functions, see Section 15.4.4 [Callbacks section], page 546.

## Display

`clipping`: `"off" | {"on"}`
If `clipping` is `"on"`, the uitoggletool is clipped in its parent axes limits.

`visible`: `"off" | {"on"}`
If `visible` is `"off"`, the uitoggletool is not rendered on screen.

## Mouse Interaction

`buttondownfcn`: string | function handle, def. `[](0x0)`
`buttondownfcn` is unused.

`contextmenu`: graphics handle, def. `[](0x0)`
: Graphics handle of the uicontextmenu object that is currently associated to this uitoggletool object.

`hittest`: `"off"` | {`"on"`}
: Specify whether uitoggletool processes mouse events or passes them to ancestors of the object. When enabled, the object may respond to mouse clicks by evaluating the `"buttondownfcn"`, showing the uicontextmenu, and eventually becoming the root `"currentobject"`. This property is only relevant when the object can accept mouse clicks which is determined by the `"pickableparts"` property. See [pickableparts property], page 541.

`pickableparts`: `"all"` | `"none"` | {`"visible"`}
: Specify whether uitoggletool will accept mouse clicks. By default, `pickableparts` is `"visible"` and only visible parts of the uitoggletool or its children may react to mouse clicks. When `pickableparts` is `"all"` both visible and invisible parts (or children) may react to mouse clicks. When `pickableparts` is `"none"` mouse clicks on the object are ignored and transmitted to any objects underneath this one. When an object is configured to accept mouse clicks the `"hittest"` property will determine how they are processed. See [hittest property], page 541.

`selected`: {`"off"`} | `"on"`
: Property indicates whether this uitoggletool is selected.

`selectionhighlight`: `"off"` | {`"on"`}
: If `selectionhighlight` is `"on"`, then the uitoggletool's selection state is visually highlighted.

`tooltipstring`: string, def. `""`
: A text string that appears in a tooltip when the mouse pointer hovers over the toggletool object.

## Object Identification

`tag`: string, def. `""`
: A user-defined string to label the graphics object.

`type` (read-only): string
: Class name of the graphics object. `type` is always `"uitoggletool"`.

`userdata`: Any Octave data, def. `[](0x0)`
: User-defined data to associate with the graphics object.

## Parent/Children

`children` (read-only): vector of graphics handles, def. `[](0x0)`
: Graphics handles of the uitoggletool's children.

`handlevisibility`: `"callback"` | `"off"` | {`"on"`}
: If `handlevisibility` is `"off"`, the uitoggletool's handle is not visible in its parent's `"children"` property.

`parent`: graphics handle
: Handle of the parent graphics object.

### Toggle Operation

`enable`: `"off"` | {`"on"`}
: Sets whether this toggletool object is active or is grayed out.

`state`: {`"off"`} | `"on"`
: The current `"on"` or `"off"` state of the toggletool object.

## 15.3.4 Searching Properties

*h* = `findobj` ()
*h* = `findobj` (*prop_name*, *prop_value*, ...)
*h* = `findobj` (*prop_name*, *prop_value*, "*-logical_op*", *prop_name*, *prop_value*)
*h* = `findobj` ("*-property*", *prop_name*)
*h* = `findobj` ("*-regexp*", *prop_name*, *pattern*)
*h* = `findobj` (*hlist*, ...)
*h* = `findobj` (*hlist*, "*flat*", ...)
*h* = `findobj` (*hlist*, "*-depth*", *d*, ...)

Find graphics objects with specified properties.

When called without arguments, return all graphic objects beginning with the root object (0) and including all of its descendants.

The simplest form for narrowing the results is

```
findobj (prop_name, prop_value)
```

which returns the handles of all objects which have a property named *prop_name* that has the value *prop_value*. If multiple property/value pairs are specified then only objects meeting all of the conditions (equivalent to `-and`) are returned.

The search can be limited to a particular set of objects and their descendants, by passing a handle or set of handles *hlist* as the first argument.

The depth of the object hierarchy to search can be limited with the `"-depth"` argument. An example of searching through only three generations of children is:

```
findobj (hlist, "-depth", 3, prop_name, prop_value)
```

Specifying a depth *d* of 0 limits the search to the set of objects passed in *hlist*. A depth of 0 is also equivalent to the `"flat"` argument. The default depth value is `Inf` which includes all descendants.

A specified logical operator may be used between *prop_name*, *prop_value* pairs. The supported logical operators are: `"-and"`, `"-or"`, `"-xor"`, `"-not"`. Example code to locate all figure and axes objects is

```
findobj ("type", "figure", "-or", "type", "axes")
```

Objects may also be matched by comparing a regular expression to the property values, where property values that match `regexp (prop_value, pattern)` are returned.

Finally, objects which have a property name can be found with the `"-property"` option. For example, code to locate objects with a `"meshstyle"` property is

```
findobj ("-property", "meshstyle")
```

Implementation Note: The search only includes objects with visible handles (HandleVisibility = `"on"`). See [findall], page 543, to search for all objects including hidden ones.

**See also:** [findall], page 543, [allchild], page 458, [get], page 457, [set], page 457.

`h = findall ()`
`h = findall (`*`prop_name`*`, `*`prop_value`*`, ...)`
`h = findall (`*`prop_name`*`, `*`prop_value`*`, "-logical_op", `*`prop_name`*`,`
`prop_value`*`)`
`h = findall ("-property", `*`prop_name`*`)`
`h = findall ("-regexp", `*`prop_name`*`, `*`pattern`*`)`
`h = findall (`*`hlist`*`, ...)`
`h = findall (`*`hlist`*`, "flat", ...)`
`h = findall (`*`hlist`*`, "-depth", `*`d`*`, ...)`
Find graphics object, including hidden ones, with specified properties.

The return value *h* is a list of handles to the found graphic objects.

`findall` performs the same search as `findobj`, but it includes hidden objects (HandleVisibility = `"off"`). For full documentation, see [`findobj`], page 542.

**See also:** [findobj], page 542, [allchild], page 458, [get], page 457, [set], page 457.

### 15.3.5 Managing Default Properties

Object properties have two classes of default values, *factory defaults* (the initial values) and *user-defined defaults*, which may override the factory defaults.

Although default values may be set for any object, they are set in parent objects and apply to child objects, of the specified object type. For example, setting the default `color` property of `line` objects to `"green"`, for the `root` object, will result in all `line` objects inheriting the `color` `"green"` as the default value.

```
set (groot, "defaultlinecolor", "green");
```

sets the default line color for all objects. The rule for constructing the property name to set a default value is

`default + `*`object-type`*` + `*`property-name`*

This rule can lead to some strange looking names, for example `defaultlinelinewidth"` specifies the default `linewidth` property for `line` objects.

The example above used the root object so the default property value will apply to all line objects. However, default values are hierarchical, so defaults set in a figure objects override those set in the root object. Likewise, defaults set in an axes object override those set in figure or root objects. For example,

```
subplot (2, 1, 1);
set (groot, "defaultlinecolor", "red");
set (1, "defaultlinecolor", "green");
set (gca (), "defaultlinecolor", "blue");
line (1:10, rand (1, 10));
subplot (2, 1, 2);
line (1:10, rand (1, 10));
figure (2)
line (1:10, rand (1, 10));
```

produces two figures. The line in first subplot window of the first figure is blue because it inherits its color from its parent axes object. The line in the second subplot window of the first figure is green because it inherits its color from its parent figure object. The line in the second figure window is red because it inherits its color from the global root object.

To remove a user-defined default setting, set the default property to the value `"remove"`. For example,

```
set (gca (), "defaultlinecolor", "remove");
```

removes the user-defined default line color setting from the current axes object. To quickly remove all user-defined defaults use the `reset` function.

By default, high level plotting functions such as `plot` reset and redefine axes properties independently from the defaults. An example of such property is the axes `box` property: it is set `on` by high level 2-D graphics functions regardless of the property `"defaultaxesbox"`. Use the `hold` function to prevent this behavior:

```
set (groot, "defaultaxesbox", "off");
subplot (2, 1, 1);
plot (1:10)
title ("Box is on anyway")
subplot (2, 1, 2);
hold on
plot (1:10)
title ("Box is off")
```

`reset (`*h*`)`

Reset the properties of the graphic object *h* to their default values.

For figures, the properties `"position"`, `"units"`, `"windowstyle"`, and `"paperunits"` are not affected. For axes, the properties `"position"` and `"units"` are not affected.

The input *h* may also be a vector of graphic handles in which case each individual object will be reset.

**See also:** [cla], page 430, [clf], page 429, [newplot], page 428.

Getting the `"default"` property of an object returns a list of user-defined defaults set for the object. For example,

```
get (gca (), "default");
```

returns a list of user-defined default values for the current axes object.

Factory default values are stored in the root object. The command

```
get (groot, "factory");
```

returns a list of factory defaults.

## 15.4 Advanced Plotting

### 15.4.1 Colors

Colors may be specified in three ways: 1) RGB triplets, 2) by name, or 3) by HTML notation.

RGB triplet

An RGB triplet is a 1x3 vector where each value is between 0 and 1 inclusive. The first value represents the percentage of Red, the second value the percentage of Green, and the third value the percentage of Blue. For example, `[1, 0, 1]` represents full Red and Blue channels resulting in the color magenta.

short or long name

Eight colors can be specified directly by name or by a single character short name.

| **Name** | **Color** |
|---|---|
| 'k', `"black"` | blacK |
| 'r', `"red"` | Red |
| 'g', `"green"` | Green |
| 'b', `"blue"` | Blue |
| 'y', `"yellow"` | Yellow |
| 'm', `"magenta"` | Magenta |
| 'c', `"cyan"` | Cyan |
| 'w', `"white"` | White |

HTML notation

HTML notation is a string that begins with the character '`#`' and is followed by either 3 or 6 hexadecimal digits. As with RGB triplets, each hexadecimal number represents the fraction of the Red, Green, and Blue channels present in the specified color. For example, `"#FF00FF"` represents the color magenta.

### 15.4.2 Line Styles

Line styles are specified by the following properties:

`linestyle`

May be one of

| | |
|---|---|
| `"-"` | Solid line. [default] |
| `"--"` | Dashed line. |
| `":"` | Dotted line. |
| `"-."` | A dash-dot line. |
| `"none"` | No line. Points will still be marked using the current Marker Style. |

`linewidth`

A number specifying the width of the line. The default is 1. A value of 2 is twice as wide as the default, etc.

### 15.4.3 Marker Styles

Marker styles are specified by the following properties:

`marker` A character indicating a plot marker to be place at each data point, or `"none"`, meaning no markers should be displayed.

`markeredgecolor`
The color of the edge around the marker, or `"auto"`, meaning that the edge color is the same as the face color. See Section 15.4.1 [Colors], page 545.

`markerfacecolor`
The color of the marker, or `"none"` to indicate that the marker should not be filled. See Section 15.4.1 [Colors], page 545.

`markersize`
A number specifying the size of the marker. The default is 1. A value of 2 is twice as large as the default, etc.

The `colstyle` function will parse a `plot`-style specification and will return the color, line, and marker values that would result.

`[`*`style`*`, `*`color`*`, `*`marker`*`, `*`msg`*`] = colstyle (`*`style`*`)`
Parse the line specification *style* and return the line style, color, and markers given.

In the case of an error, the string *msg* will return the text of the error.

### 15.4.4 Callbacks

Callback functions can be associated with graphics objects and triggered after certain events occur. The basic structure of all callback function is

```
function mycallback (hsrc, evt)
  ...
endfunction
```

where `hsrc` is a handle to the source of the callback, and `evt` gives some event specific data.

The function can be provided as a function handle to a plain Octave function, as an anonymous function, or as a string representing an Octave command. The latter syntax is not recommended since syntax errors will only occur when the string is evaluated. See Section 11.12 [Function Handles], page 248.

This can then be associated with an object either at the object's creation, or later with the `set` function. For example,

```
plot (x, "DeleteFcn", @(h, e) disp ("Window Deleted"))
```

where at the moment that the plot is deleted, the message `"Window Deleted"` will be displayed.

Additional user arguments can be passed to callback functions, and will be passed after the two default arguments. For example:

```
plot (x, "DeleteFcn", {@mycallback, "1"})
...
function mycallback (h, evt, arg1)
  fprintf ("Closing plot %d\n", arg1);
endfunction
```

**Caution:** The second argument in callback functions—`evt`—is only partially implemented in the Qt graphics toolkit:

- Mouse click events: `evt` is a class `double` value: 1 for left, 2 for middle, and 3 for right click.
- Key press events: `evt` is a structure with fields `Key` (string), `Character` (string), and `Modifier` (cell array of strings).
- Other events: `evt` is a class `double` empty matrix.

The basic callback functions that are available for all graphics objects are

- CreateFcn: called at the moment of the objects creation. It is not called if the object is altered in any way, and so it only makes sense to define this callback in the function call that defines the object. Callbacks that are added to `CreateFcn` later with the `set` function will never be executed.
- DeleteFcn: called at the moment an object is deleted.
- ButtonDownFcn: called if a mouse button is pressed while the pointer is over this object. Note, that the gnuplot interface does not implement this callback.

By default callback functions are queued (they are executed one after the other in the event queue) unless the `drawnow`, `figure`, `waitfor`, `getframe`, or `pause` functions are used. If an executing callback invokes one of those functions, it causes Octave to flush the event queue, which results in the executing callback being interrupted.

It is possible to specify that an object's callbacks should not be interrupted by setting the object's `interruptible` property to `"off"`. In this case, Octave decides what to do based on the `busyaction` property of the **interrupting** callback object:

`queue` (the default)
: The interrupting callback is executed after the executing callback has returned.

`cancel`
: The interrupting callback is discarded.

The `interruptible` property has no effect when the interrupting callback is a `deletefcn`, or a figure `resizefcn` or `closerequestfcn`. Those callbacks always interrupt the executing callback.

The handle to the object that holds the callback being executed can be obtained with the `gcbo` function. The handle to the ancestor figure of this object may be obtained using the `gcbf` function.

`h = gcbo ()`
`[h, fig] = gcbo ()`
: Return a handle to the object whose callback is currently executing.

  If no callback is executing, this function returns the empty matrix. This handle is obtained from the root object property `"CallbackObject"`.

  When called with a second output argument, return the handle of the figure containing the object whose callback is currently executing. If no callback is executing the second output is also set to the empty matrix.

  **See also:** [gcbf], page 548, [gco], page 456, [gca], page 455, [gcf], page 455, [get], page 457, [set], page 457.

*fig* = gcbf ()

Return a handle to the figure containing the object whose callback is currently executing.

If no callback is executing, this function returns the empty matrix. The handle returned by this function is the same as the second output argument of `gcbo`.

**See also:** [gcbo], page 547, [gcf], page 455, [gco], page 456, [gca], page 455, [get], page 457, [set], page 457.

Callbacks can equally be added to properties with the `addlistener` function described below.

### 15.4.5 Application-defined Data

Octave has a provision for attaching application-defined data to a graphics handle. The data can be anything which is meaningful to the application, and will be completely ignored by Octave.

setappdata (*h*, *name*, *value*)
setappdata (*h*, *name1*, *value1*, *name2*, *value2*, ...)
setappdata (*h*, {*name1*, *name2*, ...}, {*value1*, *value2*, ...})

Set the application data *name* to *value* for the graphics object with handle *h*.

*h* may also be a vector of graphics handles. If the application data with the specified *name* does not exist, it is created.

Multiple *name*/*value* pairs can be specified. Alternatively, a cell array of *names* and a corresponding cell array of *values* can be specified. For details on obtaining a list of valid application data properties, see [getappdata], page 548.

**See also:** [getappdata], page 548, [isappdata], page 549, [rmappdata], page 548, [guidata], page 1017, [get], page 457, [set], page 457, [getpref], page 1021, [setpref], page 1021.

*value* = getappdata (*h*, *name*)
*appdata* = getappdata (*h*)

Return the *value* of the application data *name* for the graphics object with handle *h*.

*h* may also be a vector of graphics handles. If no second argument *name* is given then `getappdata` returns a structure, *appdata*, whose fields correspond to the appdata properties.

**See also:** [setappdata], page 548, [isappdata], page 549, [rmappdata], page 548, [guidata], page 1017, [get], page 457, [set], page 457, [getpref], page 1021, [setpref], page 1021.

rmappdata (*h*, *name*)
rmappdata (*h*, *name1*, *name2*, ...)

Delete the application data *name* from the graphics object with handle *h*.

*h* may also be a vector of graphics handles. Multiple application data names may be supplied to delete several properties at once.

**See also:** [setappdata], page 548, [getappdata], page 548, [isappdata], page 549.

*valid* = isappdata (*h*, *name*)
: Return true if the named application data, *name*, exists for the graphics object with handle *h*.

  *h* may also be a vector of graphics handles.

  **See also:** [getappdata], page 548, [setappdata], page 548, [rmappdata], page 548, [guidata], page 1017, [get], page 457, [set], page 457, [getpref], page 1021, [setpref], page 1021.

### 15.4.6 Object Groups

A number of Octave high level plot functions return groups of other graphics objects or they return graphics objects that have their properties linked in such a way that changes to one of the properties results in changes in the others. A graphic object that groups other objects is an `hggroup`

hggroup ()
hggroup (*hax*)
hggroup (..., *property*, *value*, ...)
*h* = hggroup (...)
: Create handle graphics group object with axes parent *hax*.

  If no parent is specified, the group is created in the current axes.

  Multiple property/value pairs may be specified for the hggroup, but they must appear in pairs. The full list of properties is documented at Section 15.3.3.3 [Axes Properties], page 470.

  The optional return value *h* is a graphics handle to the created hggroup object.

  Programming Note: An hggroup is a way to group base graphics objects such as line objects or patch objects into a single unit which can react appropriately. For example, the individual lines of a contour plot are collected into a single hggroup so that they can be made visible/invisible with a single command, `set (hg_handle, "visible", "off")`.

  **See also:** [addproperty], page 549, [addlistener], page 550.

For example a simple use of a `hggroup` might be

```
x = 0:0.1:10;
hg = hggroup ();
plot (x, sin (x), "color", [1, 0, 0], "parent", hg);
hold on
plot (x, cos (x), "color", [0, 1, 0], "parent", hg);
set (hg, "visible", "off");
```

which groups the two plots into a single object and controls their visibility directly. The default properties of an `hggroup` are the same as the set of common properties for the other graphics objects. Additional properties can be added with the `addproperty` function.

addproperty (*name*, *h*, *type*)
addproperty (*name*, *h*, *type*, *arg*, ...)
: Create a new property named *name* in graphics object *h*.

*type* determines the type of the property to create. *args* usually contains the default value of the property, but additional arguments might be given, depending on the type of the property.

The supported property types are:

`string` A string property. *arg* contains the default string value.

`any` An un-typed property. This kind of property can hold any octave value. *args* contains the default value.

`radio` A string property with a limited set of accepted values. The first argument must be a string with all accepted values separated by a vertical bar ('|'). The default value can be marked by enclosing it with a '{' '}' pair. The default value may also be given as an optional second string argument.

`boolean` A boolean property. This property type is equivalent to a radio property with "on|off" as accepted values. *arg* contains the default property value.

`double` A scalar double property. *arg* contains the default value.

`handle` A handle property. This kind of property holds the handle of a graphics object. *arg* contains the default handle value. When no default value is given, the property is initialized to the empty matrix.

`data` A data (matrix) property. *arg* contains the default data value. When no default value is given, the data is initialized to the empty matrix.

`color` A color property. *arg* contains the default color value. When no default color is given, the property is set to black. An optional second string argument may be given to specify an additional set of accepted string values (like a radio property).

*type* may also be the concatenation of a core object type and a valid property name for that object type. The property created then has the same characteristics as the referenced property (type, possible values, hidden state...). This allows one to clone an existing property into the graphics object *h*.

Examples:

```
addproperty ("my_property", gcf, "string", "a string value");
addproperty ("my_radio", gcf, "radio", "val_1|val_2|{val_3}");
addproperty ("my_style", gcf, "linelinestyle", "--");
```

**See also:** [addlistener], page 550, [hggroup], page 549.

Once a property in added to an `hggroup`, it is not linked to any other property of either the children of the group, or any other graphics object. Add so to control the way in which this newly added property is used, the `addlistener` function is used to define a callback function that is executed when the property is altered.

`addlistener` (*h*, *prop*, *fcn*)

Register *fcn* as listener for the property *prop* of the graphics object *h*.

Property listeners are executed (in order of registration) when the property is set. The new value is already available when the listeners are executed.

*prop* must be a string naming a valid property in *h*.

*fcn* can be a function handle, a string or a cell array whose first element is a function handle. If *fcn* is a function handle, the corresponding function should accept at least 2 arguments, that will be set to the object handle and the empty matrix respectively. If *fcn* is a string, it must be any valid octave expression. If *fcn* is a cell array, the first element must be a function handle with the same signature as described above. The next elements of the cell array are passed as additional arguments to the function.

Example:

```
function my_listener (h, dummy, p1)
  fprintf ("my_listener called with p1=%s\n", p1);
endfunction

addlistener (gcf, "position", {@my_listener, "my string"})
```

**See also:** [dellistener], page 551, [addproperty], page 549, [hggroup], page 549.

`dellistener` (*h*, *prop*, *fcn*)

Remove the registration of *fcn* as a listener for the property *prop* of the graphics object *h*.

The function *fcn* must be the same variable (not just the same value), as was passed to the original call to `addlistener`.

If *fcn* is not defined then all listener functions of *prop* are removed.

Example:

```
function my_listener (h, dummy, p1)
  fprintf ("my_listener called with p1=%s\n", p1);
endfunction

c = {@my_listener, "my string"};
addlistener (gcf, "position", c);
dellistener (gcf, "position", c);
```

**See also:** [addlistener], page 550.

An example of the use of these two functions might be

```
x = 0:0.1:10;
hg = hggroup ();
h = plot (x, sin (x), "color", [1, 0, 0], "parent", hg);
addproperty ("linestyle", hg, "linelinestyle", get (h, "linestyle"));
addlistener (hg, "linestyle", @update_props);
hold on
plot (x, cos (x), "color", [0, 1, 0], "parent", hg);

function update_props (h, d)
  set (get (h, "children"), "linestyle", get (h, "linestyle"));
endfunction
```

that adds a `linestyle` property to the `hggroup` and propagating any changes its value to the children of the group. The `linkprop` function can be used to simplify the above to be

```
x = 0:0.1:10;
hg = hggroup ();
h1 = plot (x, sin (x), "color", [1, 0, 0], "parent", hg);
addproperty ("linestyle", hg, "linelinestyle", get (h, "linestyle"));
hold on
h2 = plot (x, cos (x), "color", [0, 1, 0], "parent", hg);
hlink = linkprop ([hg, h1, h2], "color");
```

*hlink* = linkprop (*h*, "*prop*")
*hlink* = linkprop (*h*, {"*prop1*", "*prop2*", ...})

Link graphic object properties, such that a change in one is propagated to the others.

The input *h* is a vector of graphic handles to link.

*prop* may be a string when linking a single property, or a cell array of strings for multiple properties. During the linking process all properties in *prop* will initially be set to the values that exist on the first object in the list *h*.

The function returns *hlink* which is a special object describing the link. As long as the reference *hlink* exists, the link between graphic objects will be active. This means that *hlink* must be preserved in a workspace variable, a global variable, or otherwise stored using a function such as `setappdata` or `guidata`. To unlink properties, execute `clear` *hlink*.

An example of the use of `linkprop` is

```
x = 0:0.1:10;
subplot (1,2,1);
h1 = plot (x, sin (x));
subplot (1,2,2);
h2 = plot (x, cos (x));
hlink = linkprop ([h1, h2], {"color","linestyle"});
set (h1, "color", "green");
set (h2, "linestyle", "--");
```

**See also:** [linkaxes], page 552, [addlistener], page 550.

linkaxes (*hax*)
linkaxes (*hax*, *optstr*)

Link the axis limits of 2-D plots such that a change in one is propagated to the others.

The axes handles to be linked are passed as the first argument *hax*.

The optional second argument is a string which defines which axis limits will be linked. The possible values for *optstr* are:

`"x"` Link x-axes

`"y"` Link y-axes

`"xy"` (default)
Link both axes

`"off"` Turn off linking

If unspecified the default is to link both X and Y axes.

When linking, the limits from the first axes in *hax* are applied to the other axes in the list. Subsequent changes to any one of the axes will be propagated to the others.

**See also:** [linkprop], page 552, [addproperty], page 549.

These capabilities are used in a number of basic graphics objects. The `hggroup` objects created by the functions of Octave contain one or more graphics object and are used to:

- group together multiple graphics objects,
- create linked properties between different graphics objects, and
- to hide the nominal user data, from the actual data of the objects.

For example the `stem` function creates a stem series where each `hggroup` of the stem series contains two line objects representing the body and head of the stem. The `ydata` property of the `hggroup` of the stem series represents the head of the stem, whereas the body of the stem is between the baseline and this value. For example

```
h = stem (1:4)
get (h, "xdata")
⇒ [  1   2   3   4]'
get (get (h, "children")(1), "xdata")
⇒ [  1   1 NaN   2   2 NaN   3   3 NaN   4   4 NaN]'
```

shows the difference between the `xdata` of the `hggroup` of a stem series object and the underlying line.

The basic properties of such group objects is that they consist of one or more linked `hggroup`, and that changes in certain properties of these groups are propagated to other members of the group. Whereas, certain properties of the members of the group only apply to the current member.

In addition the members of the group can also be linked to other graphics objects through callback functions. For example the baseline of the `bar` or `stem` functions is a line object, whose length and position are automatically adjusted, based on changes to the corresponding hggroup elements.

### 15.4.6.1 Data Sources in Object Groups

All of the group objects contain data source parameters. There are string parameters that contain an expression that is evaluated to update the relevant data property of the group when the `refreshdata` function is called.

`refreshdata` ()
`refreshdata` (*h*)
`refreshdata` (*h*, *workspace*)

Evaluate any '`datasource`' properties of the current figure and update the plot if the corresponding data has changed.

If the first argument *h* is a list of graphics handles to figure, axes, or graphic objects with a DataSource property, then operate on these objects rather than the current figure returned by `gcf`.

The optional second argument *workspace* can take the following values:

`"base"` Evaluate the datasource properties in the base workspace. (default).

`"caller"` Evaluate the datasource properties in the workspace of the function that called `refreshdata`.

An example of the use of `refreshdata` is:

```
x = 0:0.1:10;
y = sin (x);
plot (x, y, "ydatasource", "y");
for i = 1 : 100
  pause (0.1);
  y = sin (x + 0.1*i);
  refreshdata ();
endfor
```

Programming Note: For performance, specify *h* as the actual object(s) to be updated. If no object is supplied then Octave must search through all graphic objects of the current figure and determine which ones have DataSource properties that are not empty.

#### 15.4.6.2 Area Series

Area series objects are created by the `area` function. Each of the `hggroup` elements contains a single patch object. The properties of the area series are

`basevalue`
: The value where the base of the area plot is drawn.

`linewidth`
`linestyle`
: The line width and style of the edge of the patch objects making up the areas. See Section 15.4.2 [Line Styles], page 545.

`edgecolor`
`facecolor`
: The line and fill color of the patch objects making up the areas. See Section 15.4.1 [Colors], page 545.

`xdata`
`ydata`
: The x and y coordinates of the original columns of the data passed to `area` prior to the cumulative summation used in the `area` function.

`xdatasource`
`ydatasource`
: Data source variables.

#### 15.4.6.3 Bar Series

Bar series objects are created by the `bar` or `barh` functions. Each `hggroup` element contains a single patch object. The properties of the bar series are

`showbaseline`
`baseline`
`basevalue`
: The property `showbaseline` flags whether the baseline of the bar series is displayed (default is `"on"`). The handle of the graphics object representing the

baseline is given by the `baseline` property and the y-value of the baseline by the `basevalue` property.

Changes to any of these properties are propagated to the other members of the bar series and to the baseline itself. Equally, changes in the properties of the base line itself are propagated to the members of the corresponding bar series.

`barwidth`
`barlayout`
`horizontal`

The property `barwidth` is the width of the bar corresponding to the *width* variable passed to `bar` or *barh*. Whether the bar series is `"grouped"` or `"stacked"` is determined by the `barlayout` property and whether the bars are horizontal or vertical by the `horizontal` property.

Changes to any of these property are propagated to the other members of the bar series.

`linewidth`
`linestyle`

The line width and style of the edge of the patch objects making up the bars. See Section 15.4.2 [Line Styles], page 545.

`edgecolor`
`facecolor`

The line and fill color of the patch objects making up the bars. See Section 15.4.1 [Colors], page 545.

`xdata` The nominal x positions of the bars. Changes in this property and propagated to the other members of the bar series.

`ydata` The y value of the bars in the `hggroup`.

`xdatasource`
`ydatasource`

Data source variables.

### 15.4.6.4 Contour Groups

Contour group objects are created by the `contour`, `contourf`, and `contour3` functions. They are also one of the handles returned by the `surfc` and `meshc` functions. The properties of the contour group are

`contourmatrix`

A read only property that contains the data return by `contourc` used to create the contours of the plot.

`fill` A radio property that can have the values `"on"` or `"off"` that flags whether the contours to plot are to be filled.

`zlevelmode`
`zlevel` The radio property `zlevelmode` can have the values `"none"`, `"auto"`, or `"manual"`. When its value is `"none"` there is no z component to the plotted contours. When its value is `"auto"` the z value of the plotted contours is at

the same value as the contour itself. If the value is `"manual"`, then the z value at which to plot the contour is determined by the `zlevel` property.

`levellistmode`
`levellist`
`levelstepmode`
`levelstep`
: If `levellistmode` is `"manual"`, then the levels at which to plot the contours is determined by `levellist`. If `levellistmode` is set to `"auto"`, then the distance between contours is determined by `levelstep`. If both `levellistmode` and `levelstepmode` are set to `"auto"`, then there are assumed to be 10 equal spaced contours.

`textlistmode`
`textlist`
`textstepmode`
`textstep`
: If `textlistmode` is `"manual"`, then the labeled contours is determined by `textlist`. If `textlistmode` is set to `"auto"`, then the distance between labeled contours is determined by `textstep`. If both `textlistmode` and `textstepmode` are set to `"auto"`, then there are assumed to be 10 equal spaced labeled contours.

`showtext`
: Flag whether the contour labels are shown or not.

`labelspacing`
: The distance between labels on a single contour in points.

`linewidth`
`linestyle`
`linecolor`
: The properties of the contour lines. The properties `linewidth` and `linestyle` are similar to the corresponding properties for lines. The property `linecolor` is a color property (see Section 15.4.1 [Colors], page 545), that can also have the values of `"none"` or `"auto"`. If `linecolor` is `"none"`, then no contour line is drawn. If `linecolor` is `"auto"` then the line color is determined by the colormap.

`xdata`
`ydata`
`zdata`
: The original x, y, and z data of the contour lines.

`xdatasource`
`ydatasource`
`zdatasource`
: Data source variables.

### 15.4.6.5 Error Bar Series

Error bar series are created by the `errorbar` function. Each `hggroup` element contains two line objects representing the data and the errorbars separately. The properties of the error bar series are

`color` The RGB color or color name of the line objects of the error bars. See Section 15.4.1 [Colors], page 545.

`linewidth`
`linestyle`
The line width and style of the line objects of the error bars. See Section 15.4.2 [Line Styles], page 545.

`marker`
`markeredgecolor`
`markerfacecolor`
`markersize`
The line and fill color of the markers on the error bars. See Section 15.4.1 [Colors], page 545.

`xdata`
`ydata`
`ldata`
`udata`
`xldata`
`xudata` The original x, y, l, u, xl, xu data of the error bars.

`xdatasource`
`ydatasource`
`ldatasource`
`udatasource`
`xldatasource`
`xudatasource`
Data source variables.

#### 15.4.6.6 Line Series

Line series objects are created by the `plot` and `plot3` functions and are of the type `line`. The properties of the line series with the ability to add data sources.

`color` The RGB color or color name of the line objects. See Section 15.4.1 [Colors], page 545.

`linewidth`
`linestyle`
The line width and style of the line objects. See Section 15.4.2 [Line Styles], page 545.

`marker`
`markeredgecolor`
`markerfacecolor`
`markersize`
The line and fill color of the markers. See Section 15.4.1 [Colors], page 545.

`xdata`
`ydata`
`zdata` The original x, y and z data.

`xdatasource`
`ydatasource`
`zdatasource`
: Data source variables.

### 15.4.6.7 Quiver Group

Quiver series objects are created by the `quiver` or `quiver3` functions. Each `hggroup` element of the series contains three line objects as children representing the body and head of the arrow, together with a marker as the point of origin of the arrows. The properties of the quiver series are

`autoscale`
`autoscalefactor`
: Flag whether the length of the arrows is scaled or defined directly from the *u*, *v* and *w* data. If the arrow length is flagged as being scaled by the `autoscale` property, then the length of the autoscaled arrow is controlled by the `autoscalefactor`.

`maxheadsize`
: This property controls the size of the head of the arrows in the quiver series. The default value is 0.2.

`showarrowhead`
: Flag whether the arrow heads are displayed in the quiver plot.

`color` : The RGB color or color name of the line objects of the quiver. See Section 15.4.1 [Colors], page 545.

`linewidth`
`linestyle`
: The line width and style of the line objects of the quiver. See Section 15.4.2 [Line Styles], page 545.

`marker`
`markerfacecolor`
`markersize`
: The line and fill color of the marker objects at the original of the arrows. See Section 15.4.1 [Colors], page 545.

`xdata`
`ydata`
`zdata` : The origins of the values of the vector field.

`udata`
`vdata`
`wdata` : The values of the vector field to plot.

`xdatasource`
`ydatasource`
`zdatasource`
`udatasource`
`vdatasource`
`wdatasource`
: Data source variables.

### 15.4.6.8 Stair Group

Stair series objects are created by the `stair` function. Each `hggroup` element of the series contains a single line object as a child representing the stair. The properties of the stair series are

`color`
: The RGB color or color name of the line objects of the stairs. See Section 15.4.1 [Colors], page 545.

`linewidth`
`linestyle`
: The line width and style of the line objects of the stairs. See Section 15.4.2 [Line Styles], page 545.

`marker`
`markeredgecolor`
`markerfacecolor`
`markersize`
: The line and fill color of the markers on the stairs. See Section 15.4.1 [Colors], page 545.

`xdata`
`ydata`
: The original x and y data of the stairs.

`xdatasource`
`ydatasource`
: Data source variables.

### 15.4.6.9 Stem Series

Stem series objects are created by the `stem` or `stem3` functions. Each `hggroup` element contains a single line object as a child representing the stems. The properties of the stem series are

`showbaseline`
`baseline`
`basevalue`
: The property `showbaseline` flags whether the baseline of the stem series is displayed (default is `"on"`). The handle of the graphics object representing the baseline is given by the `baseline` property and the y-value (or z-value for `stem3`) of the baseline by the `basevalue` property.

  Changes to any of these property are propagated to the other members of the stem series and to the baseline itself. Equally changes in the properties of the base line itself are propagated to the members of the corresponding stem series.

`color` The RGB color or color name of the line objects of the stems. See Section 15.4.1 [Colors], page 545.

`linewidth`
`linestyle`
The line width and style of the line objects of the stems. See Section 15.4.2 [Line Styles], page 545.

`marker`
`markeredgecolor`
`markerfacecolor`
`markersize`
The line and fill color of the markers on the stems. See Section 15.4.1 [Colors], page 545.

`xdata`
`ydata`
`zdata` The original x, y and z data of the stems.

`xdatasource`
`ydatasource`
`zdatasource`
Data source variables.

### 15.4.6.10 Surface Group

Surface group objects are created by the `surf` or `mesh` functions, but are equally one of the handles returned by the `surfc` or `meshc` functions. The surface group is of the type `surface`.

The properties of the surface group are

`edgecolor`
`facecolor`
The RGB color or color name of the edges or faces of the surface. See Section 15.4.1 [Colors], page 545.

`linewidth`
`linestyle`
The line width and style of the lines on the surface. See Section 15.4.2 [Line Styles], page 545.

`marker`
`markeredgecolor`
`markerfacecolor`
`markersize`
The line and fill color of the markers on the surface. See Section 15.4.1 [Colors], page 545.

`xdata`
`ydata`
`zdata`
`cdata` The original x, y, z and c data.

`xdatasource`
`ydatasource`
`zdatasource`
`cdatasource`
: Data source variables.

### 15.4.7 Transform Groups

*h* = hgtransform ()
*h* = hgtransform (*property*, *value*, ...)
*h* = hgtransform (*hax*, ...)
: Create a graphics transform object.

  FIXME: Need to write documentation. FIXME: Add <makehgtform> to seealso list when it is implemented.

  **See also:** [hggroup], page 549.

### 15.4.8 Graphics Toolkits

*tkit* = graphics_toolkit ()
*tkit* = graphics_toolkit (*hlist*)
graphics_toolkit (*name*)
graphics_toolkit (*hlist*, *name*)
: Query or set the default graphics toolkit which is assigned to new figures.

  With no inputs, return the current default graphics toolkit. If the input is a list of figure graphic handles, *hlist*, then return the name of the graphics toolkit in use for each figure.

  When called with a single input *name* set the default graphics toolkit to *name*. If the toolkit is not already loaded, it is initialized by calling the function `__init_name__`. If the first input is a list of figure handles, *hlist*, then the graphics toolkit is set to *name* for these figures only.

  **See also:** [available_graphics_toolkits], page 561.

*toolkits* = available_graphics_toolkits ()
: Return a cell array of registered graphics toolkits.

  **See also:** [graphics_toolkit], page 561, [register_graphics_toolkit], page 561.

*toolkits* = loaded_graphics_toolkits ()
: Return a cell array of the currently loaded graphics toolkits.

  **See also:** [available_graphics_toolkits], page 561.

register_graphics_toolkit ("*toolkit*")
: List *toolkit* as an available graphics toolkit.

  Programming Note: No input validation is done on the input string; it is simply added to the list of possible graphics toolkits.

  **See also:** [available_graphics_toolkits], page 561.

### 15.4.8.1 Customizing Toolkit Behavior

The specific behavior of the backend toolkit may be modified using the following utility functions. Note: Not all functions apply to every graphics toolkit.

`[prog, args] = gnuplot_binary ()`
`[old_prog, old_args] = gnuplot_binary (new_prog)`
`[old_prog, old_args] = gnuplot_binary (new_prog, arg1, ...)`

Query or set the name of the program invoked by the plot command when the graphics toolkit is set to `"gnuplot"`.

Additional arguments to pass to the external plotting program may also be given. The default value is `"gnuplot"` with no additional arguments. See Appendix E [Installing Octave], page 1179.

**See also:** [graphics_toolkit], page 561.

In addition, the gnuplot program usually provides a number of different interfaces, known as terminals. Octave normally chooses a default terminal, but you can override this with the environment variable `GNUTERM`. This variable may be set in the shell before starting Octave or from within Octave before plotting for the first time. For example:

```
setenv ("GNUTERM", "wxt")
graphics_toolkit ("gnuplot")
plot (1:10)
```

### 15.4.8.2 Hardware vs Software Rendering

On some platforms, it is possible to choose between hardware (GPU-accelerated) and software (CPU-based) OpenGL rendering through the use of the Mesa 3D software implementation. The choice between hardware or software rendering only affects the OpenGL graphics toolkits (`"qt"`, `"fltk"`). Using software rendering can avoid rendering and printing issues due to the imperfect OpenGL driver implementations of diverse graphic cards from different vendors (notably integrated Intel graphics). The downside is that software rendering may be considerably slower than hardware-accelerated rendering.

On Linux and MacOS systems that have a Mesa 3-D based driver, switching to software rendering is done at Octave startup (or at least before any graphics function has been called) by setting the environment variable `LIBGL_ALWAYS_SOFTWARE`, e.g., `setenv ("LIBGL_ALWAYS_SOFTWARE", "1")`.

On Windows, Octave is shipped with `opengl32.dll`, a Mesa 3-D based software rendering library. When using the Windows installer for Octave, the user has the option to select between `"System OpenGL"` and `"Software OpenGL"` renderers. The OpenGL driver may also be switched later using the `"OpenGL Switcher"` application from the Start menu while Octave is closed. Alternatively, one can rename the following file while Octave is closed:

*octave-home*`\bin\opengl32.dll`

where *octave-home* is the directory returned by [OCTAVE_HOME], page 1074, i.e., the directory in which Octave is installed (the default is `C:\Program Files\GNU Octave\Octave\Octave-`*version*`\mingw64`). Change the file extension to `.bak` for hardware rendering or to `.dll` for software rendering.

### 15.4.8.3 Precision issues

The OpenGL graphics toolkits (`"qt"` and `"fltk"`) use single precision for rendering. This limitation in particular applies to plots of time series against serial dates as used by the `datenum`, `datestr`, `datestruct`, and `datetick` functions.

Serial dates encode timestamps as days elapsed since the year zero with hours, minutes, seconds as the fractional part. On December 31st 1999, the serial date representation was 730485. A double precision variable with this integer part allows for a resolution in its fractional part of 1.2e-10, representing about 5 microseconds. But with single precision, the resolution is reduced to about 0.06, representing 45 minutes. Any attempt to plot timestamped data with finer granularity will result in a distorted graph.

As a workaround, it is possible to use the `"gnuplot"` graphics toolkit or subtract 2000 years—i.e., `datenum (2000, 0, 0)` or 730485—from the time values. Due to the fact that the calendar structure repeats every 2000 years, the relation between year, month, day of month and day of week will stay unchanged and the ticks and ticklabels produced by the `datetick` function will still be correct. Only years will lack the millennium digit. Thus, `"2020"` will be printed as `"20"`. For example:

```
# timestamps of 24 hours in one minute steps
t = datenum (2020, 1, 1):(1/1440):datenum (2020, 1, 2);

# some example time series data
x = -cos (2*pi*t) + rand (size (t)) / 10;

subplot (1, 2, 1);
plot (t, x);
datetick ("x");
xlabel ("serial date");
title ("problem");

subplot (1, 2, 2);
plot (t - 730485, x);
datetick ("x");
xlabel ("2000 years off");
title ("workaround");
```

The result of which can be seen in Figure 15.8.

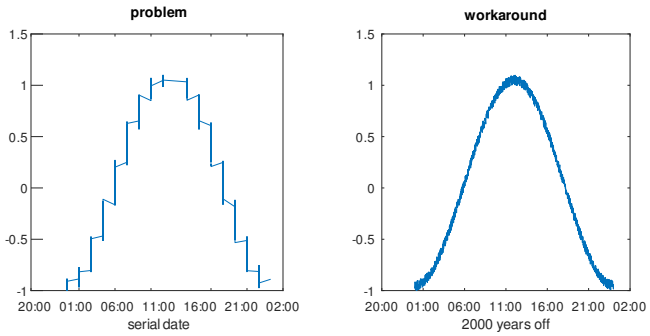



Figure 15.8: Single precision issues with OpenGL graphics toolkits

Similarly, other data can be translated or re-scaled to work around this issue.

# 16 Matrix Manipulation

There are a number of functions available for checking to see if the elements of a matrix meet some condition, and for rearranging the elements of a matrix. For example, Octave can easily tell you if all the elements of a matrix are finite, or are less than some specified value. Octave can also rotate the elements, extract the upper- or lower-triangular parts, or sort the columns of a matrix.

## 16.1 Finding Elements and Checking Conditions

The functions `any` and `all` are useful for determining whether any or all of the elements of a matrix satisfy some condition. The `find` function is also useful in determining which elements of a matrix meet a specified condition.

*tf* = `any` (*x*)
*tf* = `any` (*x*, *dim*)
*tf* = `any` (*x*, *vecdim*)
*tf* = `any` (*x*, "*all*")

Return true (logical 1) if any elements are nonzero or true.

If *x* is a vector, then `any (x)` returns true (logical 1) if any elements of the vector are nonzero.

If *x* is a matrix, then `any (x)` returns a row vector of logical ones and zeros where each element indicates whether any of the elements of the corresponding column of the matrix are nonzero.

If *x* is an array, then `any(x)` operates along the first non-singleton dimension of *x* and returns a logical array, whose size is equal to *x* except for the operating dimension which becomes 1.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension in *x*, including any dimension exceeding `ndims (x)`, will return *x*.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored. The size of the dimensions specified by *vecdim* become 1 in the returned logical array.

Specifying the dimension as `"all"` will cause `any` to operate on all elements of *x*, and is equivalent to `any (x(:))`.

**See also:** [all], page 565.

*tf* = `all` (*x*)
*tf* = `all` (*x*, *dim*)
*tf* = `all` (*x*, *vecdim*)
*tf* = `all` (*x*, "*all*")

Return true (logical 1) if all elements are nonzero or true.

If *x* is a vector, then `all (x)` returns true (logical 1) if all elements of the vector are nonzero.

If *x* is a matrix, then `all (x)` returns a row vector of logical ones and zeros where each element indicates whether all of the elements of the corresponding column of the matrix are nonzero.

If *x* is an array, then `all(x)` operates along the first non-singleton dimension of *x* and returns a logical array, whose size is equal to *x* except for the operating dimension which becomes 1.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension in *x*, including any dimension exceeding `ndims (x)`, will return *x*.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored. The size of the dimensions specified by *vecdim* become 1 in the returned logical array.

Specifying the dimension as `"all"` will cause `all` to operate on all elements of *x*, and is equivalent to `all (x(:))`.

**See also:** [any], page 565.

Since the comparison operators (see Section 8.4 [Comparison Ops], page 174) return matrices of ones and zeros, it is easy to test a matrix for many things, not just whether the elements are nonzero. For example,

```
all (all (rand (5) < 0.9))
     ⇒ 0
```

tests a random 5 by 5 matrix to see if all of its elements are less than 0.9.

Note that in conditional contexts (like the test clause of `if` and `while` statements) Octave treats the test as if you had typed `all (all (condition))`.

*z* = `xor` (*x*, *y*)
*z* = `xor` (*x1*, *x2*, ...)

Return the *exclusive or* of *x* and *y*.

For boolean expressions *x* and *y*, `xor (x, y)` is true if and only if one of *x* or *y* is true. Otherwise, if *x* and *y* are both true or both false, `xor` returns false.

The truth table for the xor operation is

| *x* | *y* | *z* |
|---|---|---|
| - | - | - |
| 0 | 0 | 0 |
| 1 | 0 | 1 |
| 0 | 1 | 1 |
| 1 | 1 | 0 |

If more than two arguments are given the xor operation is applied cumulatively from left to right:

```
(...((x1 XOR x2) XOR x3) XOR ...)
```

**See also:** [and], page 176, [or], page 177, [not], page 177.

`y = diff (x)`
`y = diff (x, k)`
`y = diff (x, k, dim)`

If $x$ is a vector of length $n$, `diff (x)` is the vector of first differences $x_2 - x_1, \ldots, x_n - x_{n-1}$.

If $x$ is a matrix, `diff (x)` is the matrix of column differences along the first non-singleton dimension.

The second argument is optional. If supplied, `diff (x, k)`, where $k$ is a non-negative integer, returns the $k$-th differences. It is possible that $k$ is larger than the first non-singleton dimension of the matrix. In this case, `diff` continues to take the differences along the next non-singleton dimension.

The dimension along which to take the difference can be explicitly stated with the optional variable *dim*. In this case the $k$-th order differences are calculated along this dimension. In the case where $k$ exceeds `size (x, dim)` an empty matrix is returned.

**See also:** [sort], page 575, [merge], page 178.

`tf = isinf (x)`

Return a logical array which is true where the elements of $x$ are infinite and false where they are not.

For example:

```
isinf ([13, Inf, NA, NaN])
      ⇒ [ 0, 1, 0, 0 ]
```

**See also:** [isfinite], page 567, [isnan], page 567, [isna], page 46.

`tf = isnan (x)`

Return a logical array which is true where the elements of $x$ are NaN values and false where they are not.

NA values are also considered NaN values. For example:

```
isnan ([13, Inf, NA, NaN])
      ⇒ [ 0, 0, 1, 1 ]
```

**See also:** [isna], page 46, [isinf], page 567, [isfinite], page 567.

`tf = isfinite (x)`

Return a logical array which is true where the elements of $x$ are finite values and false where they are not.

For example:

```
isfinite ([13, Inf, NA, NaN])
     ⇒ [ 1, 0, 0, 0 ]
```

**See also:** [isinf], page 567, [isnan], page 567, [isna], page 46.

`[err, yi, ...] = common_size (xi, ...)`

Determine if all input arguments are either scalar or of common size.

If true, *err* is zero, and *yi* is a matrix of the common size with all entries equal to *xi* if this is a scalar or *xi* otherwise. If the inputs cannot be brought to a common size, *err* is 1, and *yi* is *xi*. For example:

```
[err, a, b] = common_size ([1 2; 3 4], 5)
     ⇒ err = 0
     ⇒ a = [ 1, 2; 3, 4 ]
     ⇒ b = [ 5, 5; 5, 5 ]
```

This is useful for implementing functions where arguments can either be scalars or of common size.

**See also:** [size], page 48, [size_equal], page 49, [numel], page 47, [ndims], page 46.

*idx* = `find` (*x*)
*idx* = `find` (*x*, *n*)
*idx* = `find` (*x*, *n*, *direction*)
`[i, j] = find (...)`
`[i, j, v] = find (...)`

Return a vector of indices of nonzero elements of a matrix, as a row if *x* is a row vector or as a column otherwise.

To obtain a single index for each matrix element, Octave pretends that the columns of a matrix form one long vector (like Fortran arrays are stored). For example:

```
find (eye (2))
  ⇒ [ 1; 4 ]
```

If two inputs are given, *n* indicates the maximum number of elements to find from the beginning of the matrix or vector.

If three inputs are given, *direction* should be one of `"first"` or `"last"`, requesting only the first or last *n* indices, respectively. However, the indices are always returned in ascending order.

If two outputs are requested, `find` returns the row and column indices of nonzero elements of a matrix. For example:

```
[i, j] = find (2 * eye (2))
    ⇒ i = [ 1; 2 ]
    ⇒ j = [ 1; 2 ]
```

If three outputs are requested, `find` also returns a vector containing the nonzero values. For example:

```
[i, j, v] = find (3 * eye (2))
       ⇒ i = [ 1; 2 ]
       ⇒ j = [ 1; 2 ]
       ⇒ v = [ 3; 3 ]
```

If *x* is a multi-dimensional array of size m x n x p x ..., *j* contains the column locations as if *x* was flattened into a two-dimensional matrix of size m x (n + p + ...).

Note that this function is particularly useful for sparse matrices, as it extracts the nonzero elements as vectors, which can then be used to create the original matrix. For example:

```
sz = size (a);
[i, j, v] = find (a);
b = sparse (i, j, v, sz(1), sz(2));
```

**See also:** [nonzeros], page 734.

*idx* = lookup (*table*, *y*)
*idx* = lookup (*table*, *y*, *opt*)

Lookup values in a **sorted** table.

This function is usually used as a prelude to interpolation.

If table is increasing, of length N and `idx = lookup (table, y)`, then `table(idx(i)) <= y(i) < table(idx(i+1))` for all `y(i)` within the table. If `y(i) < table(1)` then `idx(i)` is 0. If `y(i) >= table(end)` or `isnan (y(i))` then `idx(i)` is N.

If the table is decreasing, then the tests are reversed. For non-strictly monotonic tables, empty intervals are always skipped. The result is undefined if *table* is not monotonic, or if *table* contains a NaN.

The complexity of the lookup is O(M*log(N)) where M is the size of *y*. In the special case when *y* is also sorted, the complexity is O(min (M*log(N), M+N)).

*table* and *y* can also be cell arrays of strings (or *y* can be a single string). In this case, string lookup is performed using lexicographical comparison.

If *opts* is specified, it must be a string with letters indicating additional options.

`m` Match. `table(idx(i)) == y(i)` if `y(i)` occurs in table; otherwise, `idx(i)` is zero.

`b` Boolean. `idx(i)` is a logical 1 or 0, indicating whether `y(i)` is contained in table or not.

`l` Left. For numeric lookups the leftmost subinterval shall be extended to minus infinity (i.e., all indices at least 1).

`r` Right. For numeric lookups the rightmost subinterval shall be extended to infinity (i.e., all indices at most N-1).

**Note**: If *table* is not sorted the results from `lookup` will be unpredictable.

If you wish to check if a variable exists at all, instead of properties its elements may have, see Section 7.3 [Status of Variables], page 150.

## 16.2 Rearranging Matrices

*B* = fliplr (*A*)

Flip array left to right.

Return a copy of *A* with the order of the columns reversed. In other words, *A* is flipped left-to-right about a vertical axis. For example:

```
fliplr ([1, 2; 3, 4])
     ⇒  2  1
        4  3
```

**See also:** [flipud], page 570, [flip], page 570, [rot90], page 570, [rotdim], page 571.

`B = flipud (A)`

Flip array upside down.

Return a copy of $A$ with the order of the rows reversed. In other words, $A$ is flipped upside-down about a horizontal axis. For example:

```
flipud ([1, 2; 3, 4])
     ⇒  3  4
        1  2
```

**See also:** [fliplr], page 569, [flip], page 570, [rot90], page 570, [rotdim], page 571.

`B = flip (A)`
`B = flip (A, dim)`

Return a copy of array $A$ flipped across dimension *dim*.

If *dim* is unspecified it defaults to the first non-singleton dimension.

Examples:

```
## row vector
flip ([1  2  3  4])
      ⇒  4  3  2  1

## column vector
flip ([1; 2; 3; 4])
      ⇒  4
         3
         2
         1

## 2-D matrix along dimension 1
flip ([1 2; 3 4])
      ⇒  3  4
         1  2

## 2-D matrix along dimension 2
flip ([1 2; 3 4], 2)
      ⇒  2  1
         4  3
```

**See also:** [fliplr], page 569, [flipud], page 570, [rot90], page 570, [rotdim], page 571, [permute], page 572, [transpose], page 174.

`B = rot90 (A)`
`B = rot90 (A, k)`

Rotate array by 90 degree increments.

Return a copy of $A$ with the elements rotated counterclockwise in 90-degree increments.

The second argument is optional, and specifies how many 90-degree rotations are to be applied (the default value is 1). Negative values of $k$ rotate the matrix in a clockwise direction. For example,

```
rot90 ([1, 2; 3, 4], -1)
    ⇒  3  1
       4  2
```

rotates the given matrix clockwise by 90 degrees. The following are all equivalent statements:

```
rot90 ([1, 2; 3, 4], -1)
rot90 ([1, 2; 3, 4], 3)
rot90 ([1, 2; 3, 4], 7)
```

The rotation is always performed on the plane of the first two dimensions, i.e., rows and columns. To perform a rotation on any other plane, use `rotdim`.

**See also:** [rotdim], page 571, [fliplr], page 569, [flipud], page 570, [flip], page 570.

`B = rotdim (A)`
`B = rotdim (A, n)`
`B = rotdim (A, n, plane)`

Return a copy of $A$ with the elements rotated counterclockwise in 90-degree increments.

The second argument $n$ is optional, and specifies how many 90-degree rotations are to be applied (the default value is 1). Negative values of $n$ rotate the matrix in a clockwise direction.

The third argument is also optional and defines the plane of the rotation. If present, *plane* is a two element vector containing two different valid dimensions of the matrix. When *plane* is not given the first two non-singleton dimensions are used.

For example,

```
rotdim ([1, 2; 3, 4], -1, [1, 2])
     ⇒  3  1
        4  2
```

rotates the given matrix clockwise by 90 degrees. The following are all equivalent statements:

```
rotdim ([1, 2; 3, 4], -1, [1, 2])
rotdim ([1, 2; 3, 4], 3, [1, 2])
rotdim ([1, 2; 3, 4], 7, [1, 2])
```

**See also:** [rot90], page 570, [fliplr], page 569, [flipud], page 570, [flip], page 570.

`A = cat (dim, array1, array2, ..., arrayN)`

Return the concatenation of N-D array objects, *array1*, *array2*, ..., *arrayN* along dimension *dim*.

```
A = ones (2, 2);
B = zeros (2, 2);
cat (2, A, B)
  ⇒ 1 1 0 0
    1 1 0 0
```

Alternatively, we can concatenate $A$ and $B$ along the second dimension in the following way:

```
[A, B]
```

*dim* can be larger than the dimensions of the N-D array objects and the result will thus have *dim* dimensions as the following example shows:

```
cat (4, ones (2, 2), zeros (2, 2))
  ⇒ ans(:,:,1,1) =

      1 1
      1 1

    ans(:,:,1,2) =

      0 0
      0 0
```

**See also:** [horzcat], page 572, [vertcat], page 572.

`A = horzcat (`*array1*`, `*array2*`, ..., `*arrayN*`)`

Return the horizontal concatenation of N-D array objects, *array1*, *array2*, ..., *arrayN* along dimension 2.

Arrays may also be concatenated horizontally using the syntax for creating new matrices. For example:

```
A = [ array1, array2, ... ]
```

This syntax is slightly more efficient because the Octave parser can concatenate the arrays without the overhead of a function call.

**See also:** [cat], page 571, [vertcat], page 572.

`A = vertcat (`*array1*`, `*array2*`, ..., `*arrayN*`)`

Return the vertical concatenation of N-D array objects, *array1*, *array2*, ..., *arrayN* along dimension 1.

Arrays may also be concatenated vertically using the syntax for creating new matrices. For example:

```
A = [ array1; array2; ... ]
```

This syntax is slightly more efficient because the Octave parser can concatenate the arrays without the overhead of a function call.

**See also:** [cat], page 571, [horzcat], page 572.

`B = permute (`*A*`, `*perm*`)`

Return the generalized transpose for an N-D array object *A*.

The permutation vector *perm* must contain the elements `1:ndims (A)` (in any order, but each element must appear only once). The *N*th dimension of *A* gets remapped to dimension `PERM(N)`. For example:

```
x = zeros ([2, 3, 5, 7]);
size (x)
   ⇒  2   3   5   7

size (permute (x, [2, 1, 3, 4]))
   ⇒  3   2   5   7

size (permute (x, [1, 3, 4, 2]))
   ⇒  2   5   7   3

## The identity permutation
size (permute (x, [1, 2, 3, 4]))
   ⇒  2   3   5   7
```

**See also:** [ipermute], page 573.

`A = ipermute (B, iperm)`

The inverse of the `permute` function.

The expression

```
ipermute (permute (A, perm), perm)
```

returns the original array $A$.

**See also:** [permute], page 572.

`B = reshape (A, m, n, ...)`
`B = reshape (A, [m n ...])`
`B = reshape (A, ..., [], ...)`
`B = reshape (A, size)`

Return a matrix with the specified dimensions ($m$, $n$, ...) whose elements are taken from the matrix $A$.

The elements of the matrix are accessed in column-major order (like Fortran arrays are stored).

The following code demonstrates reshaping a 1x4 row vector into a 2x2 square matrix.

```
reshape ([1, 2, 3, 4], 2, 2)
      ⇒  1  3
         2  4
```

Note that the total number of elements in the original matrix (`prod (size (A))`) must match the total number of elements in the new matrix (`prod ([m n ...])`).

A single dimension of the return matrix may be left unspecified and Octave will determine its size automatically. An empty matrix ([]) is used to flag the unspecified dimension.

**See also:** [resize], page 573, [vec], page 579, [postpad], page 579, [cat], page 571, [squeeze], page 49.

`B = resize (A, m)`
`B = resize (A, m, n, ...)`

*B* = resize (*A*, [*m n* ...])

Resize *A* cutting off elements as necessary.

In the result, element with certain indices is equal to the corresponding element of *A* if the indices are within the bounds of *A*; otherwise, the element is set to zero.

In other words, the statement

```
B = resize (A, dv)
```

is equivalent to the following code:

```
B = zeros (dv, class (A));
sz = min (dv, size (A));
for i = 1:length (sz)
  idx{i} = 1:sz(i);
endfor
B(idx{:}) = A(idx{:});
```

but is performed more efficiently.

If only *m* is supplied, and it is a scalar, the dimension of the result is *m*-by-*m*. If *m*, *n*, ... are all scalars, then the dimensions of the result are *m*-by-*n*-by-... If given a vector as input, then the dimensions of the result are given by the elements of that vector.

An object can be resized to more dimensions than it has; in such case the missing dimensions are assumed to be 1. Resizing an object to fewer dimensions is not possible.

**See also:** [reshape], page 573, [postpad], page 579, [prepad], page 579, [cat], page 571.

*y* = circshift (*x*, *n*)
*y* = circshift (*x*, *n*, *dim*)

Circularly shift the values of the array *x*.

*n* must be a vector of integers no longer than the number of dimensions in *x*. The values of *n* can be either positive or negative, which determines the direction in which the values of *x* are shifted. If an element of *n* is zero, then the corresponding dimension of *x* will not be shifted. If *n* is a scalar and no *dim* is specified then the shift is applied to the first non-singular dimension.

If a scalar *dim* is given then operate along the specified dimension. In this case *n* must be a scalar as well.

Examples:

```
x = [1, 2, 3;
     4, 5, 6;
     7, 8, 9];
## positive shift on rows (1st non-singular dim)
circshift (x, 1)
  ⇒
      7   8   9
      1   2   3
      4   5   6
## negative shift on rows (1st non-singular dim)
circshift (x, -2)
```

```
  ⇒
       7   8   9
       1   2   3
       4   5   6
## no shift of rows, shift columns by 1 (2nd dimension)
circshift (x, [0,1])
  ⇒
       3   1   2
       6   4   5
       9   7   8
## shift columns (2nd dimension)
circshift (x, 1, 2)
  ⇒
       3   1   2
       6   4   5
       9   7   8
```

**See also:** [permute], page 572, [ipermute], page 573, [shiftdim], page 575.

`y = shiftdim (x, n)`
`[y, ns] = shiftdim (x)`

Shift the dimensions of $x$ by $n$, where $n$ must be an integer scalar.

When $n$ is positive, the dimensions of $x$ are shifted to the left, with the leading dimensions circulated to the end. If $n$ is negative, then the dimensions of $x$ are shifted to the right, with $n$ leading singleton dimensions added.

Called with a single argument, `shiftdim`, removes the leading singleton dimensions, returning the number of dimensions removed in the second output argument *ns*.

For example:

```
x = ones (1, 2, 3);
size (shiftdim (x, -1))
  ⇒   1   1   2   3
size (shiftdim (x, 1))
  ⇒   2   3
[b, ns] = shiftdim (x)
  ⇒ b =
       1   1   1
       1   1   1
  ⇒ ns = 1
```

**See also:** [reshape], page 573, [permute], page 572, [ipermute], page 573, [circshift], page 574, [squeeze], page 49.

`[s, i] = sort (x)`
`[s, i] = sort (x, dim)`
`[s, i] = sort (x, mode)`
`[s, i] = sort (x, dim, mode)`

Return a copy of $x$ with the elements arranged in increasing order.

For matrices, `sort` orders the elements within columns

For example:

```
sort ([1, 2; 2, 3; 3, 1])
   ⇒  1  1
      2  2
      3  3
```

If the optional argument *dim* is given, then the matrix is sorted along the dimension defined by *dim*. The optional argument *mode* defines the order in which the values will be sorted. Valid values of *mode* are `"ascend"` or `"descend"`.

The `sort` function may also be used to produce a matrix containing the original row indices of the elements in the sorted matrix. For example:

```
[s, i] = sort ([1, 2; 2, 3; 3, 1])
  ⇒ s = 1  1
        2  2
        3  3
  ⇒ i = 1  3
        2  1
        3  2
```

For equal elements, the indices are such that equal elements are listed in the order in which they appeared in the original list.

Sorting of complex entries is done first by magnitude (`abs (z)`) and for any ties by phase angle (`angle (z)`). For example:

```
sort ([1+i; 1; 1-i])
    ⇒ 1 + 0i
      1 - 1i
      1 + 1i
```

NaN values are treated as being greater than any other value and are sorted to the end of the list.

The `sort` function may also be used to sort strings and cell arrays of strings, in which case ASCII dictionary order (uppercase 'A' precedes lowercase 'a') of the strings is used.

The algorithm used in `sort` is optimized for the sorting of partially ordered lists.

**See also:** [sortrows], page 576, [issorted], page 577.

`[s, i] = sortrows (A)`
`[s, i] = sortrows (A, c)`

Sort the rows of the matrix $A$ according to the order of the columns specified in $c$.

By default ($c$ omitted, or a particular column unspecified in $c$) an ascending sort order is used. However, if elements of $c$ are negative then the corresponding column is sorted in descending order. If the elements of $A$ are strings then a lexicographical sort is used.

Example: sort by column 2 in descending order, then 3 in ascending order

```
x = [ 7, 1, 4;
      8, 3, 5;
      9, 3, 6 ];
sortrows (x, [-2, 3])
   ⇒ 8  3  5
     9  3  6
     7  1  4
```

**See also:** [sort], page 575.

*tf* = issorted (*A*)
*tf* = issorted (*A*, *mode*)
*tf* = issorted (*A*, "*rows*", *mode*)

Return true if the vector *A* is sorted according to *mode*, which may be either `"ascend"`, `"descend"`, `"either"`, or `"monotonic"` (`"either"` and `"monotonic"` are equivalent).

By default, *mode* is `"ascend"`. NaNs are treated in the same manner as `sort`.

If the optional argument `"rows"` is supplied, check whether the matrix is sorted by rows as output by the function `sortrows` (with no options). *Note:* the `"rows"` argument can not be used with sparse matrices.

**See also:** [sort], page 575, [sortrows], page 576.

*nel* = nth_element (*x*, *n*)
*nel* = nth_element (*x*, *n*, *dim*)

Select the n-th smallest element of a vector, using the ordering defined by `sort`.

The result is equivalent to `sort(x)(n)`.

*n* can also be a contiguous range, either ascending `l:u` or descending `u:-1:l`, in which case a range of elements is returned.

If *x* is an array, `nth_element` operates along the dimension defined by *dim*, or the first non-singleton dimension if *dim* is not given.

Programming Note: nth_element encapsulates the C++ standard library algorithms nth_element and partial_sort. On average, the complexity of the operation is O(M*log(K)), where `M = size (x, dim)` and `K = length (n)`. This function is intended for cases where the ratio K/M is small; otherwise, it may be better to use `sort`.

**See also:** [sort], page 575, [min], page 618, [max], page 617.

*A_LO* = tril (*A*)
*A_LO* = tril (*A*, *k*)
*A_LO* = tril (*A*, *k*, *pack*)

Return a new matrix formed by extracting the lower triangular part of the matrix *A*, and setting all other elements to zero.

The optional second argument specifies how many diagonals above or below the main diagonal should also be set to zero. The default value of *k* is zero which includes the main diagonal as part of the result. If the value of *k* is a nonzero integer then the selection of elements starts at an offset of *k* diagonals above the main diagonal for

positive $k$ or below the main diagonal for negative $k$. The absolute value of $k$ may not be greater than the number of subdiagonals or superdiagonals.

Example 1 : exclude main diagonal

```
tril (ones (3), -1)
     ⇒  0  0  0
        1  0  0
        1  1  0
```

Example 2 : include first superdiagonal

```
tril (ones (3), 1)
     ⇒  1  1  0
        1  1  1
        1  1  1
```

If the optional third argument `"pack"` is given then the extracted elements are not inserted into a matrix, but instead stacked column-wise one above another, and returned as a column vector.

**See also:** [triu], page 578, [istril], page 72, [diag], page 579.

*A_UP* = **triu** (*A*)
*A_UP* = **triu** (*A*, *k*)
*A_UP* = **triu** (*A*, *k*, *pack*)

Return a new matrix formed by extracting the upper triangular part of the matrix $A$, and setting all other elements to zero.

The optional second argument specifies how many diagonals above or below the main diagonal should also be set to zero. The default value of $k$ is zero which includes the main diagonal as part of the result. If the value of $k$ is a nonzero integer then the selection of elements starts at an offset of $k$ diagonals above the main diagonal for positive $k$ or below the main diagonal for negative $k$. The absolute value of $k$ may not be greater than the number of subdiagonals or superdiagonals.

Example 1 : exclude main diagonal

```
triu (ones (3), 1)
     ⇒  0  1  1
        0  0  1
        0  0  0
```

Example 2 : include first subdiagonal

```
triu (ones (3), -1)
     ⇒  1  1  1
        1  1  1
        0  1  1
```

If the optional third argument `"pack"` is given then the extracted elements are not inserted into a matrix, but instead stacked column-wise one above another, and returned as a column vector.

**See also:** [tril], page 577, [istriu], page 72, [diag], page 579.

`v = vec (x)`
`v = vec (x, dim)`
: Return the vector obtained by stacking the columns of the matrix *x* one above the other.

  Without *dim* this is equivalent to `x(:)`.

  If *dim* is supplied, the dimensions of *v* are set to *dim* with all elements along the last dimension. This is equivalent to `shiftdim (x(:), 1-dim)`.

  **See also:** [vech], page 579, [resize], page 573, [cat], page 571.

`v = vech (x)`
: Return the vector obtained by eliminating all superdiagonal elements of the square matrix *x* and stacking the result one column above the other.

  This has uses in matrix calculus where the underlying matrix is symmetric and it would be pointless to keep values above the main diagonal.

  **See also:** [vec], page 579.

`B = prepad (A, l)`
`B = prepad (A, l, c)`
`B = prepad (A, l, c, dim)`
: Prepend the scalar value *c* to the vector *A* until it is of length *l*. If *c* is not given, a value of 0 is used.

  If `length (A) > l`, elements from the beginning of *A* are removed until a vector of length *l* is obtained.

  If *A* is a matrix, elements are prepended or removed from each row.

  If the optional argument *dim* is given, operate along this dimension.

  If *dim* is larger than the dimensions of *A*, the result will have *dim* dimensions.

  **See also:** [postpad], page 579, [cat], page 571, [resize], page 573.

`B = postpad (A, l)`
`B = postpad (A, l, c)`
`B = postpad (A, l, c, dim)`
: Append the scalar value *c* to the vector *A* until it is of length *l*. If *c* is not given, a value of 0 is used.

  If `length (A) > l`, elements from the end of *A* are removed until a vector of length *l* is obtained.

  If *A* is a matrix, elements are appended or removed from each row.

  If the optional argument *dim* is given, operate along this dimension.

  If *dim* is larger than the dimensions of *A*, the result will have *dim* dimensions.

  **See also:** [prepad], page 579, [cat], page 571, [resize], page 573.

`M = diag (v)`
`M = diag (v, k)`
`M = diag (v, m, n)`
`v = diag (M)`

`v = diag (M, k)`

Return a diagonal matrix with vector $v$ on diagonal $k$.

The second argument is optional. If it is positive, the vector is placed on the $k$-th superdiagonal. If it is negative, it is placed on the -$k$-th subdiagonal. The default value of $k$ is 0, and the vector is placed on the main diagonal. For example:

```
diag ([1, 2, 3], 1)
   ⇒  0  1  0  0
      0  0  2  0
      0  0  0  3
      0  0  0  0
```

The 3-input form returns a diagonal matrix with vector $v$ on the main diagonal and the resulting matrix being of size $m$ rows x $n$ columns.

Given a matrix argument, instead of a vector, `diag` extracts the $k$-th diagonal of the matrix.

`M = blkdiag (A, B, C, ...)`

Build a block diagonal matrix from $A$, $B$, $C$, ...

All arguments must be numeric and either two-dimensional matrices or scalars. If any argument is of type sparse, the output will also be sparse.

**See also:** [diag], page 579, [horzcat], page 572, [vertcat], page 572, [sparse], page 732.

## 16.3 Special Utility Matrices

`I = eye ()`
`I = eye (n)`
`I = eye (m, n)`
`I = eye ([m, n])`
`I = eye (..., class)`

Return an identity matrix.

If called with no arguments, return the scalar value `1`.

If invoked with a single scalar argument $n$, return a square NxN identity matrix.

If supplied two scalar arguments ($m$, $n$), or a 2-element vector `[m, n]`, return an MxN identity matrix with $m$ rows and $n$ columns.

The optional argument *class* specifies the return type of the matrix and defaults to `"double"`.

Example 1 : 1-input, square identity matrix

```
eye (3)
 ⇒  1  0  0
    0  1  0
    0  0  1
```

Example 2 : following expressions all produce 2x2 identity matrix

```
eye (2) ≡ eye (2, 2) ≡ eye (size ([1, 2; 3, 4]))
 ⇒  1  0
    0  1
```

Example 3 : 2x2 uint8 identity matrix

```
I = eye (2, "uint8")
```

Programming Note: Calling `eye` with no arguments is equivalent to calling it with an argument of `1`. Any negative dimensions are treated as zero. These odd definitions are for compatibility with MATLAB.

**See also:** [speye], page 730, [ones], page 581, [zeros], page 581.

`x = ones ()`
`x = ones (n)`
`x = ones (m, n, ...)`
`x = ones ([m, n, ...])`
`x = ones (..., class)`
`x = ones (..., "like", var)`

Return a scalar, matrix, or N-dimensional array whose elements are all `1`.

If called with no arguments, return the scalar value `1`.

If invoked with a single scalar integer argument *n*, return a square NxN matrix.

If invoked with two or more scalar integer arguments, or a vector of integer values, return an array with the given dimensions.

The optional argument *class* specifies the class of the return array and defaults to `"double"`.

If a variable *var* is specified after `"like"`, the output *val* will have the same data type, complexity, and sparsity as *var*.

Example 1 : MxN matrix of constant value *val*

```
C = val * ones (m, n)
```

Example 2 : MxN matrix of uint8

```
C = ones (m, n, "uint8")
```

Programming Note: Any negative dimensions are treated as zero, and any zero dimensions will result in an empty matrix. This odd behavior is for MATLAB compatibility.

**See also:** [zeros], page 581, [true], page 66, [false], page 67.

`x = zeros ()`
`x = zeros (n)`
`x = zeros (m, n, ...)`
`x = zeros ([m, n, ...])`
`x = zeros (..., class)`
`x = zeros (..., "like", var)`

Return a scalar, matrix, or N-dimensional array whose elements are all `0`.

If called with no arguments, return the scalar value `0`.

If invoked with a single scalar integer argument *n*, return a square NxN matrix.

If invoked with two or more scalar integer arguments, or a vector of integer values, return an array with the given dimensions.

The optional argument *class* specifies the class of the return array and defaults to `"double"`.

If a variable *var* is specified after `"like"`, the output *val* will have the same data type, complexity, and sparsity as *var*.

Example : MxN matrix of uint8

```
C = ones (m, n, "uint8")
```

Programming Note: Any negative dimensions are treated as zero, and any zero dimensions will result in an empty matrix. This odd behavior is for MATLAB compatibility.

**See also:** [ones], page 581, [true], page 66, [false], page 67.

`B = repmat (A, m)`
`B = repmat (A, m, n)`
`B = repmat (A, m, n, p ...)`
`B = repmat (A, [m n])`
`B = repmat (A, [m n p ...])`

Repeat matrix or N-D array.

Form a block matrix of size $m$ by $n$, with a copy of matrix $A$ as each element.

If $n$ is not specified, form an $m$ by $m$ block matrix. For copying along more than two dimensions, specify the number of times to copy across each dimension $m$, $n$, $p$, ..., in a vector in the second argument.

**See also:** [bsxfun], page 692, [kron], page 674, [repelems], page 582.

`y = repelems (x, r)`

Construct a vector of repeated elements from $x$.

$r$ is a 2x$N$ integer matrix specifying which elements to repeat and how often to repeat each element. Entries in the first row, $r$(1,j), select an element to repeat. The corresponding entry in the second row, $r$(2,j), specifies the repeat count. If $x$ is a matrix then the columns of $x$ are imagined to be stacked on top of each other for purposes of the selection index. A row vector is always returned.

Conceptually the result is calculated as follows:

```
y = [];
for i = 1:columns (r)
  y = [y, x(r(1,i)*ones(1, r(2,i)))];
endfor
```

**See also:** [repmat], page 582, [cat], page 571.

`xxx = repelem (x, R)`
`xxx = repelem (x, R_1, ..., R_n)`

Construct an array of repeated elements from $x$ and repeat instructions *R_1*, ...

$x$ must be a scalar, vector, or N-dimensional array.

A repeat instruction *R_j* must either be a scalar or a vector. If the instruction is a scalar then each component of $x$ in dimension $j$ is repeated *R_j* times. If the instruction is a vector then it must have the same number of elements as the corresponding dimension $j$ of $x$. In this case, the $k$th component of dimension $j$ is repeated `R_j(k)` times.

If $x$ is a scalar or vector then `repelem` may be called with just a single repeat instruction $R$ and `repelem` will return a vector with the same orientation as the input.

If $x$ is a matrix then at least two $R_j$s must be specified.

Note: Using `repelem` with a vector $x$ and a vector for $R_j$ is equivalent to Run Length Decoding.

Examples:

```
A = [1 2 3 4 5];
B = [2 1 0 1 2];
repelem (A, B)
  ⇒   1   1   2   4   5   5
```

```
A = magic (3)
  ⇒ A =
       8   1   6
       3   5   7
       4   9   2
B1 = [1 2 3];
B2 = 2;
repelem (A, B1, B2)
  ⇒     8   8   1   1   6   6
        3   3   5   5   7   7
        3   3   5   5   7   7
        4   4   9   9   2   2
        4   4   9   9   2   2
        4   4   9   9   2   2
```

More $R_j$ may be specified than the number of dimensions of $x$. Any excess $R_j$ must be scalars (because $x$'s size in those dimensions is only 1), and $x$ will be replicated in those dimensions accordingly.

```
A = [1 2 3 4 5];
B1 = 2;
B2 = [2 1 3 0 2];
B3 = 3;
repelem (A, B1, B2, B3)
  ⇒    ans(:,:,1) =
          1   1   2   3   3   3   5   5
          1   1   2   3   3   3   5   5

       ans(:,:,2) =

          1   1   2   3   3   3   5   5
          1   1   2   3   3   3   5   5

       ans(:,:,3) =
          1   1   2   3   3   3   5   5
          1   1   2   3   3   3   5   5
```

$R_j$ must be specified in order. A placeholder of 1 may be used for dimensions which do not need replication.

```
repelem ([-1, 0; 0, 1], 1, 2, 1, 2)
  ⇒  ans(:,:,1,1) =
       -1  -1   0   0
        0   0   1   1

     ans(:,:,1,2) =
       -1  -1   0   0
        0   0   1   1
```

If fewer *R_j* are given than the number of dimensions in *x*, `repelem` will assume *R_j* is 1 for those dimensions.

```
A = cat (3, [-1 0; 0 1], [-1 0; 0 1])
  ⇒  ans(:,:,1) =
       -1   0
        0   1

     ans(:,:,2) =
       -1   0
        0   1

repelem (A,2,3)
  ⇒  ans(:,:,1) =
       -1  -1  -1   0   0   0
       -1  -1  -1   0   0   0
        0   0   0   1   1   1
        0   0   0   1   1   1

     ans(:,:,2) =
       -1  -1  -1   0   0   0
       -1  -1  -1   0   0   0
        0   0   0   1   1   1
        0   0   0   1   1   1
```

`repelem` preserves the class of *x*, and works with strings, cell arrays, NA, and NAN inputs. If any *R_j* is 0 the output will be an empty array.

```
repelem ("Octave", 2, 3)
  ⇒    OOOcccttttaaavvveee
       OOOcccttttaaavvveee

repelem ([1 2 3; 1 2 3], 2, 0)
  ⇒    [](4x0)
```

**See also:** [cat], page 571, [kron], page 674, [repmat], page 582.

The functions `linspace` and `logspace` make it very easy to create vectors with evenly or logarithmically spaced elements. See Section 4.2 [Ranges], page 56.

`y = linspace (`*start*`, `*end*`)`
`y = linspace (`*start*`, `*end*`, `*n*`)`

Return a row vector with *n* linearly spaced elements between *start* and *end*.

If the number of elements $n$ is greater than one, then the endpoints *start* and *end* are always included in the range. If *start* is greater than *end*, the elements are stored in decreasing order. If the number of points $n$ is not specified, a value of 100 is used.

The `linspace` function returns a row vector when both *start* and *end* are scalars. If one, or both, inputs are vectors, then `linspace` transforms them to column vectors and returns a matrix where each row is an independent sequence between `start(row_n), end(row_n)`.

Programming Notes: For compatibility with MATLAB, return the second argument (*end*) when a single value ($n = 1$) is requested. If $n$ is not an integer then `floor (n)` is used to round the number of elements. If $n$ is zero or negative then an empty 1x0 matrix is returned.

**See also:** [colon], page 982, [logspace], page 585.

`y = logspace (a, b)`
`y = logspace (a, b, n)`
`y = logspace (a, pi)`
`y = logspace (a, pi, n)`

Return a row vector with $n$ elements logarithmically spaced from $10^a$ to $10^b$.

If the number of elements $n$ is unspecified it defaults to 50.

If $b$ is equal to $\pi$, the points are between $10^a$ and $\pi$, *not* $10^a$ and $10^\pi$, which is useful in digital signal processing.

Programming Notes: For compatibility with MATLAB, return the right-hand side of the range ($10^b$) when a single value ($n = 1$) is requested. If $n$ is not an integer then `floor (n)` is used to round the number of elements. If $n$ is zero or negative then an empty 1x0 matrix is returned.

**See also:** [linspace], page 584.

`x = rand ()`
`x = rand (n)`
`x = rand (m, n, ...)`
`x = rand ([m, n, ...])`
`x = rand (..., class)`
`v = rand ("state")`
`rand ("state", v)`
`rand ("state", "reset")`
`v = rand ("seed")`
`rand ("seed", v)`
`rand ("seed", "reset")`

Return a scalar, matrix, or N-dimensional array whose elements are random numbers uniformly distributed on the interval (0, 1).

If called with no arguments, return a scalar random value.

If invoked with a single scalar integer argument $n$, return a square NxN matrix.

If invoked with two or more scalar integer arguments, or a vector of integer values, return an array with the given dimensions.

The optional argument *class* specifies the class of the return array. The only valid options are `"double"` (default) or `"single"`.

Programming Note: Any negative dimensions are treated as zero, and any zero dimensions will result in an empty matrix. This odd behavior is for MATLAB compatibility.

The additional calling forms provide an interface to the underlying random number generator.

You can query the state of the random number generator using the form

```
v = rand ("state")
```

This returns a column vector $v$ of length 625. Later, you can restore the random number generator to the state $v$ using the form

```
rand ("state", v)
```

You may also initialize the state vector from an arbitrary vector of length $\leq 625$ for $v$. The new state will be a hash based on the value of $v$, not $v$ itself.

By default, the generator is initialized by contributing entropy from the wall clock time, the CPU time, the current fraction of a second, the process ID and—if available—up to 1024 bits from the C++ random number source `random_device`, which might be non-deterministic (implementation specific). Note that this differs from MATLAB, which always initializes the random number generator to the same state at startup. To obtain behavior comparable to MATLAB, initialize with a deterministic state vector in Octave's startup files (see Section 2.1.2 [Startup Files], page 19).

Programming Notes: To compute the pseudo-random sequence, `rand` uses the Mersenne Twister with a period of $2^{19937} - 1$ (See M. Matsumoto and T. Nishimura, "Mersenne Twister: A 623-dimensionally equidistributed uniform pseudorandom number generator", *ACM Trans. on Modeling and Computer Simulation*, Vol. 8, No. 1, pp. 3–30, January 1998, http://www.math.sci.hiroshima-u.ac.jp/~m-mat/MT/emt.html). Do **not** use for cryptography without securely hashing several returned values together, otherwise the generator state can be learned after reading 624 consecutive values.

Older versions of Octave used a different random number generator. The new generator is used by default as it is significantly faster than the old generator, and produces random numbers with a significantly longer cycle time. However, in some circumstances it might be desirable to obtain the same random sequences as produced by the old generators. To do this the keyword `"seed"` is used to specify that the old generators should be used, as in

```
rand ("seed", val)
```

which sets the seed of the generator to *val*. The seed of the generator can be queried with

```
s = rand ("seed")
```

However, it should be noted that querying the seed will not cause `rand` to use the old generators, only setting the seed will. To cause `rand` to once again use the new generators, the keyword `"state"` must be used to reset `rand` back to using the default generator.

The state or seed of the generator can be reset to a new random value using the `"reset"` keyword.

**See also:** [randi], page 587, [randn], page 587, [rande], page 588, [randg], page 590, [randp], page 589.

```
R = randi (imax)
R = randi (imax, n)
R = randi (imax, m, n, ...)
R = randi (imax, [m, n, ...])
R = randi ([imin, imax], ...)
R = randi (..., "class")
```

Return a scalar, matrix, or N-dimensional array whose elements are random integers in the range [1, *imax*].

If called with no size arguments, return a scalar random value.

If invoked with a single scalar size argument *n*, return a square NxN matrix.

If invoked with two or more scalar integer size arguments, or a vector of integer values, return an array with the given dimensions.

The integer range may optionally be described by a two-element matrix with a lower and upper bound in which case the returned integers will be on the interval [*imin*, *imax*].

The optional argument *class* will return a matrix of the requested type. The default is `"double"`. Supported classes are: `"double"`, `"single"`, `"int8"`, `"int16"`, `"int32"`, `"uint8"`, `"uint16"`, `"uint32"`, and `"logical"`.

Programming Notes: `randi` relies internally on `rand` which uses class `"double"` to represent numbers. This limits the maximum integer (*imax*) and range (*imax* - *imin*) to the value returned by the `flintmax` function. For IEEE 754 floating point numbers this value is $2^{53} - 1$.

When the output class is `"logical"` the function constructs an array of random integers as specified and then applies the test `x > 0` to determine which elements will be true. Because the one-input form of the function uses `1` for *imin* `randi` will always return a matrix of all ones.

Example: 150 integers in the range 1–10.

```
ri = randi (10, 150, 1)
```

**See also:** [rand], page 585, [randn], page 587, [rande], page 588, [randg], page 590, [randp], page 589.

```
x = randn ()
x = randn (n)
x = randn (m, n, ...)
x = randn ([m, n, ...])
x = randn (..., class)
v = randn ("state")
randn ("state", v)
randn ("state", "reset")
v = randn ("seed")
```

`randn` ("*seed*", *v*)
`randn` ("*seed*", "*reset*")

Return a scalar, matrix, or N-dimensional array whose elements are random numbers from the standard normal distribution having a mean of `0` and a variance of `1`.

If called with no arguments, return a scalar random value.

If invoked with a single scalar integer argument *n*, return a square NxN matrix.

If invoked with two or more scalar integer arguments, or a vector of integer values, return an array with the given dimensions.

The optional argument *class* specifies the class of the return array. The only valid options are `"double"` (default) or `"single"`.

Programming Note: Any negative dimensions are treated as zero, and any zero dimensions will result in an empty matrix. This odd behavior is for MATLAB compatibility.

The additional calling forms provide an interface to the underlying random number generator. See [`rand`], page 585, for documentation on querying and controlling the random number generator.

Programming Note: By default, `randn` uses the Marsaglia and Tsang "Ziggurat technique" to transform from a uniform to a normal distribution.

Reference: G. Marsaglia and W.W. Tsang, "Ziggurat Method for Generating Random Variables", *J. Statistical Software*, Vol. 5, 2000, `https://www.jstatsoft.org/v05/i08/`

**See also:** [rand], page 585, [randi], page 587, [rande], page 588, [randg], page 590, [randp], page 589.

*x* = `rande` ()
*x* = `rande` (*n*)
*x* = `rande` (*m*, *n*, ...)
*x* = `rande` ([*m*, *n*, ...])
*x* = `rande` (..., *class*)
*v* = `rande` ("*state*")
`rande` ("*state*", *v*)
`rande` ("*state*", "*reset*")
*v* = `rande` ("*seed*")
`rande` ("*seed*", *v*)
`rande` ("*seed*", "*reset*")

Return a scalar, matrix, or N-dimensional array whose elements are random numbers from the exponential distribution with rate parameter `0`.

If called with no arguments, return a scalar random value.

If invoked with a single scalar integer argument *n*, return a square NxN matrix.

If invoked with two or more scalar integer arguments, or a vector of integer values, return an array with the given dimensions.

The optional argument *class* specifies the class of the return array. The only valid options are `"double"` (default) or `"single"`.

Programming Note: Any negative dimensions are treated as zero, and any zero dimensions will result in an empty matrix. This odd behavior is for MATLAB compatibility.

The additional calling forms provide an interface to the underlying random number generator. See [rand], page 585, for documentation on querying and controlling the random number generator.

Programming Note: By default, `rande` uses the Marsaglia and Tsang "Ziggurat technique" to transform from a uniform to an exponential distribution.

Reference: G. Marsaglia and W.W. Tsang, "Ziggurat Method for Generating Random Variables", *J. Statistical Software*, Vol 5, 2000, https://www.jstatsoft.org/v05/i08/

**See also:** [rand], page 585, [randi], page 587, [randn], page 587, [randg], page 590, [randp], page 589.

```
x = randp (l)
x = randp (l, n)
x = randp (l, m, n, ...)
x = randp (l, [m, n, ...])
x = randp (..., class)
v = randp ("state")
randp ("state", v)
randp ("state", "reset")
v = randp ("seed")
randp ("seed", v)
randp ("seed", "reset")
```

Return a scalar, matrix, or N-dimensional array whose elements are random numbers from the Poisson distribution with mean value parameter *l*.

If called with no size arguments, return a scalar random value.

If invoked with a single scalar size argument *n*, return a square NxN matrix.

If invoked with two or more scalar integer size arguments, or a vector of integer values, return an array with the given dimensions.

The optional argument *class* specifies the class of the return array. The only valid options are `"double"` (default) or `"single"`.

Programming Note: Any negative dimensions are treated as zero, and any zero dimensions will result in an empty matrix. This odd behavior is for MATLAB compatibility.

The additional calling forms provide an interface to the underlying random number generator. See [rand], page 585, for documentation on querying and controlling the random number generator.

Programming Notes: Five different algorithms are used depending on the range of *l* and whether *l* is a scalar or a matrix.

For scalar $l \leq 12$, use direct method.
W.H. Press, et al., *Numerical Recipes in C*, Cambridge University Press, 1992.

For scalar $l > 12$, use rejection method.[1]
W.H. Press, et al., *Numerical Recipes in C*, Cambridge University Press, 1992.

For matrix $l \leq 10$, use inversion method.[2]
E. Stadlober, et al., WinRand source code, available via FTP.

For matrix $l > 10$, use patchwork rejection method.
E. Stadlober, et al., WinRand source code, available via FTP, or H. Zechner, *Efficient sampling from continuous and discrete unimodal distributions*, Doctoral Dissertation, pp. 156, Technical University Graz, Austria, 1994.

For $l >$ 1e8, use normal approximation.
L. Montanet, et al., "Review of Particle Properties", *Physical Review D*, 50, p. 1284, 1994.

**See also:** [rand], page 585, [randi], page 587, [randn], page 587, [rande], page 588, [randg], page 590.

```
x = randg (a)
x = randg (a, n)
x = randg (a, m, n, ...)
x = randg (a, [m, n, ...])
x = randg (..., class)
v = randg ("state")
randg ("state", v)
randg ("state", "reset")
v = randg ("seed")
randg ("seed", v)
randg ("seed", "reset")
```

Return a scalar, matrix, or N-dimensional array whose elements are random numbers from the gamma distribution `gamma (a,1)`.

If called with no size arguments, return a scalar random value.

If invoked with a single scalar size argument $n$, return a square NxN matrix.

If invoked with two or more scalar integer size arguments, or a vector of integer values, return an array with the given dimensions.

The optional argument *class* specifies the class of the return array. The only valid options are `"double"` (default) or `"single"`.

Programming Note: Any negative dimensions are treated as zero, and any zero dimensions will result in an empty matrix. This odd behavior is for MATLAB compatibility.

The additional calling forms provide an interface to the underlying random number generator. See [rand], page 585, for documentation on querying and controlling the random number generator.

Programming Notes: The gamma distribution can be used to generate many other distributions:

`gamma (a, b)` for `a > -1`, `b > 0`

```
r = b * randg (a)
```

`beta (a, b)` for `a > -1`, `b > -1`

```
r1 = randg (a, 1)
r = r1 / (r1 + randg (b, 1))
```

`Erlang (a, n)`

```
r = a * randg (n)
```

`chisq (df)` for `df > 0`

```
r = 2 * randg (df / 2)
```

`t (df)` for `0 < df < inf` (use randn if df is infinite)

```
r = randn () / sqrt (2 * randg (df / 2) / df)
```

`F (n1, n2)` for `0 < n1`, `0 < n2`

```
## r1 equals 1 if n1 is infinite
r1 = 2 * randg (n1 / 2) / n1
## r2 equals 1 if n2 is infinite
r2 = 2 * randg (n2 / 2) / n2
r = r1 / r2
```

negative `binomial (n, p)` for `n > 0`, `0 < p <= 1`

```
r = randp ((1 - p) / p * randg (n))
```

non-central `chisq (df, L)`, for `df >= 0` and `L > 0`
(use chisq if `L = 0`)

```
r = randp (L / 2)
r(r > 0) = 2 * randg (r(r > 0))
r(df > 0) += 2 * randg (df(df > 0)/2)
```

`Dirichlet (a1, ... ak)`

```
r = (randg (a1), ..., randg (ak))
r = r / sum (r)
```

**See also:** [rand], page 585, [randi], page 587, [randn], page 587, [rande], page 588, [randp], page 589.

`rng (`*seed*`)`
`rng (`*seed*`, "`*generator*`")`
`rng ("`*shuffle*`")`
`rng ("`*shuffle*`", "`*generator*`")`
`rng ("`*default*`")`
*s* `= rng ()`
`rng (`*s*`)`
*s* `= rng (...)`

Set or query the seed of the random number generator used by `rand` and `randn`.

The input *seed* is a scalar numeric value used to initialize the state vector of the random number generator.

The optional string *generator* specifies the type of random number generator to be used. Its value can be `"twister"`, `"v5uniform"`, or `"v5normal"`. The `"twister"` keyword is described below. `"v5uniform"` and `"v5normal"` refer to older versions of Octave that used to use a different random number generator.

The state or seed of the random number generator can be reset to a new random value using the `"shuffle"` keyword.

The random number generator can be reset to default values using the `"default"` keyword. The default values are to use the Mersenne Twister generator with a seed of 0.

The optional return value *s* contains the state of the random number generator at the time the function is called (i.e., before it might be modified according to the input arguments). It is encoded as a structure variable with three fields: `"Type"`, `"Seed"`, and `"State"`. The random number generator can be restored to the state *s* using `rng (s)`. This is useful when the identical sequence of pseudo-random numbers is required for an algorithm.

By default, and with the `"twister"` option, pseudo-random sequences are computed using the Mersenne Twister with a period of $2^{19937} - 1$ (See M. Matsumoto and T. Nishimura, "Mersenne Twister: A 623-dimensionally equidistributed uniform pseudorandom number generator", *ACM Trans. on Modeling and Computer Simulation*, Vol. 8, No. 1, pp. 3–30, January 1998, http://www.math.sci.hiroshima-u.ac.jp/~m-mat/MT/emt.html). Do **not** use for cryptography without securely hashing several returned values together, otherwise the generator state can be learned after reading 624 consecutive values.

**See also:** [rand], page 585, [randn], page 587.

The generators operate in the new or old style together, it is not possible to mix the two. Initializing any generator with `"state"` or `"seed"` causes the others to switch to the same style for future calls.

The state of each generator is independent and calls to different generators can be interleaved without affecting the final result. For example,

```
rand ("state", [11, 22, 33]);
randn ("state", [44, 55, 66]);
u = rand (100, 1);
n = randn (100, 1);
```

and

```
rand ("state", [11, 22, 33]);
randn ("state", [44, 55, 66]);
u = zeros (100, 1);
n = zeros (100, 1);
for i = 1:100
  u(i) = rand ();
  n(i) = randn ();
end
```

produce equivalent results. When the generators are initialized in the old style with `"seed"` only `rand` and `randn` are independent, because the old `rande`, `randg` and `randp` generators make calls to `rand` and `randn`.

The generators are initialized with random states at start-up, so that the sequences of random numbers are not the same each time you run Octave.[1] If you really do need to reproduce a sequence of numbers exactly, you can set the state or seed to a specific value.

[1] The old versions of `rand` and `randn` obtain their initial seeds from the system clock.

If invoked without arguments, `rand` and `randn` return a single element of a random sequence.

The original `rand` and `randn` functions use Fortran code from RANLIB, a library of Fortran routines for random number generation, compiled by Barry W. Brown and James Lovato of the Department of Biomathematics at The University of Texas, M.D. Anderson Cancer Center, Houston, TX 77030.

`v = randperm (n)`
`v = randperm (n, m)`
Return a row vector containing a random permutation of `1:n`.

If $m$ is supplied, return $m$ unique entries, sampled without replacement from `1:n`.

The complexity is O($n$) in memory and O($m$) in time, unless $m < n/5$, in which case O($m$) memory is used as well. The randomization is performed using rand(). All permutations are equally likely.

**See also:** [perms], page 861.

## 16.4 Famous Matrices

The following functions return famous matrix forms.

`gallery (name)`
`gallery (name, args)`
Create interesting matrices for testing.

`c = gallery ("cauchy", x)`
`c = gallery ("cauchy", x, y)`
Create a Cauchy matrix.

`c = gallery ("chebspec", n)`
`c = gallery ("chebspec", n, k)`
Create a Chebyshev spectral differentiation matrix.

`c = gallery ("chebvand", p)`
`c = gallery ("chebvand", m, p)`
Create a Vandermonde-like matrix for the Chebyshev polynomials.

`a = gallery ("chow", n)`
`a = gallery ("chow", n, alpha)`
`a = gallery ("chow", n, alpha, delta)`
Create a Chow matrix – a singular Toeplitz lower Hessenberg matrix.

`c = gallery ("circul", v)`
Create a circulant matrix.

`a = gallery ("clement", n)`
`a = gallery ("clement", n, k)`
Create a tridiagonal matrix with zero diagonal entries.

`c = gallery ("compar", a)`
`c = gallery ("compar", a, k)`
Create a comparison matrix.

`a = gallery ("condex", n)`
`a = gallery ("condex", n, k)`
`a = gallery ("condex", n, k, theta)`
Create a "counterexample" matrix to a condition estimator.

`a = gallery ("cycol", [m n])`
`a = gallery ("cycol", n)`
`a = gallery (..., k)`
Create a matrix whose columns repeat cyclically.

`[c, d, e] = gallery ("dorr", n)`
`[c, d, e] = gallery ("dorr", n, theta)`
`a = gallery ("dorr", ...)`
Create a diagonally dominant, ill-conditioned, tridiagonal matrix.

`a = gallery ("dramadah", n)`
`a = gallery ("dramadah", n, k)`
Create a (0, 1) matrix whose inverse has large integer entries.

`a = gallery ("fiedler", c)`
Create a symmetric Fiedler matrix.

`a = gallery ("forsythe", n)`
`a = gallery ("forsythe", n, alpha)`
`a = gallery ("forsythe", n, alpha, lambda)`
Create a Forsythe matrix (a perturbed Jordan block).

`f = gallery ("frank", n)`
`f = gallery ("frank", n, k)`
Create a Frank matrix (ill-conditioned eigenvalues).

`c = gallery ("gcdmat", n)`
Create a greatest common divisor matrix.

$c$ is an $n$-by-$n$ matrix whose values correspond to the greatest common divisor of its coordinate values, i.e., $c$(i,j) correspond `gcd (i, j)`.

`a = gallery ("gearmat", n)`
`a = gallery ("gearmat", n, i)`
`a = gallery ("gearmat", n, i, j)`
Create a Gear matrix.

`g = gallery ("grcar", n)`
`g = gallery ("grcar", n, k)`
Create a Toeplitz matrix with sensitive eigenvalues.

`a = gallery ("hanowa", n)`
`a = gallery ("hanowa", n, d)`
Create a matrix whose eigenvalues lie on a vertical line in the complex plane.

`v = gallery ("house", x)`
`[v, beta] = gallery ("house", x)`
Create a householder matrix.

`a = gallery ("integerdata", imax, [M N ...], j)`
`a = gallery ("integerdata", imax, M, N, ..., j)`
`a = gallery ("integerdata", [imin, imax], [M N ...], j)`
`a = gallery ("integerdata", [imin, imax], M, N, ..., j)`
`a = gallery ("integerdata", ..., "class")`
Create a matrix with random integers in the range [1, *imax*]. If *imin* is given then the integers are in the range [*imin*, *imax*].

The second input is a matrix of dimensions describing the size of the output. The dimensions can also be input as comma-separated arguments.

The input *j* is an integer index in the range [0, 2^32-1]. The values of the output matrix are always exactly the same (reproducibility) for a given size input and *j* index.

The final optional argument determines the class of the resulting matrix. Possible values for *class*: `"uint8"`, `"uint16"`, `"uint32"`, `"int8"`, `"int16"`, int32", `"single"`, `"double"`. The default is `"double"`.

`a = gallery ("invhess", x)`
`a = gallery ("invhess", x, y)`
Create the inverse of an upper Hessenberg matrix.

`a = gallery ("invol", n)`
Create an involutory matrix.

`a = gallery ("ipjfact", n)`
`a = gallery ("ipjfact", n, k)`
Create a Hankel matrix with factorial elements.

`a = gallery ("jordbloc", n)`
`a = gallery ("jordbloc", n, lambda)`
Create a Jordan block.

`u = gallery ("kahan", n)`
`u = gallery ("kahan", n, theta)`
`u = gallery ("kahan", n, theta, pert)`
Create a Kahan matrix (upper trapezoidal).

`a = gallery ("kms", n)`
`a = gallery ("kms", n, rho)`
Create a Kac-Murdock-Szego Toeplitz matrix.

`b = gallery ("krylov", a)`
`b = gallery ("krylov", a, x)`
`b = gallery ("krylov", a, x, j)`
Create a Krylov matrix.

`a = gallery ("lauchli", n)`
`a = gallery ("lauchli", n, mu)`
Create a Lauchli matrix (rectangular).

`a = gallery ("lehmer", n)`
Create a Lehmer matrix (symmetric positive definite).

`t = gallery ("lesp", n)`
Create a tridiagonal matrix with real, sensitive eigenvalues.

`a = gallery ("lotkin", n)`
Create a Lotkin matrix.

`a = gallery ("minij", n)`
Create a symmetric positive definite matrix MIN(i,j).

`a = gallery ("moler", n)`
`a = gallery ("moler", n, alpha)`
Create a Moler matrix (symmetric positive definite).

`[a, t] = gallery ("neumann", n)`
Create a singular matrix from the discrete Neumann problem (sparse).

`a = gallery ("normaldata", [M N ...], j)`
`a = gallery ("normaldata", M, N, ..., j)`
`a = gallery ("normaldata", ..., "class")`
Create a matrix with random samples from the standard normal distribution (mean = 0, std = 1).

The first input is a matrix of dimensions describing the size of the output. The dimensions can also be input as comma-separated arguments.

The input *j* is an integer index in the range [0, 2^32-1]. The values of the output matrix are always exactly the same (reproducibility) for a given size input and *j* index.

The final optional argument determines the class of the resulting matrix. Possible values for *class*: `"single"`, `"double"`. The default is `"double"`.

`q = gallery ("orthog", n)`
`q = gallery ("orthog", n, k)`
Create orthogonal and nearly orthogonal matrices.

`a = gallery ("parter", n)`
Create a Parter matrix (a Toeplitz matrix with singular values near pi).

`p = gallery ("pei", n)`
`p = gallery ("pei", n, alpha)`
Create a Pei matrix.

`a = gallery ("poisson", n)`
Create a block tridiagonal matrix from Poisson's equation (sparse).

`a = gallery ("prolate", n)`
`a = gallery ("prolate", n, w)`
Create a prolate matrix (symmetric, ill-conditioned Toeplitz matrix).

`h = gallery ("randhess", x)`
Create a random, orthogonal upper Hessenberg matrix.

`a = gallery ("rando", n)`
`a = gallery ("rando", n, k)`
Create a random matrix with elements -1, 0 or 1.

`a = gallery ("randsvd", n)`
`a = gallery ("randsvd", n, kappa)`
`a = gallery ("randsvd", n, kappa, mode)`
`a = gallery ("randsvd", n, kappa, mode, kl)`
`a = gallery ("randsvd", n, kappa, mode, kl, ku)`
Create a random matrix with pre-assigned singular values.

`a = gallery ("redheff", n)`
Create a zero and ones matrix of Redheffer associated with the Riemann hypothesis.

`a = gallery ("riemann", n)`
Create a matrix associated with the Riemann hypothesis.

`a = gallery ("ris", n)`
Create a symmetric Hankel matrix.

`a = gallery ("smoke", n)`
`a = gallery ("smoke", n, k)`
Create a complex matrix, with a "smoke ring" pseudospectrum.

`t = gallery ("toeppd", n)`
`t = gallery ("toeppd", n, m)`
`t = gallery ("toeppd", n, m, w)`
`t = gallery ("toeppd", n, m, w, theta)`
Create a symmetric positive definite Toeplitz matrix.

`p = gallery ("toeppen", n)`
`p = gallery ("toeppen", n, a)`
`p = gallery ("toeppen", n, a, b)`
`p = gallery ("toeppen", n, a, b, c)`
`p = gallery ("toeppen", n, a, b, c, d)`
`p = gallery ("toeppen", n, a, b, c, d, e)`
Create a pentadiagonal Toeplitz matrix (sparse).

`a = gallery ("tridiag", x, y, z)`
`a = gallery ("tridiag", n)`
`a = gallery ("tridiag", n, c, d, e)`
Create a tridiagonal matrix (sparse).

`t = gallery ("triw", n)`
`t = gallery ("triw", n, alpha)`
`t = gallery ("triw", n, alpha, k)`
Create an upper triangular matrix discussed by Kahan, Golub, and Wilkinson.

`a = gallery ("uniformdata", [M N ...], j)`
`a = gallery ("uniformdata", M, N, ..., j)`
`a = gallery ("uniformdata", ..., "class")`
Create a matrix with random samples from the standard uniform distribution (range [0,1]).

The first input is a matrix of dimensions describing the size of the output. The dimensions can also be input as comma-separated arguments.

The input *j* is an integer index in the range [0, 2^32-1]. The values of the output matrix are always exactly the same (reproducibility) for a given size input and *j* index.

The final optional argument determines the class of the resulting matrix. Possible values for *class*: `"single"`, `"double"`. The default is `"double"`.

`a = gallery ("wathen", nx, ny)`
`a = gallery ("wathen", nx, ny, k)`
Create the Wathen matrix.

`[a, b] = gallery ("wilk", n)`
Create various specific matrices devised/discussed by Wilkinson.

`h = hadamard (n)`
Construct a Hadamard matrix (Hn) of size *n*-by-*n*.

The size *n* must be of the form $2^k * p$ in which p is one of 1, 12, 20 or 28. The returned matrix is normalized, meaning `Hn(:,1) == 1` and `Hn(1,:) == 1`.

Some of the properties of Hadamard matrices are:

- `kron (Hm, Hn)` is a Hadamard matrix of size *m*-by-*n*.
- `Hn * Hn' = n * eye (n)`.
- The rows of Hn are orthogonal.
- `det (A) <= abs (det (Hn))` for all *A* with `abs (A(i, j)) <= 1`.
- Multiplying any row or column by -1 and the matrix will remain a Hadamard matrix.

**See also:** [compan], page 878, [hankel], page 598, [toeplitz], page 600.

`h = hankel (c)`
`h = hankel (c, r)`
Return the Hankel matrix constructed from the first column *c*, and (optionally) the last row *r*.

If the last element of *c* is not the same as the first element of *r*, the last element of *c* is used. If the second argument is omitted, it is assumed to be a vector of zeros with the same size as *c*.

A Hankel matrix formed from an m-vector $c$, and an n-vector $r$, has the elements

$$H(i,j) = \begin{cases} c_{i+j-1}, & i+j-1 \le m; \\ r_{i+j-m}, & \text{otherwise.} \end{cases}$$

**See also:** [hadamard], page 598, [toeplitz], page 600.

`h = hilb (n)`

Return the Hilbert matrix of order $n$.

The $i,j$ element of a Hilbert matrix is defined as

$$H(i,j) = \frac{1}{(i+j-1)}$$

Hilbert matrices are close to being singular which make them difficult to invert with numerical routines. Comparing the condition number of a random matrix 5x5 matrix with that of a Hilbert matrix of order 5 reveals just how difficult the problem is.

```
cond (rand (5))
   ⇒ 14.392
cond (hilb (5))
   ⇒ 4.7661e+05
```

**See also:** [invhilb], page 599.

`hinv = invhilb (n)`

Return the inverse of the Hilbert matrix of order $n$.

This can be computed exactly using

$$\begin{aligned} A_{ij} &= -1^{i+j}(i+j-1)\binom{n+i-1}{n-j}\binom{n+j-1}{n-i}\binom{i+j-2}{i-2}^2 \\ &= \frac{p(i)p(j)}{(i+j-1)} \end{aligned}$$

where

$$p(k) = -1^k \binom{k+n-1}{k-1}\binom{n}{k}$$

The validity of this formula can easily be checked by expanding the binomial coefficients in both formulas as factorials. It can be derived more directly via the theory of Cauchy matrices. See J. W. Demmel, *Applied Numerical Linear Algebra*, p. 92.

Compare this with the numerical calculation of `inv (hilb (n))`, which suffers from the ill-conditioning of the Hilbert matrix, and the finite precision of your computer's floating point arithmetic.

**See also:** [hilb], page 599.

`M = magic (n)`

Create an $n$-by-$n$ magic square.

A magic square is an arrangement of the integers `1:n^2` such that the row sums, column sums, and diagonal sums are all equal to the same value.

Note: $n$ must be a scalar greater than or equal to 3. If you supply $n$ less than 3, magic returns either a nonmagic square, or else the degenerate magic squares 1 and [].

`P = pascal (n)`
`P = pascal (n, t)`

Return the Pascal matrix of order $n$ if `t = 0`.

The default value of $t$ is 0.

When `t = 1`, return the pseudo-lower triangular Cholesky factor of the Pascal matrix (The sign of some columns may be negative). This matrix is its own inverse, that is `pascal (n, 1) ^ 2 == eye (n)`.

If `t = -1`, return the true Cholesky factor with strictly positive values on the diagonal.

If `t = 2`, return a transposed and permuted version of `pascal (n, 1)`, which is the cube root of the identity matrix. That is, `pascal (n, 2) ^ 3 == eye (n)`.

**See also:** [chol], page 657.

`R = rosser ()`

Return the Rosser matrix.

This is a difficult test case used to evaluate eigenvalue algorithms.

**See also:** [wilkinson], page 601, [eig], page 649.

`T = toeplitz (c)`
`T = toeplitz (c, r)`

Return the Toeplitz matrix constructed from the first column $c$, and optionally the first row $r$.

If the second argument is omitted, the first row is taken to be the same as the first column. If the first element of $r$ is not the same as the first element of $c$, the first element of $c$ is used.

A Toeplitz, or diagonal-constant, matrix has the same value along each diagonal. Although it need not be square, it often is. An MxN Toeplitz matrix has the form:

$$\begin{bmatrix} c_1 & r_2 & r_3 & \cdots & r_n \\ c_2 & c_1 & r_2 & \cdots & r_{n-1} \\ c_3 & c_2 & c_1 & \cdots & r_{n-2} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ c_m & c_{m-1} & c_{m-2} & \ldots & cm-n+1 \end{bmatrix}$$

**See also:** [hankel], page 598.

`V = vander (c)`
`V = vander (c, n)`

Return the Vandermonde matrix whose next to last column is $c$.

If $n$ is specified, it determines the number of columns; otherwise, $n$ is taken to be equal to the length of $c$.

A Vandermonde matrix has the form:

$$\begin{bmatrix} c_1^{n-1} & \cdots & c_1^2 & c_1 & 1 \\ c_2^{n-1} & \cdots & c_2^2 & c_2 & 1 \\ \vdots & \ddots & \vdots & \vdots & \vdots \\ c_n^{n-1} & \cdots & c_n^2 & c_n & 1 \end{bmatrix}$$

**See also:** [polyfit], page 883.

`W = wilkinson (n)`

Return the Wilkinson matrix of order *n*.

Wilkinson matrices are symmetric and tridiagonal with pairs of nearly, but not exactly, equal eigenvalues. They are useful in testing the behavior and performance of eigenvalue solvers.

**See also:** [rosser], page 600, [eig], page 649.

# 17 Arithmetic

Unless otherwise noted, all of the functions described in this chapter will work for real and complex scalar, vector, or matrix arguments. Functions described as *mapping functions* apply the given operation individually to each element when given a matrix argument. For example:

```
sin ([1, 2; 3, 4])
     ⇒  0.84147   0.90930
        0.14112  -0.75680
```

## 17.1 Exponents and Logarithms

`y = exp (x)`
Compute $e^x$ for each element of $x$.

To compute the matrix exponential, see Chapter 18 [Linear Algebra], page 647.

**See also:** [log], page 603.

`y = expm1 (x)`
Compute $e^x - 1$ accurately in the neighborhood of zero.

**See also:** [exp], page 603.

`y = log (x)`
Compute the natural logarithm, $\ln(x)$, for each element of $x$.

To compute the matrix logarithm, see Chapter 18 [Linear Algebra], page 647.

**See also:** [exp], page 603, [log1p], page 603, [log2], page 603, [log10], page 603, [logspace], page 585.

`y = reallog (x)`
Return the real-valued natural logarithm of each element of $x$.

If any element results in a complex return value `reallog` aborts and issues an error.

**See also:** [log], page 603, [realpow], page 604, [realsqrt], page 604.

`y = log1p (x)`
Compute $\ln(1 + x)$ accurately in the neighborhood of zero.

**See also:** [log], page 603, [exp], page 603, [expm1], page 603.

`y = log10 (x)`
Compute the base-10 logarithm of each element of $x$.

**See also:** [log], page 603, [log2], page 603, [logspace], page 585, [exp], page 603.

`y = log2 (x)`
`[f, e] = log2 (x)`
Compute the base-2 logarithm of each element of $x$.

If called with one output, compute the base-2 logarithm such that $2^y = x$.

If called with two output arguments, split $x$ into binary mantissa ($f$) and exponent ($e$) such that $x = f \cdot 2^e$ where $\frac{1}{2} \leq |f| < 1$ and $e$ is an integer. If $x = 0$, $f = e = 0$.

**See also:** [pow2], page 604, [log], page 603, [log10], page 603, [exp], page 603.

`y = pow2 (x)`
`y = pow2 (f, e)`

With one input argument, compute $y = 2^x$ for each element of $x$.

With two input arguments, return $y = f \cdot 2^e$, where for complex inputs only the real part of both inputs is regarded and from $e$ only the real integer part. This calling form corresponds to C/C++ standard function `ldexp()`.

**See also:** [log2], page 603, [nextpow2], page 604, [power], page 173.

`n = nextpow2 (x)`

Compute the exponent of the next power of two not smaller than the input.

For each element in the input array $x$, return the smallest integer $n$ such that $2^n \geq |x|$. For input elements equal to zero, return zero.

**See also:** [pow2], page 604, [log2], page 603.

`z = realpow (x, y)`

Compute the real-valued, element-by-element power operator.

This is equivalent to `x .^ y`, except that `realpow` reports an error if any return value is complex.

**See also:** [power], page 173, [reallog], page 603, [realsqrt], page 604.

`y = sqrt (x)`

Compute the square root of each element of $x$.

If $x$ is negative, a complex result is returned.

To compute the matrix square root, see Chapter 18 [Linear Algebra], page 647.

**See also:** [realsqrt], page 604, [nthroot], page 604.

`y = realsqrt (x)`

Return the real-valued square root of each element of $x$.

If any element results in a complex return value `realsqrt` aborts and issues an error.

**See also:** [sqrt], page 604, [realpow], page 604, [reallog], page 603.

`y = cbrt (x)`

Compute the real-valued cube root of each element of $x$.

Unlike `x^(1/3)`, the result will be negative if $x$ is negative.

If any element of $x$ is complex, `cbrt` aborts with an error.

**See also:** [nthroot], page 604.

`y = nthroot (x, n)`

Compute the real (non-complex) $n$-th root of $x$.

$x$ must have all real entries and $n$ must be a scalar. If $n$ is an even integer and $x$ has negative entries then `nthroot` aborts and issues an error.

Example:

```
nthroot (-1, 3)
⇒ -1
(-1) ^ (1 / 3)
⇒ 0.50000 - 0.86603i
```

**See also:** [realsqrt], page 604, [sqrt], page 604, [cbrt], page 604.

## 17.2 Complex Arithmetic

In the descriptions of the following functions, $z$ is the complex number $x + iy$, where $i$ is defined as $\sqrt{-1}$.

`z = abs (x)`

Compute the magnitude of $x$.

The magnitude is defined as $|z| = \sqrt{x^2 + y^2}$.

For example:

```
abs (3 + 4i)
     ⇒ 5
```

**See also:** [arg], page 605.

`theta = arg (z)`
`theta = angle (z)`

Compute the argument, i.e., angle of $z$.

This is defined as, $\theta = atan2(y, x)$, in radians.

For example:

```
arg (3 + 4i)
     ⇒ 0.92730
```

**See also:** [abs], page 605.

`zc = conj (z)`

Return the complex conjugate of $z$.

The complex conjugate is defined as $\bar{z} = x - iy$.

**See also:** [real], page 606, [imag], page 606.

`zsort = cplxpair (z)`
`zsort = cplxpair (z, tol)`
`zsort = cplxpair (z, tol, dim)`

Sort the numbers $z$ into complex conjugate pairs ordered by increasing real part.

The negative imaginary complex numbers are placed first within each pair. All real numbers (those with `abs (imag (z)) / abs (z) < tol`) are placed after the complex pairs.

*tol* is a weighting factor in the range [0, 1) which determines the tolerance of the matching. The default value is `100 * eps` and the resulting tolerance for a given complex pair is `tol * abs (z(i))`.

By default the complex pairs are sorted along the first non-singleton dimension of $z$. If *dim* is specified, then the complex pairs are sorted along this dimension.

Signal an error if some complex numbers could not be paired. Signal an error if all complex numbers are not exact conjugates (to within *tol*). Note that there is no defined order for pairs with identical real parts but differing imaginary parts.

```
cplxpair (exp (2i*pi*[0:4]'/5)) == exp (2i*pi*[3; 2; 4; 1; 0]/5)
```

`y = imag (z)`
Return the imaginary part of $z$ as a real number.

**See also:** [real], page 606, [conj], page 605.

`x = real (z)`
Return the real part of $z$.

**See also:** [imag], page 606, [conj], page 605.

## 17.3 Trigonometry

Octave provides the following trigonometric functions where angles are specified in radians. To convert from degrees to radians multiply by $\pi/180$ or use the `deg2rad` function. For example, `sin (30 * pi/180)` returns the sine of 30 degrees. As an alternative, Octave provides a number of trigonometric functions which work directly on an argument specified in degrees. These functions are named after the base trigonometric function with a '`d`' suffix. As an example, `sin` expects an angle in radians while `sind` expects an angle in degrees.

Octave uses the C library trigonometric functions. It is expected that these functions are defined by the ISO/IEC 9899 Standard. This Standard is available at: http://www.open-std.org/jtc1/sc22/wg14/www/docs/n1124.pdf. Section F.9.1 deals with the trigonometric functions. The behavior of most of the functions is relatively straightforward. However, there are some exceptions to the standard behavior. Many of the exceptions involve the behavior for -0. The most complex case is atan2. Octave exactly implements the behavior given in the Standard. Including $atan2(\pm 0, -0)$ returns $\pm\pi$.

It should be noted that MATLAB uses different definitions which apparently do not distinguish -0.

*rad* = `deg2rad (`*deg*`)`
Convert degrees to radians.

The input *deg* must be a scalar, vector, or N-dimensional array of double or single floating point values. *deg* may be complex in which case the real and imaginary components are converted separately.

The output *rad* is the same size and shape as *deg* with degrees converted to radians using the conversion constant `pi/180`.

Example:

```
deg2rad ([0, 90, 180, 270, 360])
  ⇒  0.00000   1.57080   3.14159   4.71239   6.28319
```

**See also:** [rad2deg], page 606.

*deg* = `rad2deg (`*rad*`)`
Convert radians to degrees.

The input *rad* must be a scalar, vector, or N-dimensional array of double or single floating point values. *rad* may be complex in which case the real and imaginary components are converted separately.

The output *deg* is the same size and shape as *rad* with radians converted to degrees using the conversion constant `180/pi`.

Example:

```
rad2deg ([0, pi/2, pi, 3/2*pi, 2*pi])
  ⇒  0    90   180   270   360
```

**See also:** [deg2rad], page 606.

`y = sin (x)`
Compute the sine for each element of $x$ in radians.

**See also:** [asin], page 607, [sind], page 609, [sinh], page 608.

`y = cos (x)`
Compute the cosine for each element of $x$ in radians.

**See also:** [acos], page 607, [cosd], page 609, [cosh], page 608.

`y = tan (z)`
Compute the tangent for each element of $x$ in radians.

**See also:** [atan], page 607, [tand], page 609, [tanh], page 608.

`y = sec (x)`
Compute the secant for each element of $x$ in radians.

**See also:** [asec], page 607, [secd], page 609, [sech], page 608.

`y = csc (x)`
Compute the cosecant for each element of $x$ in radians.

**See also:** [acsc], page 608, [cscd], page 610, [csch], page 608.

`y = cot (x)`
Compute the cotangent for each element of $x$ in radians.

**See also:** [acot], page 608, [cotd], page 610, [coth], page 608.

`y = asin (x)`
Compute the inverse sine in radians for each element of $x$.

**See also:** [sin], page 607, [asind], page 610.

`y = acos (x)`
Compute the inverse cosine in radians for each element of $x$.

**See also:** [cos], page 607, [acosd], page 610.

`y = atan (x)`
Compute the inverse tangent in radians for each element of $x$.

**See also:** [tan], page 607, [atand], page 610.

`y = asec (x)`
Compute the inverse secant in radians for each element of $x$.

**See also:** [sec], page 607, [asecd], page 610.

`y = acsc (x)`
Compute the inverse cosecant in radians for each element of *x*.

**See also:** [csc], page 607, [acscd], page 610.

`y = acot (x)`
Compute the inverse cotangent in radians for each element of *x*.

**See also:** [cot], page 607, [acotd], page 610.

`y = sinh (x)`
Compute the hyperbolic sine for each element of *x*.

**See also:** [asinh], page 608, [cosh], page 608, [tanh], page 608.

`y = cosh (x)`
Compute the hyperbolic cosine for each element of *x*.

**See also:** [acosh], page 608, [sinh], page 608, [tanh], page 608.

`y = tanh (x)`
Compute hyperbolic tangent for each element of *x*.

**See also:** [atanh], page 608, [sinh], page 608, [cosh], page 608.

`y = sech (x)`
Compute the hyperbolic secant of each element of *x*.

**See also:** [asech], page 608.

`y = csch (x)`
Compute the hyperbolic cosecant of each element of *x*.

**See also:** [acsch], page 609.

`y = coth (x)`
Compute the hyperbolic cotangent of each element of *x*.

**See also:** [acoth], page 609.

`y = asinh (x)`
Compute the inverse hyperbolic sine for each element of *x*.

**See also:** [sinh], page 608.

`y = acosh (x)`
Compute the inverse hyperbolic cosine for each element of *x*.

**See also:** [cosh], page 608.

`y = atanh (x)`
Compute the inverse hyperbolic tangent for each element of *x*.

**See also:** [tanh], page 608.

`y = asech (x)`
Compute the inverse hyperbolic secant of each element of *x*.

**See also:** [sech], page 608.

*y* = acsch (*x*)
: Compute the inverse hyperbolic cosecant of each element of *x*.

  **See also:** [csch], page 608.

*y* = acoth (*x*)
: Compute the inverse hyperbolic cotangent of each element of *x*.

  **See also:** [coth], page 608.

*angle* = atan2 (*y*, *x*)
: Compute atan (*y* / *x*) for corresponding elements of *y* and *x*.

  *y* and *x* must match in size and orientation. The signs of elements of *y* and *x* are used to determine the quadrants of each resulting value.

  This function is equivalent to `arg (complex (x, y))`.

  **See also:** [tan], page 607, [tand], page 609, [tanh], page 608, [atanh], page 608.

Octave provides the following trigonometric functions where angles are specified in degrees. These functions produce true zeros at the appropriate intervals rather than the small round-off error that occurs when using radians. For example:

```
cosd (90)
     ⇒ 0
cos (pi/2)
     ⇒ 6.1230e-17
```

*y* = sind (*x*)
: Compute the sine for each element of *x* in degrees.

  The function is more accurate than `sin` for large values of *x* and for multiples of 180 degrees (`x/180` is an integer) where `sind` returns 0 rather than a small value on the order of eps.

  **See also:** [asind], page 610, [sin], page 607.

*y* = cosd (*x*)
: Compute the cosine for each element of *x* in degrees.

  The function is more accurate than `cos` for large values of *x* and for multiples of 90 degrees (`x = 90 + 180*n` with n an integer) where `cosd` returns 0 rather than a small value on the order of eps.

  **See also:** [acosd], page 610, [cos], page 607.

*y* = tand (*x*)
: Compute the tangent for each element of *x* in degrees.

  Returns zero for elements where `x/180` is an integer and `Inf` for elements where `(x-90)/180` is an integer.

  **See also:** [atand], page 610, [tan], page 607.

*y* = secd (*x*)
: Compute the secant for each element of *x* in degrees.

  **See also:** [asecd], page 610, [sec], page 607.

`y = cscd (x)`
Compute the cosecant for each element of $x$ in degrees.

**See also:** [acscd], page 610, [csc], page 607.

`y = cotd (x)`
Compute the cotangent for each element of $x$ in degrees.

**See also:** [acotd], page 610, [cot], page 607.

`y = asind (x)`
Compute the inverse sine in degrees for each element of $x$.

**See also:** [sind], page 609, [asin], page 607.

`y = acosd (x)`
Compute the inverse cosine in degrees for each element of $x$.

**See also:** [cosd], page 609, [acos], page 607.

`y = atand (x)`
Compute the inverse tangent in degrees for each element of $x$.

**See also:** [tand], page 609, [atan], page 607.

`d = atan2d (y, x)`
Compute atan ($y$ / $x$) in degrees for corresponding elements from $y$ and $x$.

**See also:** [tand], page 609, [atan2], page 609.

`y = asecd (x)`
Compute the inverse secant in degrees for each element of $x$.

**See also:** [secd], page 609, [asec], page 607.

`y = acscd (x)`
Compute the inverse cosecant in degrees for each element of $x$.

**See also:** [cscd], page 610, [acsc], page 608.

`y = acotd (x)`
Compute the inverse cotangent in degrees for each element of $x$.

**See also:** [cotd], page 610, [acot], page 608.

Finally, there are two trigonometric functions that calculate special arguments with increased accuracy.

`y = sinpi (x)`
Compute sine ($x$ * pi) for each element of $x$ accurately.

The ordinary `sin` function uses IEEE 754 floating point numbers and may produce results that are very close (within a few eps) of the correct value, but which are not exact. The `sinpi` function is more accurate and returns 0 exactly for integer values of $x$ and +1/-1 for half-integer values (e.g., . . . , -3/2, -1/2, 1/2, 3/2, . . . ).

Example
comparison of `sin` and `sinpi` for integer values of *x*

```
sin ([0, 1, 2, 3] * pi)
⇒
     0   1.2246e-16  -2.4493e-16   3.6739e-16

sinpi ([0, 1, 2, 3])
⇒
       0   0   0   0
```

**See also:** [cospi], page 611, [sin], page 607.

*y* = `cospi (`*x*`)`

Compute cosine (*x* * pi) for each element of *x* accurately.

The ordinary `cos` function uses IEEE 754 floating point numbers and may produce results that are very close (within a few eps) of the correct value, but which are not exact. The `cospi` function is more accurate and returns 0 exactly for half-integer values of *x* (e.g., . . . , -3/2, -1/2, 1/2, 3/2, . . . ), and +1/-1 for integer values.

Example
comparison of `cos` and `cospi` for half-integer values of *x*

```
cos ([-3/2, -1/2, 1/2, 3/2] * pi)
⇒
     -1.8370e-16   6.1232e-17   6.1232e-17  -1.8370e-16

cospi ([-3/2, -1/2, 1/2, 3/2])
⇒
       0   0   0   0
```

**See also:** [sinpi], page 610, [cos], page 607.

## 17.4 Sums and Products

*y* = `sum (`*x*`)`
*y* = `sum (`*x*`, `*dim*`)`
*y* = `sum (`*x*`, `*vecdim*`)`
*y* = `sum (`*x*`, "`*all*`")`
*y* = `sum (...,` *outtype*`)`
*y* = `sum (...,` *nanflag*`)`

Compute the sum of the elements of *x*.

If *x* is a vector, then `sum (x)` returns the sum of the elements in *x*.

If *x* is a matrix, then `sum (x)` returns a row vector with each element containing the sum of the corresponding column in *x*.

If *x* is an array, then `sum(x)` computes the sum along the first non-singleton dimension of *x*.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension in *x*, including any dimension exceeding `ndims (x)`, will return *x*.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored.

Specifying the dimension as `"all"` will cause `sum` to operate on all elements of *x*, and is equivalent to `cumsum (x(:))`.

The optional input *outtype* specifies the data type that is returned as well as the class of the variable used for calculations. *outtype* can take the following values:

`"default"`
: Operations on floating point inputs (double or single) are performed in their native data type, while operations on integer, logical, and character data types are performed using doubles. Output is of type double, unless the input is single in which case the output is of type single.

`"double"`
: Operations are performed in double precision even for single precision inputs. Output is of type double.

`"extra"`
: For double precision inputs, `sum` will use a more accurate algorithm than straightforward summation. For single precision inputs, `"extra"` is the same as `"double"`. For all other data types, `"extra"` has no effect.

`"native"`
: Operations are performed in their native data types and output is of the same type as the input as reported by (`class (x)`). When the input is logical, `sum (x, "native")` is equivalent to `any (x)`.

The optional variable *nanflag* specifies whether to include or exclude NaN values from the calculation using any of the previously specified input argument combinations. The default value for *nanflag* is `"includenan"` which keeps NaN values in the calculation. To exclude NaN values set the value of *nanflag* to `"omitnan"`. The output will be *0*, if *x* consists of all NaN values in the operating dimension.

**See also:** [cumsum], page 613, [sumsq], page 615, [prod], page 612.

`y = prod (x)`
`y = prod (x, dim)`
`y = prod (x, vecdim)`
`y = prod (x, "all")`
`y = prod (..., outtype)`
`y = prod (..., nanflag)`

Compute the product of the elements of *x*.

If *x* is a vector, then `prod (x)` returns the product of the elements in *x*.

If *x* is a matrix, then `prod (x)` returns a row vector with each element containing the product of the corresponding column in *x*.

If *x* is an array, then `prod(x)` computes the product along the first non-singleton dimension of *x*.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension in *x*, including any dimension exceeding `ndims (x)`, will return *x*.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored.

Specifying the dimension as `"all"` will cause `prod` to operate on all elements of *x*, and is equivalent to `prod (x(:))`.

The optional input *outtype* specifies the data type that is returned as well as the class of the variable used for calculations. *outtype* can take the following values:

`"default"`
: Operations on floating point inputs (double or single) are performed in their native data type; while operations on integer, logical, and character data types are performed using doubles. Output is of type double, unless the input is single in which case the output is of type single.

`"double"`
: Operations are performed in double precision even for single precision inputs. Output is of type double.

`"native"`
: Operations are performed in their native data types and output is of the same type as the input as reported by (`class (x)`). When the input is logical, `prod (x, "native")` is equivalent to `all (x)`.

The optional variable *nanflag* specifies whether to include or exclude NaN values from the calculation using any of the previously specified input argument combinations. The default value for *nanflag* is `"includenan"` which keeps NaN values in the calculation. To exclude NaN values set the value of *nanflag* to `"omitnan"`. The output will be *1*, if *x* consists of all NaN values in the operating dimension.

**See also:** [cumprod], page 614, [sum], page 611.

```
y = cumsum (x)
y = cumsum (x, dim)
y = cumsum (x, vecdim)
y = cumsum (..., "all")
y = cumsum (..., direction)
y = cumsum (..., nanflag)
```

Compute the cumulative sum of elements in *x*.

If *x* is a vector, then `cumsum (x)` returns a vector of the same size with the cumulative sum of *x*.

If *x* is a matrix, then `cumsum (x)` returns a matrix of the same size with the cumulative sum along each column of *x*.

If *x* is an array, then `cumsum(x)` returns an array of the same size with the cumulative sum along the first non-singleton dimension of *x*.

The class of output *y* is the same as the class of input *x*, unless *x* is logical, in which case *y* is double.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension in *x*, including any dimension exceeding `ndims (x)`, will return *x*.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored.

Specifying the dimension as `"all"` will cause `cumsum` to operate on all elements of *x*, and is equivalent to `cumsum (x(:))`.

The optional input *direction* specifies how the operating dimension is traversed and can take the following values:

`"forward"` (default)
: The cumulative sum is computed from beginning (index 1) to end along the operating dimension.

`"reverse"`
: The cumulative sum is computed from end to beginning along the operating dimension.

The optional variable *nanflag* specifies whether to include or exclude NaN values from the calculation using any of the previously specified input argument combinations. The default value for *nanflag* is `"includenan"` which keeps NaN values in the calculation. To exclude NaN values set the value of *nanflag* to `"omitnan"`. The output will still contain NaN values if *x* consists of all NaN values in the operating dimension.

**See also:** [sum], page 611, [cumprod], page 614.

`y = cumprod (x)`
`y = cumprod (x, dim)`
`y = cumprod (x, vecdim)`
`y = cumprod (..., "all")`
`y = cumprod (..., direction)`
`y = cumprod (..., nanflag)`

Compute the cumulative product of elements in *x*.

If *x* is a vector, then `cumprod (x)` returns a vector of the same size with the cumulative product of *x*.

If *x* is a matrix, then `cumprod (x)` returns a matrix of the same size with the cumulative product along each column of *x*.

If *x* is an array, then `cumprod(x)` returns an array of the same size with the cumulative product along the first non-singleton dimension of *x*.

The class of output *y* is the same as the class of input *x*, unless *x* is logical, in which case *y* is double.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension in *x*, including any dimension exceeding `ndims (x)`, will return *x*.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored.

Specifying the dimension as `"all"` will cause `cumprod` to operate on all elements of *x*, and is equivalent to `cumprod (x(:))`.

The optional input *direction* specifies how the operating dimension is traversed and can take the following values:

`"forward"` (default)
: The cumulative product is computed from beginning (index 1) to end along the operating dimension.

`"reverse"`
: The cumulative product is computed from end to beginning along the operating dimension.

The optional variable *nanflag* specifies whether to include or exclude NaN values from the calculation using any of the previously specified input argument combinations. The default value for *nanflag* is `"includenan"` which keeps NaN values in the calculation. To exclude NaN values set the value of *nanflag* to `"omitnan"`. The output will still contain NaN values if *x* consists of all NaN values in the operating dimension.

**See also:** [prod], page 612, [cumsum], page 613.

`y = sumsq (x)`
`y = sumsq (x, dim)`
`y = sumsq (x, vecdim)`
`y = sumsq (x, "all")`
`y = sumsq (..., outtype)`
`y = sumsq (..., nanflag)`

Compute the sum of squares of the elements of *x*.

If *x* is a vector, then `sumsq (x)` returns the sum of the squares of the elements in *x*.

If *x* is a matrix, then `sumsq (x)` returns a row vector with each element containing the sum of squares of the corresponding column in *x*.

If *x* is an array, then `sumsq(x)` computes the sum of squares along the first non-singleton dimension of *x*.

This function is conceptually equivalent to computing

```
sum (x .* conj (x))
```

but it uses less memory and avoids calling `conj` if *x* is real.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension in *x*, including any dimension exceeding `ndims (x)`, will return a sum of squares equal to `x.^2`.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored.

Specifying the dimension as `"all"` will cause `prod` to operate on all elements of *x*, and is equivalent to `sumsq (x(:))`.

The optional input *outtype* specifies the data type that is returned as well as the class of the variable used for calculations. *outtype* can take the following values:

`"default"`
: Operations on floating point inputs (double or single) are performed in their native data type; while operations on integer, logical, and character data types are performed using doubles. Output is of type double, unless the input is single in which case the output is of type single.

`"double"`
: Operations are performed in double precision even for single precision inputs. Output is of type double.

`"native"`
: Operations are performed in their native data types and output is of the same type as the input as reported by (`class (x)`). When the input is logical, `sumsq (x, "native")` is equivalent to `all (x)`.

The optional variable *nanflag* specifies whether to include or exclude NaN values from the calculation using any of the previously specified input argument combinations. The default value for *nanflag* is `"includenan"` which keeps NaN values in the calculation. To exclude NaN values set the value of *nanflag* to `"omitnan"`. The output will be *0*, if *x* consists of all NaN values in the operating dimension.

**See also:** [sum], page 611, [prod], page 612.

## 17.5 Utility Functions

`y = ceil (x)`

Return the smallest integer not less than *x*.

This is equivalent to rounding towards positive infinity.

If *x* is complex, return `ceil (real (x)) + ceil (imag (x)) * I`.

```
ceil ([-2.7, 2.7])
    ⇒ -2    3
```

**See also:** [floor], page 616, [round], page 617, [fix], page 616.

`y = fix (x)`

Truncate fractional portion of *x* and return the integer portion.

This is equivalent to rounding towards zero. If *x* is complex, return `fix (real (x)) + fix (imag (x)) * I`.

```
fix ([-2.7, 2.7])
   ⇒ -2    2
```

**See also:** [ceil], page 616, [floor], page 616, [round], page 617.

`y = floor (x)`

Return the largest integer not greater than *x*.

This is equivalent to rounding towards negative infinity. If *x* is complex, return `floor (real (x)) + floor (imag (x)) * I`.

```
floor ([-2.7, 2.7])
     ⇒ -3    2
```

**See also:** [ceil], page 616, [round], page 617, [fix], page 616.

`y = round (x)`

Return the integer nearest to *x*.

If *x* is complex, return `round (real (x)) + round (imag (x)) * I`. If there are two nearest integers, return the one further away from zero.

```
round ([-2.7, 2.7])
     ⇒ -3    3
```

**See also:** [ceil], page 616, [floor], page 616, [fix], page 616, [roundb], page 617.

`y = roundb (x)`

Return the integer nearest to *x*. If there are two nearest integers, return the even one (banker's rounding).

If *x* is complex, return `roundb (real (x)) + roundb (imag (x)) * I`.

**See also:** [round], page 617.

`m = max (x)`
`m = max (x, [], dim)`
`m = max (x, [], vecdim)`
`m = max (x, [], "all")`
`m = max (x, [], nanflag)`
`[m, im] = max (...)`
`[m, im] = max (..., "linear")`
`m = max (x, y)`
`m = max (x, y, nanflag)`
`... = max (..., "ComparisonMethod", method)`

Find maximum values in the array *x*.

If *x* is a vector, then `max (x)` returns the maximum value of the elements in *x*.

If *x* is a matrix, then `max (x)` returns a row vector with each element containing the maximum value of the corresponding column in *x*.

If *x* is an array, then `max (x)` computes the maximum value along the first non-singleton dimension of *x*.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension of *x*, including any dimension exceeding `ndims (x)`, will return *x*.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored.

Specifying the dimension as `"all"` will cause `max` to operate on on all elements of *x*, and is equivalent to `max (x(:))`.

If called with two output arguments, `max` also returns the first index of the maximum value(s) in *x*. The second output argument is only valid when `max` operates on a single input array. Setting the `"linear"` flag returns the linear index to the corresponding minimum values in *x*.

If called with two input arrays (*x* and *y*), `max` return the pairwise maximum according to the rules for Section 19.2 [Broadcasting], page 691.

The optional variable *nanflag* specifies whether to include or exclude `NaN` values from the calculation using any of the previously specified input argument combinations. The default value for *nanflag* is `"omitnan"` which ignores `NaN` values in the calculation. To include `NaN` values, set the value of *nanflag* to `"includenan"`, in which case `max` will return `NaN`, if any element along the operating dimension is `NaN`.

The optional `"ComparisonMethod"` paired argument specifies the comparison method for numeric input and it applies to both one input and two input arrays. *method* can take any of the following values:

`'auto'` This is the default method, which compares elements by `real (x)` when $x$ is real, and by `abs (x)` when $x$ is complex.

`'real'` Compares elements by `real (x)` when $x$ is real or complex. For elements with equal real parts, a second comparison by `imag (x)` is performed.

`'abs'` Compares elements by `abs (x)` when $x$ is real or complex. For elements with equal magnitude, a second comparison by `angle (x)` in the interval from $-\pi$ to $\pi$ is performed.

**See also:** [min], page 618, [cummax], page 619, [cummin], page 620.

`m = min (x)`
`m = min (x, [], dim)`
`m = min (x, [], vecdim)`
`m = min (x, [], "all")`
`m = min (x, [], nanflag)`
`[m, im] = min (...)`
`[m, im] = min (..., "linear")`
`m = min (x, y)`
`m = min (x, y, nanflag)`
`... = min (..., "ComparisonMethod", method)`

Find minimum values in the array $x$.

If $x$ is a vector, then `min (x)` returns the minimum value of the elements in $x$.

If $x$ is a matrix, then `min (x)` returns a row vector with each element containing the minimum value of the corresponding column in $x$.

If $x$ is an array, then `min (x)` computes the minimum value along the first non-singleton dimension of $x$.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension of $x$, including any dimension exceeding `ndims (x)`, will return $x$.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of $x$, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored.

Specifying the dimension as `"all"` will cause `min` to operate on on all elements of $x$, and is equivalent to `min (x(:))`.

If called with two output arguments, `min` also returns the first index of the minimum value(s) in $x$. The second output argument is only valid when `min` operates on a single

input array. Setting the `"linear"` flag returns the linear index to the corresponding minimum values in *x*.

If called with two input arrays (*x* and *y*), `min` return the pairwise minimum according to the rules for Section 19.2 [Broadcasting], page 691.

The optional variable *nanflag* specifies whether to include or exclude `NaN` values from the calculation using any of the previously specified input argument combinations. The default value for *nanflag* is `"omitnan"` which ignores `NaN` values in the calculation. To include `NaN` values, set the value of *nanflag* to `"includenan"`, in which case `min` will return `NaN`, if any element along the operating dimension is `NaN`.

The optional `"ComparisonMethod"` paired argument specifies the comparison method for numeric input and it applies to both one input and two input arrays. *method* can take any of the following values:

`'auto'` This is the default method, which compares elements by `real (x)` when *x* is real, and by `abs (x)` when *x* is complex.

`'real'` Compares elements by `real (x)` when *x* is real or complex. For elements with equal real parts, a second comparison by `imag (x)` is performed.

`'abs'` Compares elements by `abs (x)` when *x* is real or complex. For elements with equal magnitude, a second comparison by `angle (x)` in the interval from $-\pi$ to $\pi$ is performed.

**See also:** [max], page 617, [cummin], page 620, [cummax], page 619.

```
M = cummax (x)
m = cummax (x, dim)
m = cummax (x, vecdim)
m = cummax (x, "all")
m = cummax (..., direction)
m = cummax (x, nanflag)
[m, im] = cummax (...)
[m, im] = cummax (..., "linear")
... = cummax (..., "ComparisonMethod", method)
```

Return the cumulative maximum values from the array *x*.

If *x* is a vector, then `cummax (x)` returns a vector of the same size with the cumulative maximum values of *x*.

If *x* is a matrix, then `cummax (x)` returns a matrix of the same size with the cumulative maximum values along each column of *x*.

If *x* is an array, then `cummax(x)` returns an array of the same size with the cumulative product along the first non-singleton dimension of *x*.

The class of output *y* is the same as the class of input *x*, unless *x* is logical, in which case *y* is double.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension in *x*, including any dimension exceeding `ndims (x)`, will return *x*.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all

dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored.

Specifying the dimension as `"all"` will cause `cummax` to operate on all elements of *x*, and is equivalent to `cummax (x(:))`.

The optional input *direction* specifies how the operating dimension is traversed and can take the following values:

`"forward"` (default)
: The cumulative maximum values are computed from beginning (index 1) to end along the operating dimension.

`"reverse"`
: The cumulative maximum values are computed from end to beginning along the operating dimension.

The optional variable *nanflag* specifies whether to include or exclude `NaN` values from the calculation using any of the previously specified input argument combinations. The default value for *nanflag* is `"omitnan"` which ignores `NaN` values in the calculation. To include `NaN` values, set the value of *nanflag* to `"includenan"`, in which case `max` will return `NaN`, if any element along the operating dimension is `NaN`.

If called with two output arguments, `cummax` also returns the first index of the maximum value(s) in *x*. Setting the `"linear"` flag returns the linear index to the corresponding minimum values in *x*.

The optional `"ComparisonMethod"` paired argument specifies the comparison method for numeric input and it applies to both one input and two input arrays. *method* can take any of the following values:

`'auto'` (default)
: Compare elements by `real (x)` when *x* is real, and by `abs (x)` when *x* is complex.

`'real'`
: Compare elements by `real (x)` when *x* is real or complex. For elements with equal real parts, a second comparison by `imag (x)` is performed.

`'abs'`
: Compare elements by `abs (x)` when *x* is real or complex. For elements with equal magnitude, a second comparison by `angle (x)` in the interval from $-\pi$ to $\pi$ is performed.

**See also:** [cummin], page 620, [max], page 617, [min], page 618.

```
M = cummin (x)
m = cummin (x, dim)
m = cummin (x, vecdim)
m = cummin (x, "all")
m = cummin (..., direction)
m = cummin (x, nanflag)
[m, im] = cummin (...)
[m, im] = cummin (..., "linear")
... = cummin (..., "ComparisonMethod", method)
```

Return the cumulative minimum values from the array *x*.

If *x* is a vector, then `cummin (x)` returns a vector of the same size with the cumulative minimum values of *x*.

If *x* is a matrix, then `cummin (x)` returns a matrix of the same size with the cumulative minimum values along each column of *x*.

If *x* is an array, then `cummin(x)` returns an array of the same size with the cumulative product along the first non-singleton dimension of *x*.

The class of output *y* is the same as the class of input *x*, unless *x* is logical, in which case *y* is double.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension in *x*, including any dimension exceeding `ndims (x)`, will return *x*.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored.

Specifying the dimension as `"all"` will cause `cummin` to operate on all elements of *x*, and is equivalent to `cummin (x(:))`.

The optional input *direction* specifies how the operating dimension is traversed and can take the following values:

`"forward"` (default)
: The cumulative minimum values are computed from beginning (index 1) to end along the operating dimension.

`"reverse"`
: The cumulative minimum values are computed from end to beginning along the operating dimension.

The optional variable *nanflag* specifies whether to include or exclude `NaN` values from the calculation using any of the previously specified input argument combinations. The default value for *nanflag* is `"omitnan"` which ignores `NaN` values in the calculation. To include `NaN` values, set the value of *nanflag* to `"includenan"`, in which case `max` will return `NaN`, if any element along the operating dimension is `NaN`.

If called with two output arguments, `cummin` also returns the first index of the minimum value(s) in *x*. Setting the `"linear"` flag returns the linear index to the corresponding minimum values in *x*.

The optional `"ComparisonMethod"` paired argument specifies the comparison method for numeric input and it applies to both one input and two input arrays. *method* can take any of the following values:

`'auto'` (default)
: Compare elements by `real (x)` when *x* is real, and by `abs (x)` when *x* is complex.

`'real'`
: Compare elements by `real (x)` when *x* is real or complex. For elements with equal real parts, a second comparison by `imag (x)` is performed.

`'abs'` Compare elements by `abs (x)` when $x$ is real or complex. For elements with equal magnitude, a second comparison by `angle (x)` in the interval from $-\pi$ to $\pi$ is performed.

**See also:** [cummax], page 619, [min], page 618, [max], page 617.

`h = hypot (x, y)`
`h = hypot (x, y, z, ...)`

Compute the element-by-element square root of the sum of the squares of $x$ and $y$.

This is equivalent to `sqrt (x.^2 + y.^2)`, but is calculated in a manner that avoids overflows for large values of $x$ or $y$.

`hypot` can also be called with more than 2 arguments; in this case, the arguments are accumulated from left to right:

```
hypot (hypot (x, y), z)
hypot (hypot (hypot (x, y), z), w), etc.
```

`dx = gradient (m)`
`[dx, dy, dz, ...] = gradient (m)`
`[...] = gradient (m, s)`
`[...] = gradient (m, sx, sy, sz, ...)`
`[...] = gradient (f, x0)`
`[...] = gradient (f, x0, s)`
`[...] = gradient (f, x0, sx, sy, ...)`

Calculate the gradient of sampled data or a function.

$$\text{grad } F(x,y,z) \equiv \nabla F = \frac{\partial F_x}{\partial x}\hat{i} + \frac{\partial F_y}{\partial y}\hat{j} + \frac{\partial F_z}{\partial z}\hat{k}$$

If $m$ is a vector, calculate the one-dimensional numerical gradient of $m$. If $m$ is a matrix the gradient is calculated for each dimension. The return argument(s) are the estimated partial derivatives for each dimension at the specified sample points.

The default spacing of between data points is 1. A constant spacing between points can be specified with the $s$ parameter. If $s$ is a scalar, the single spacing value is used for all dimensions. Otherwise, separate values of the spacing can be supplied by the $sx$, ... arguments. Scalar values specify an equidistant spacing. Vector values for the $sx$, ... arguments specify the coordinate for that dimension. The length must match the respective dimension of $m$.

If the first argument $f$ is a function handle, the gradient of the function is calculated for the points in *x0*. As with sampled data, the spacing values between the points from which the gradient is estimated can be set via the $s$ or $dx$, $dy$, ... arguments. By default a spacing of 1 is used, however this is normally overly large unless the function is very slowly varying, and it is often necessary to specify a smaller sample spacing.

Example: numerical gradient of `cos` (analytically = `-sin`)

```
gradient (@cos, pi/2, .1)
⇒ -0.9983
-sin (pi/2)
⇒ -1
```

Programming Notes: The value for interior data points is approximated using the central difference.

```
y'(i) = 1/2 * (y(i+1) - y(i-1)).
```

At boundary points a linear extrapolation is applied.

```
y'(1) = y(2) - y(1).
```

**See also:** [diff], page 567, [del2], page 624.

`z = dot (x, y)`
`z = dot (x, y, dim)`

Compute the dot product of two vectors.

If $x$ and $y$ are matrices, calculate the dot products along the first non-singleton dimension.

If the optional argument *dim* is given, calculate the dot products along this dimension.

Implementation Note: This is equivalent to `sum (conj (X) .* Y, dim)`, but avoids forming a temporary array and is faster. When $X$ and $Y$ are column vectors, the result is equivalent to `X' * Y`. Although, `dot` is defined for integer arrays, the output may differ from the expected result due to the limited range of integer objects.

**See also:** [cross], page 623, [divergence], page 624, [tensorprod], page 675.

`z = cross (x, y)`
`z = cross (x, y, dim)`

Compute the vector cross product of two 3-dimensional vectors $x$ and $y$.

If $x$ and $y$ are arrays, the cross product is applied along the first dimension with three elements.

The optional argument *dim* forces the cross product to be calculated along the specified dimension. An error will be produced if the specified dimension is not three elements in size.

In the case of a complex output, orthogonality of the output with respect to the inputs is also satisfied, and the condition

```
dot (conj (z), x) ≡ dot (conj (z), y) = 0
```

is met. `dot (z, x) = 0` and `dot (z, y) = 0` will not hold. Also note that instead of using the `dot` function, the inner product

```
z(:).' * x(:) ≡ z(:).' * y(:) = 0
```

will meet the orthogonality condition for vector input.

Example Code:

```
cross ([1, 1, 0], [0, 1, 1])
  ⇒
       1  -1   1

cross (magic (3), eye (3), 2)
  ⇒
       0   6  -1
      -7   0   3
       9  -4   0
```

**See also:** [dot], page 623, [curl], page 624, [divergence], page 624, [conj], page 605.

`div = divergence (x, y, z, fx, fy, fz)`
`div = divergence (fx, fy, fz)`
`div = divergence (x, y, fx, fy)`
`div = divergence (fx, fy)`

Calculate divergence of a vector field given by the arrays *fx*, *fy*, and *fz* or *fx*, *fy* respectively.

$$\text{div } \mathbf{F}(x, y, z) \equiv \nabla \cdot \mathbf{F} = \frac{\partial F_x}{\partial x} + \frac{\partial F_y}{\partial y} + \frac{\partial F_z}{\partial z}$$

The coordinates of the vector field can be given by the arguments *x*, *y*, *z* or *x*, *y* respectively.

**See also:** [curl], page 624, [gradient], page 622, [del2], page 624, [dot], page 623.

`[cx, cy, cz, v] = curl (x, y, z, fx, fy, fz)`
`[cz, v] = curl (x, y, fx, fy)`
`[...] = curl (fx, fy, fz)`
`[...] = curl (fx, fy)`
`v = curl (...)`

Calculate curl of vector field given by the arrays *fx*, *fy*, and *fz* or *fx*, *fy* respectively.

$$\text{curl } \mathbf{F}(x, y, z) \equiv \nabla \times \mathbf{F} = \left( \frac{\partial F_z}{\partial y} - \frac{\partial F_y}{\partial z}, \frac{\partial F_x}{\partial z} - \frac{\partial F_z}{\partial x}, \frac{\partial F_y}{\partial x} - \frac{\partial F_x}{\partial y} \right)$$

The coordinates of the vector field can be given by the arguments *x*, *y*, *z* or *x*, *y* respectively. *v* calculates the scalar component of the angular velocity vector in direction of the z-axis for two-dimensional input. For three-dimensional input the scalar rotation is calculated at each grid point in direction of the vector field at that point.

**See also:** [divergence], page 624, [gradient], page 622, [del2], page 624, [cross], page 623.

`L = del2 (M)`
`L = del2 (M, h)`
`L = del2 (M, dx, dy, ...)`

Calculate the discrete Laplace operator ($\nabla^2$).

For a 2-dimensional matrix $M(x, y)$ this is defined as

$$L = \frac{1}{4} \left( \frac{\partial^2 M}{\partial x^2} + \frac{\partial^2 M}{\partial y^2} \right)$$

For N-dimensional arrays the sum in parentheses is expanded to include second derivatives over the additional higher dimensions.

The spacing between evaluation points may be defined by *h*, which is a scalar defining the equidistant spacing in all dimensions. Alternatively, the spacing in each dimension may be defined separately by *dx*, *dy*, etc. A scalar spacing argument defines equidistant spacing, whereas a vector argument can be used to specify variable spacing. The length of the spacing vectors must match the respective dimension of *M*. The default spacing value is 1.

Dimensions with fewer than 3 data points are skipped. Boundary points are calculated from the linear extrapolation of interior points.

Example: Second derivative of 2*x^3

```
f = @(x) 2*x.^3;
dd = @(x) 12*x;
x = 1:6;
L = 4*del2 (f(x));
assert (L, dd (x));
```

**See also:** [gradient], page 622, [diff], page 567.

*f* = factorial (*n*)

Return the factorial of *n* where *n* is a real non-negative integer.

If *n* is a scalar, this is equivalent to `prod (1:n)`. For vector or matrix arguments, return the factorial of each element in the array.

For non-integers see the generalized factorial function `gamma`. Note that the factorial function grows large quite quickly, and even with double precision values overflow will occur if *n* > 171. For such cases consider `gammaln`.

**See also:** [prod], page 612, [gamma], page 637, [gammaln], page 639.

*pf* = factor (*q*)
[*pf*, *n*] = factor (*q*)

Return the prime factorization of *q*.

The prime factorization is defined as `prod (pf) == q` where every element of *pf* is a prime number. If `q == 1`, return 1. The output *pf* is of the same numeric class as the input.

With two output arguments, return the unique prime factors *pf* and their multiplicities. That is, `prod (pf .^ n) == q`.

Implementation Note: If the input *q* is `single` or `double`, then it must not exceed the corresponding `flintmax`. For larger inputs, cast them to `uint64` if they're less than 2^64:

```
factor (uint64 (18446744073709011493))
   ⇒      571111    761213  42431951
```

For even larger inputs, use `sym` if you have the Symbolic package installed and loaded:

```
factor (sym ('9444733049654361449941'))
   ⇒ (sym)
                  1           1
  1099511627689 ·8589934669
```

**See also:** [gcd], page 625, [lcm], page 626, [isprime], page 72, [primes], page 627.

*g* = gcd (*a1*, *a2*, ...)
[*g*, *v1*, ...] = gcd (*a1*, *a2*, ...)

Compute the greatest common divisor of *a1*, *a2*, ...

All arguments must be the same size or scalar. For arrays, the greatest common divisor is calculated for each element individually. All elements must be ordinary or

Gaussian (complex) integers. Note that for Gaussian integers, the gcd is only unique up to a phase factor (multiplication by 1, -1, i, or -i), so an arbitrary greatest common divisor among the four possible is returned.

Optional return arguments *v1*, ..., contain integer vectors such that, $g = v_1 a_1 + v_2 a_2 + \cdots$

Example code:

```
gcd ([15, 9], [20, 18])
   ⇒  5  9
```

Programming tip: To find the GCD of all the elements of a single array, use `num2cell` instead of nested calls or a loop:

```
x = [30 42 70 105];    # vector or array of inputs
gcd (num2cell (x) {:})
   ⇒     1
```

**See also:** [lcm], page 626, [factor], page 625, [isprime], page 72.

`l = lcm (x, y)`
`l = lcm (x, y, ...)`

Compute the least common multiple of $x$ and $y$, or of the list of all arguments.

All inputs must be of the same size, or scalar. All elements must be real integer or Gaussian (complex) integer. For complex inputs, the result is unique only up to a phase factor (multiplication by +1, +i, -1, or -i), and one of the four is returned arbitrarily.

Example code:

```
lcm (5:8, 9:12)
   ⇒  45  30  77  24
```

Programming tip: To find the LCM of all the elements of a single array, use `num2cell` instead of nested calls or a loop:

```
x = 1:10;    # vector or array of inputs
lcm (num2cell (x) {:})
   ⇒     2520
```

**See also:** [factor], page 625, [gcd], page 625, [isprime], page 72.

`r = rem (x, y)`

Return the remainder of the division `x / y`.

The remainder is computed using the expression

```
x - y .* fix (x ./ y)
```

An error message is printed if the dimensions of the arguments do not agree, or if either argument is complex.

Programming Notes: When calculating with floating point numbers (double, single), values within a few eps of an integer will be rounded to that integer before computation for compatibility with MATLAB. Any floating point integers greater than `flintmax` (2^53 for double) will not compute correctly. For larger integer values convert the input to `uint64` before calling this function.

By convention,

```
rem (x, 0) = NaN  if x is a floating point variable
rem (x, 0) = 0    if x is an integer variable
rem (x, y)  returns a value with the signbit from x
```

For the opposite conventions see the `mod` function. In general, `rem` is best when computing the remainder after division of two *positive* numbers. For negative numbers, or when the values are periodic, `mod` is a better choice.

**See also:** [mod], page 627.

`m = mod (x, y)`

Compute the modulo of $x$ and $y$.

Conceptually this is given by

```
x - y .* floor (x ./ y)
```

and is written such that the correct modulus is returned for integer types. This function handles negative values correctly. That is, `mod (-1, 3)` is 2, not -1, as `rem (-1, 3)` returns.

An error results if the dimensions of the arguments do not agree, or if either of the arguments is complex.

Programming Notes: When calculating with floating point numbers (double, single), values within a few eps of an integer will be rounded to that integer before computation for compatibility with MATLAB. Any floating point integers greater than `flintmax` (2^53 for double) will not compute correctly. For larger integer values convert the input to `uint64` before calling this function.

By convention,

```
mod (x, 0) = x
mod (x, y)      returns a value with the signbit from y
```

For the opposite conventions see the `rem` function. In general, `mod` is a better choice than `rem` when any of the inputs are negative numbers or when the values are periodic.

**See also:** [rem], page 626.

`p = primes (n)`

Return all primes up to $n$.

The output data class (double, single, uint32, etc.) is the same as the input class of $n$. The algorithm used is the Sieve of Eratosthenes.

Note: For a specific number $n$ of primes, call `list_primes (n)`. Alternatively, call `primes (n*log (k*n))(1:n)` where $k$ is about 5 or 6. This works because the distance from one prime to the next is proportional to the logarithm of the prime, on average. On integrating, there are about $n$ primes less than `n * log (5*n)`.

**See also:** [list_primes], page 627, [isprime], page 72.

`p = list_primes ()`
`p = list_primes (n)`

List the first $n$ primes.

If $n$ is unspecified, the first 25 primes are listed.

**See also:** [primes], page 627, [isprime], page 72.

`y = sign (x)`

Compute the *signum* function.

This is defined as

$$\text{sign}(x) = \begin{cases} 1, & x > 0; \\ 0, & x = 0; \\ -1, & x < 0. \end{cases}$$

For complex arguments, `sign` returns `x ./ abs (x)`.

Note that `sign (-0.0)` is 0. Although IEEE 754 floating point allows zero to be signed, 0.0 and -0.0 compare equal. If you must test whether zero is signed, use the `signbit` function.

**See also:** [signbit], page 628.

`y = signbit (x)`

Return logical true if the value of $x$ has its sign bit set and false otherwise.

This behavior is consistent with the other logical functions. See Section 4.6 [Logical Values], page 66. The behavior differs from the C language function which returns nonzero if the sign bit is set.

This is not the same as `x < 0.0`, because IEEE 754 floating point allows zero to be signed. The comparison `-0.0 < 0.0` is false, but `signbit (-0.0)` will return a nonzero value.

**See also:** [sign], page 628.

## 17.6 Special Functions

`a = airy (z)`
`a = airy (k, z)`
`a = airy (k, z, scale)`
`[a, ierr] = airy (...)`

Compute Airy functions of the first and second kind, and their derivatives.

```
 K   Function   Scale factor (if scale is true)
---  --------   ------------------------------------
 0   Ai (Z)     exp ((2/3) * Z * sqrt (Z))
 1   dAi(Z)/dZ  exp ((2/3) * Z * sqrt (Z))
 2   Bi (Z)     exp (-abs (real ((2/3) * Z * sqrt (Z))))
 3   dBi(Z)/dZ  exp (-abs (real ((2/3) * Z * sqrt (Z))))
```

The function call `airy (z)` is equivalent to `airy (0, z)`.

The optional third input *scale* determines whether to apply scaling as described above. It is false by default.

The result $a$ is the same size as $z$.

The optional output *ierr* contains the following status information and is the same size as the result.

0. Normal return.
1. Input error, return `NaN`.

2. Overflow, return `Inf`.
3. Loss of significance by argument reduction results in less than half of machine accuracy.
4. Loss of significance by argument reduction, output may be inaccurate.
5. Error—no computation, algorithm termination condition not met, return `NaN`.

`J = besselj (alpha, x)`
`J = besselj (alpha, x, opt)`
`[J, ierr] = besselj (...)`

Compute Bessel functions of the first kind.

The order of the Bessel function *alpha* must be real. The points for evaluation *x* may be complex.

If the optional argument *opt* is 1 or true, the result *J* is multiplied by `exp (-abs (imag (x)))`.

If *alpha* is a scalar, the result is the same size as *x*. If *x* is a scalar, the result is the same size as *alpha*. If *alpha* is a row vector and *x* is a column vector, the result is a matrix with `length (x)` rows and `length (alpha)` columns. Otherwise, *alpha* and *x* must conform and the result will be the same size.

If requested, *ierr* contains the following status information and is the same size as the result.

0. Normal return.
1. Input error, return `NaN`.
2. Overflow, return `Inf`.
3. Loss of significance by argument reduction results in less than half of machine accuracy.
4. Loss of significance by argument reduction, output may be inaccurate.
5. Error—no computation, algorithm termination condition not met, return `NaN`.

**See also:** [bessely], page 629, [besseli], page 630, [besselk], page 630, [besselh], page 631.

`Y = bessely (alpha, x)`
`Y = bessely (alpha, x, opt)`
`[Y, ierr] = bessely (...)`

Compute Bessel functions of the second kind.

The order of the Bessel function *alpha* must be real. The points for evaluation *x* may be complex.

If the optional argument *opt* is 1 or true, the result *Y* is multiplied by `exp (-abs (imag (x)))`.

If *alpha* is a scalar, the result is the same size as *x*. If *x* is a scalar, the result is the same size as *alpha*. If *alpha* is a row vector and *x* is a column vector, the result is a matrix with `length (x)` rows and `length (alpha)` columns. Otherwise, *alpha* and *x* must conform and the result will be the same size.

If requested, *ierr* contains the following status information and is the same size as the result.

0. Normal return.
1. Input error, return `NaN`.
2. Overflow, return `Inf`.
3. Loss of significance by argument reduction results in less than half of machine accuracy.
4. Complete loss of significance by argument reduction, return `NaN`.
5. Error—no computation, algorithm termination condition not met, return `NaN`.

**See also:** [besselj], page 629, [besseli], page 630, [besselk], page 630, [besselh], page 631.

`I = besseli (alpha, x)`
`I = besseli (alpha, x, opt)`
`[I, ierr] = besseli (...)`

Compute modified Bessel functions of the first kind.

The order of the Bessel function *alpha* must be real. The points for evaluation *x* may be complex.

If the optional argument *opt* is 1 or true, the result *I* is multiplied by `exp (-abs (real (x)))`.

If *alpha* is a scalar, the result is the same size as *x*. If *x* is a scalar, the result is the same size as *alpha*. If *alpha* is a row vector and *x* is a column vector, the result is a matrix with `length (x)` rows and `length (alpha)` columns. Otherwise, *alpha* and *x* must conform and the result will be the same size.

If requested, *ierr* contains the following status information and is the same size as the result.

0. Normal return.
1. Input error, return `NaN`.
2. Overflow, return `Inf`.
3. Loss of significance by argument reduction results in less than half of machine accuracy.
4. Complete loss of significance by argument reduction, return `NaN`.
5. Error—no computation, algorithm termination condition not met, return `NaN`.

**See also:** [besselk], page 630, [besselj], page 629, [bessely], page 629, [besselh], page 631.

`K = besselk (alpha, x)`
`K = besselk (alpha, x, opt)`
`[K, ierr] = besselk (...)`

Compute modified Bessel functions of the second kind.

The order of the Bessel function *alpha* must be real. The points for evaluation *x* may be complex.

If the optional argument *opt* is 1 or true, the result *K* is multiplied by `exp (x)`.

If *alpha* is a scalar, the result is the same size as *x*. If *x* is a scalar, the result is the same size as *alpha*. If *alpha* is a row vector and *x* is a column vector, the result is a matrix with `length (x)` rows and `length (alpha)` columns. Otherwise, *alpha* and *x* must conform and the result will be the same size.

If requested, *ierr* contains the following status information and is the same size as the result.

0. Normal return.
1. Input error, return `NaN`.
2. Overflow, return `Inf`.
3. Loss of significance by argument reduction results in less than half of machine accuracy.
4. Complete loss of significance by argument reduction, return `NaN`.
5. Error—no computation, algorithm termination condition not met, return `NaN`.

**See also:** [besseli], page 630, [besselj], page 629, [bessely], page 629, [besselh], page 631.

`H = besselh (alpha, x)`
`H = besselh (alpha, k, x)`
`H = besselh (alpha, k, x, opt)`
`[H, ierr] = besselh (...)`

Compute Bessel functions of the third kind (Hankel functions).

The order of the Bessel function *alpha* must be real. The kind of Hankel function is specified by *k* and may be either first ($k = 1$) or second ($k = 2$). The default is Hankel functions of the first kind. The points for evaluation *x* may be complex.

If the optional argument *opt* is 1 or true, the result is multiplied by `exp (-I*x)` for $k = 1$ or `exp (I*x)` for $k = 2$.

If *alpha* is a scalar, the result is the same size as *x*. If *x* is a scalar, the result is the same size as *alpha*. If *alpha* is a row vector and *x* is a column vector, the result is a matrix with `length (x)` rows and `length (alpha)` columns. Otherwise, *alpha* and *x* must conform and the result will be the same size.

If requested, *ierr* contains the following status information and is the same size as the result.

0. Normal return.
1. Input error, return `NaN`.
2. Overflow, return `Inf`.
3. Loss of significance by argument reduction results in less than half of machine accuracy.
4. Complete loss of significance by argument reduction, return `NaN`.
5. Error—no computation, algorithm termination condition not met, return `NaN`.

**See also:** [besselj], page 629, [bessely], page 629, [besseli], page 630, [besselk], page 630.

`y = beta (a, b)`

Compute the Beta function for real inputs *a* and *b*.

The Beta function definition is

$$B(a,b) = \frac{\Gamma(a)\Gamma(b)}{\Gamma(a+b)}.$$

The Beta function can grow quite large and it is often more useful to work with the logarithm of the output rather than the function directly. See [betaln], page 633, for computing the logarithm of the Beta function in an efficient manner.

**See also:** [betaln], page 633, [betainc], page 632, [betaincinv], page 632.

`I = betainc (x, a, b)`
`I = betainc (x, a, b, tail)`

Compute the incomplete beta function.

This is defined as

$$I_x(a,b) = \frac{1}{B(a,b)} \int_0^x t^{a-1}(1-t)^{b-1} dt$$

with real $x$ in the range [0,1]. The inputs $a$ and $b$ must be real and strictly positive ($> 0$). If one of the inputs is not a scalar then the other inputs must be scalar or of compatible dimensions.

By default, *tail* is `"lower"` and the incomplete beta function integrated from 0 to $x$ is computed. If *tail* is `"upper"` then the complementary function integrated from $x$ to 1 is calculated. The two choices are related by

betainc ($x$, $a$, $b$, `"upper"`) = 1 - betainc ($x$, $a$, $b$, `"lower"`).

`betainc` uses a more sophisticated algorithm than subtraction to get numerically accurate results when the `"lower"` value is small.

Reference: A. Cuyt, V. Brevik Petersen, B. Verdonk, H. Waadeland, W.B. Jones, *Handbook of Continued Fractions for Special Functions*, ch. 18.

**See also:** [beta], page 631, [betaincinv], page 632, [betaln], page 633.

`x = betaincinv (y, a, b)`
`x = betaincinv (y, a, b, "`*lower*`")`
`x = betaincinv (y, a, b, "`*upper*`")`

Compute the inverse of the normalized incomplete beta function.

The normalized incomplete beta function is defined as

$$I_x(a,b) = \frac{1}{B(a,b)} \int_0^x t^{a-1}(1-t)^{b-1} dt$$

If two inputs are scalar, then `betaincinv (y, a, b)` is returned for each of the other inputs.

If two or more inputs are not scalar, the sizes of them must agree, and `betaincinv` is applied element-by-element.

The variable $y$ must be in the interval [0,1], while $a$ and $b$ must be real and strictly positive.

By default, *tail* is `"lower"` and the inverse of the incomplete beta function integrated from 0 to $x$ is computed. If *tail* is `"upper"` then the complementary function integrated from $x$ to 1 is inverted.

The function is computed by standard Newton's method, by solving

$$y - I_x(a, b) = 0$$

**See also:** [betainc], page 632, [beta], page 631, [betaln], page 633.

*lnb* = `betaln (a, b)`

Compute the natural logarithm of the Beta function for real inputs $a$ and $b$.

`betaln` is defined as

$$\text{betaln}(a, b) = \ln(B(a, b)) \equiv \ln(\frac{\Gamma(a)\Gamma(b)}{\Gamma(a + b)}).$$

and is calculated in a way to reduce the occurrence of underflow.

The Beta function can grow quite large and it is often more useful to work with the logarithm of the output rather than the function directly.

**See also:** [beta], page 631, [betainc], page 632, [betaincinv], page 632, [gammaln], page 639.

*b* = `bincoeff (n, k)`

Return the binomial coefficient of $n$ and $k$.

The binomial coefficient is defined as

$$\binom{n}{k} = \frac{n(n-1)(n-2)\cdots(n-k+1)}{k!}$$

For example:

```
bincoeff (5, 2)
   ⇒ 10
```

In most cases, the `nchoosek` function is faster for small scalar integer arguments. It also warns about loss of precision for big arguments.

**See also:** [nchoosek], page 860.

*k* = `commutation_matrix (m, n)`

Return the commutation matrix $K_{m,n}$ which is the unique $mn \times mn$ matrix such that $K_{m,n} \cdot \text{vec}(A) = \text{vec}(A^T)$ for all $m \times n$ matrices $A$.

If only one argument $m$ is given, $K_{m,m}$ is returned.

See Magnus and Neudecker, *Matrix Differential Calculus with Applications in Statistics and Econometrics*, 1988.

`y = cosint (x)`

Compute the cosine integral function:

$$\mathrm{Ci}(x) = -\int_x^\infty \frac{\cos(t)}{t} dt$$

An equivalent definition is

$$\mathrm{Ci}(x) = \gamma + \log(x) + \int_0^x \frac{\cos(t) - 1}{t} dt$$

Reference:

M. Abramowitz and I.A. Stegun, *Handbook of Mathematical Functions*, 1964.

**See also:** [sinint], page 639, [expint], page 636, [cos], page 607.

`d = duplication_matrix (n)`

Return the duplication matrix $D_n$ which is the unique $n^2 \times n(n+1)/2$ matrix such that $D_n * \mathrm{vech}(A) = \mathrm{vec}(A)$ for all symmetric $n \times n$ matrices $A$.

See Magnus and Neudecker, *Matrix Differential Calculus with Applications in Statistics and Econometrics*, 1988.

`v = dawson (z)`

Compute the Dawson (scaled imaginary error) function.

The Dawson function is defined as

$$\frac{\sqrt{\pi}}{2} e^{-z^2} \mathrm{erfi}(z) \equiv -i \frac{\sqrt{\pi}}{2} e^{-z^2} \mathrm{erf}(iz)$$

**See also:** [erfc], page 635, [erf], page 635, [erfcx], page 635, [erfi], page 636, [erfinv], page 636, [erfcinv], page 636.

`[sn, cn, dn, err] = ellipj (u, m)`
`[sn, cn, dn, err] = ellipj (u, m, tol)`

Compute the Jacobi elliptic functions *sn*, *cn*, and *dn* of complex argument $u$ and real parameter $m$.

If $m$ is a scalar, the results are the same size as $u$. If $u$ is a scalar, the results are the same size as $m$. If $u$ is a column vector and $m$ is a row vector, the results are matrices with `length (u)` rows and `length (m)` columns. Otherwise, $u$ and $m$ must conform in size and the results will be the same size as the inputs.

The value of $u$ may be complex. The value of $m$ must be $0 \leq m \leq 1$.

The optional input *tol* is currently ignored (MATLAB uses this to allow faster, less accurate approximation).

If requested, *err* contains the following status information and is the same size as the result.

0. Normal return.
1. Error—no computation, algorithm termination condition not met, return `NaN`.

Reference: Milton Abramowitz and Irene A Stegun, *Handbook of Mathematical Functions*, Chapter 16 (Sections 16.4, 16.13, and 16.15), Dover, 1965.

**See also:** [ellipke], page 635.

`k = ellipke (m)`
`k = ellipke (m, tol)`
`[k, e] = ellipke (...)`

Compute complete elliptic integrals of the first K($m$) and second E($m$) kind.

$m$ must be a scalar or real array with -Inf $\leq m \leq 1$.

The optional input *tol* controls the stopping tolerance of the algorithm and defaults to `eps (class (m))`. The tolerance can be increased to compute a faster, less accurate approximation.

When called with one output only elliptic integrals of the first kind are returned.

Mathematical Note:

Elliptic integrals of the first kind are defined as

$$\mathrm{K}(m) = \int_0^1 \frac{dt}{\sqrt{(1-t^2)(1-mt^2)}}$$

Elliptic integrals of the second kind are defined as

$$\mathrm{E}(m) = \int_0^1 \frac{\sqrt{1-mt^2}}{\sqrt{1-t^2}} dt$$

Reference: Milton Abramowitz and Irene A. Stegun, *Handbook of Mathematical Functions*, Chapter 17, Dover, 1965.

**See also:** [ellipj], page 634.

`v = erf (z)`

Compute the error function.

The error function is defined as

$$\mathrm{erf}(z) = \frac{2}{\sqrt{\pi}} \int_0^z e^{-t^2} dt$$

**See also:** [erfc], page 635, [erfcx], page 635, [erfi], page 636, [dawson], page 634, [erfinv], page 636, [erfcinv], page 636.

`v = erfc (z)`

Compute the complementary error function.

The complementary error function is defined as $1 - \mathrm{erf}(z)$.

**See also:** [erfcinv], page 636, [erfcx], page 635, [erfi], page 636, [dawson], page 634, [erf], page 635, [erfinv], page 636.

`v = erfcx (z)`

Compute the scaled complementary error function.

The scaled complementary error function is defined as

$$e^{z^2}\mathrm{erfc}(z) \equiv e^{z^2}(1 - \mathrm{erf}(z))$$

**See also:** [erfc], page 635, [erf], page 635, [erfi], page 636, [dawson], page 634, [erfinv], page 636, [erfcinv], page 636.

`v = erfi (z)`
Compute the imaginary error function.

The imaginary error function is defined as

$$-i\mathrm{erf}(iz)$$

**See also:** [erfc], page 635, [erf], page 635, [erfcx], page 635, [dawson], page 634, [erfinv], page 636, [erfcinv], page 636.

`y = erfinv (x)`
Compute the inverse error function.

The inverse error function is defined such that

```
erf (y) == x
```

**See also:** [erf], page 635, [erfc], page 635, [erfcx], page 635, [erfi], page 636, [dawson], page 634, [erfcinv], page 636.

`y = erfcinv (x)`
Compute the inverse complementary error function.

The inverse complementary error function is defined such that

```
erfc (y) == x
```

**See also:** [erfc], page 635, [erf], page 635, [erfcx], page 635, [erfi], page 636, [dawson], page 634, [erfinv], page 636.

`y = expint (x)`
Compute the exponential integral.

The exponential integral is defined as:

$$\mathrm{E}_1(x) = \int_x^\infty \frac{e^{-t}}{t} dt$$

Note: For compatibility, this function uses the MATLAB definition of the exponential integral. Most other sources refer to this particular value as $E_1(x)$, and the exponential integral as

$$\mathrm{Ei}(x) = -\int_{-x}^\infty \frac{e^{-t}}{t} dt.$$

The two definitions are related, for positive real values of $x$, by $E_1(-x) = -\mathrm{Ei}(x) - i\pi$.

References:

M. Abramowitz and I.A. Stegun, *Handbook of Mathematical Functions*, 1964.

N. Bleistein and R.A. Handelsman, *Asymptotic expansions of integrals*, 1986.

**See also:** [cosint], page 634, [sinint], page 639, [exp], page 603.

`v = gamma (z)`

Compute the Gamma function.

The Gamma function is defined as

$$\Gamma(z) = \int_0^\infty t^{z-1} e^{-t} dt.$$

Programming Note: The gamma function can grow quite large even for small input values. In many cases it may be preferable to use the natural logarithm of the gamma function (`gammaln`) in calculations to minimize loss of precision. The final result is then `exp (result_using_gammaln)`.

**See also:** [gammainc], page 637, [gammaln], page 639, [factorial], page 625.

`y = gammainc (x, a)`
`y = gammainc (x, a, tail)`

Compute the normalized incomplete gamma function.

This is defined as

$$\gamma(x, a) = \frac{1}{\Gamma(a)} \int_0^x t^{a-1} e^{-t} dt$$

with the limiting value of 1 as $x$ approaches infinity. The standard notation is $P(a, x)$, e.g., Abramowitz and Stegun (6.5.1).

If $a$ is scalar, then `gammainc (x, a)` is returned for each element of $x$ and vice versa.

If neither $x$ nor $a$ is scalar then the sizes of $x$ and $a$ must agree, and `gammainc` is applied element-by-element. The elements of $a$ must be non-negative.

By default, *tail* is `"lower"` and the incomplete gamma function integrated from 0 to $x$ is computed. If *tail* is `"upper"` then the complementary function integrated from $x$ to infinity is calculated.

If *tail* is `"scaledlower"`, then the lower incomplete gamma function is multiplied by $\Gamma(a+1)\exp(x)x^{-a}$. If *tail* is `"scaledupper"`, then the upper incomplete gamma function is multiplied by the same quantity.

References:

M. Abramowitz and I.A. Stegun, *Handbook of mathematical functions*, Dover publications, Inc., 1972.

W. Gautschi, "A computational procedure for incomplete gamma functions", *ACM Trans. Math Software*, Vol. 5, No. 4, pp. 466–481, 2012.

W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in Fortran 77*, Vol. 1, ch. 6.2, 1992.

**See also:** [gamma], page 637, [gammaincinv], page 637, [gammaln], page 639.

`x = gammaincinv (y, a)`
`x = gammaincinv (y, a, tail)`

Compute the inverse of the normalized incomplete gamma function.

The normalized incomplete gamma function is defined as

$$\gamma(x, a) = \frac{1}{\Gamma(a)} \int_0^x t^{a-1} e^{-t} dt$$

and `gammaincinv (gammainc (x, a), a) = x` for each non-negative value of $x$. If $a$ is scalar then `gammaincinv (y, a)` is returned for each element of $y$ and vice versa.

If neither $y$ nor $a$ is scalar then the sizes of $y$ and $a$ must agree, and `gammaincinv` is applied element-by-element. The variable $y$ must be in the interval $[0, 1]$ while $a$ must be real and positive.

By default, *tail* is `"lower"` and the inverse of the incomplete gamma function integrated from 0 to $x$ is computed. If *tail* is `"upper"`, then the complementary function integrated from $x$ to infinity is inverted.

The function is computed with Newton's method by solving

$$y - \gamma(x, a) = 0$$

Reference: A. Gil, J. Segura, and N. M. Temme, "Efficient and accurate algorithms for the computation and inversion of the incomplete gamma function ratios", *SIAM J. Sci. Computing*, Vol. 34, pp. A2965–A2981, 2012.

**See also:** [gammainc], page 637, [gamma], page 637, [gammaln], page 639.

`l = legendre (n, x)`
`l = legendre (n, x, normalization)`

Compute the associated Legendre function of degree $n$ and order $m = 0 \ldots n$.

The value $n$ must be a real non-negative integer.

$x$ is a vector with real-valued elements in the range [-1, 1].

The optional argument *normalization* may be one of `"unnorm"`, `"sch"`, or `"norm"`. The default if no normalization is given is `"unnorm"`.

When the optional argument *normalization* is `"unnorm"`, compute the associated Legendre function of degree $n$ and order $m$ and return all values for $m = 0 \ldots n$. The return value has one dimension more than $x$.

The associated Legendre function of degree $n$ and order $m$:

$$P_n^m(x) = (-1)^m (1 - x^2)^{m/2} \frac{d^m}{dx^m} P_n(x)$$

with Legendre polynomial of degree $n$:

$$P(x) = \frac{1}{2^n n!} \left( \frac{d^n}{dx^n} (x^2 - 1)^n \right)$$

`legendre (3, [-1.0, -0.9, -0.8])` returns the matrix:

```
   x  |   -1.0   |   -0.9   |   -0.8
  ------------------------------------
  m=0 | -1.00000 | -0.47250 | -0.08000
  m=1 |  0.00000 | -1.99420 | -1.98000
  m=2 |  0.00000 | -2.56500 | -4.32000
  m=3 |  0.00000 | -1.24229 | -3.24000
```

When the optional argument *normalization* is `"sch"`, compute the Schmidt semi-normalized associated Legendre function. The Schmidt semi-normalized associated Legendre function is related to the unnormalized Legendre functions by the following:

For Legendre functions of degree $n$ and order 0:

$$SP_n^0(x) = P_n^0(x)$$

For Legendre functions of degree n and order m:

$$SP_n^m(x) = P_n^m(x)(-1)^m \left(\frac{2(n-m)!}{(n+m)!}\right)^{0.5}$$

When the optional argument *normalization* is `"norm"`, compute the fully normalized associated Legendre function. The fully normalized associated Legendre function is related to the unnormalized associated Legendre functions by the following:

For Legendre functions of degree $n$ and order $m$

$$NP_n^m(x) = P_n^m(x)(-1)^m \left(\frac{(n+0.5)(n-m)!}{(n+m)!}\right)^{0.5}$$

`y = gammaln (x)`
`y = lgamma (x)`

Return the natural logarithm of the gamma function of $x$.

Programming Note: `lgamma` is an alias for `gammaln` and either name can be used in Octave.

**See also:** [gamma], page 637, [gammainc], page 637.

`y = psi (z)`
`y = psi (k, z)`

Compute the psi (polygamma) function.

The polygamma functions are the $k$th derivative of the logarithm of the gamma function. If unspecified, $k$ defaults to zero. A value of zero computes the digamma function, a value of 1, the trigamma function, and so on.

The digamma function is defined:

$$\Psi(z) = \frac{d(log(\Gamma(z)))}{dx}$$

When computing the digamma function (when $k$ equals zero), $z$ can have any value real or complex value. However, for polygamma functions ($k$ higher than 0), $z$ must be real and non-negative.

**See also:** [gamma], page 637, [gammainc], page 637, [gammaln], page 639.

`y = sinint (x)`

Compute the sine integral function:

$$\mathrm{Si}(x) = \int_0^x \frac{\sin(t)}{t} dt$$

Reference: M. Abramowitz and I.A. Stegun, *Handbook of Mathematical Functions*, 1964.

**See also:** [cosint], page 634, [expint], page 636, [sin], page 607.

## 17.7 Rational Approximations

`s = rat (x)`
`s = rat (x, tol)`
`[n, d] = rat (...)`

Find a rational approximation of $x$ to within the tolerance defined by *tol*.

If unspecified, the default tolerance is `1e-6 * norm (x(:), 1)`.

When called with one output argument, return a string containing a continued fraction expansion (multiple terms).

When called with two output arguments, return numeric matrices for the numerator and denominator of a fractional representation of $x$ such that `x = n ./ d`.

For example:

```
s = rat (pi)
⇒ s = 3 + 1/(7 + 1/16)

[n, d] = rat (pi)
⇒ n =  355
⇒ d =  113

n / d - pi
⇒ 2.6676e-07
```

Complex inputs are similar:

```
s = rat (0.5 + i * pi)
⇒ s = complex (1 + 1/(-2), 3 + 1/(7 + 1/16))

[n, d] = rat (0.5 + i * pi)
⇒ n =  113 + 710i
⇒ d =  226

n / d - (0.5 + i * pi)
⇒ 0 + 2.6676e-07i
```

Programming Notes:

1. With one output `rat` produces a string which is a continued fraction expansion. To produce a string which is a simple fraction (one numerator, one denominator) use `rats`.

2. The string output produced by `rat` can be passed to `eval` to get back the original input up to the tolerance used.

**See also:** [rats], page 640, [format], page 288.

`s = rats (x)`
`s = rats (x, len)`

Convert $x$ into a rational approximation represented as a string.

A rational approximation to a floating point number is a simple fraction with numerator $N$ and denominator $D$ such that `x = N/D`.

The optional second argument defines the maximum length of the string representing the elements of *x*. By default, *len* is 13.

If the length of the smallest possible rational approximation exceeds *len*, an asterisk (*) padded with spaces will be returned instead.

Example conversion from matrix to string, and back again.

```
r = rats (hilb (4));
x = str2num (r)
```

**See also:** [rat], page 640, [format], page 288.

## 17.8 Coordinate Transformations

```
[theta, r] = cart2pol (x, y)
[theta, r, z] = cart2pol (x, y, z)
[theta, r] = cart2pol (C)
[theta, r, z] = cart2pol (C)
```

Transform Cartesian coordinates to polar or cylindrical coordinates.

The inputs *x*, *y* (, and *z*) must be the same shape, or scalar. If called with a single matrix argument then each row of *C* represents the Cartesian coordinate pair (*x*, *y*) or triplet (*x*, *y*, *z*).

The outputs *theta*, *r* (, and *z*) match the shape of the inputs. For a matrix input *C* the outputs will be column vectors with rows corresponding to the rows of the input matrix.

*theta* describes the angle relative to the positive x-axis measured in the xy-plane.

*r* is the distance to the z-axis (0, 0, z).

*z*, if present, is unchanged by the transformation.

The coordinate transformation is computed using:

$$\theta = \arctan\left(\frac{y}{x}\right)$$

$$r = \sqrt{x^2 + y^2}$$

$$z = z$$

**See also:** [pol2cart], page 641, [cart2sph], page 642, [sph2cart], page 642.

```
[x, y] = pol2cart (theta, r)
[x, y, z] = pol2cart (theta, r, z)
[x, y] = pol2cart (P)
[x, y, z] = pol2cart (P)
```

Transform polar or cylindrical coordinates to Cartesian coordinates.

The inputs *theta*, *r*, (and *z*) must be the same shape, or scalar. If called with a single matrix argument then each row of *P* represents the polar coordinate pair (*theta*, *r*) or the cylindrical triplet (*theta*, *r*, *z*).

The outputs $x$, $y$ (, and $z$) match the shape of the inputs. For a matrix input $P$ the outputs will be column vectors with rows corresponding to the rows of the input matrix.

*theta* describes the angle relative to the positive x-axis measured in the xy-plane.

$r$ is the distance to the z-axis (0, 0, z).

$z$, if present, is unchanged by the transformation.

The coordinate transformation is computed using:

$$x = r \cos \theta$$

$$y = r \sin \theta$$

$$z = z$$

**See also:** [cart2pol], page 641, [sph2cart], page 642, [cart2sph], page 642.

`[theta, phi, r] = cart2sph (x, y, z)`
`[theta, phi, r] = cart2sph (C)`

Transform Cartesian coordinates to spherical coordinates.

The inputs $x$, $y$, and $z$ must be the same shape, or scalar. If called with a single matrix argument then each row of $C$ must represent a Cartesian coordinate triplet ($x$, $y$, $z$).

The outputs *theta*, *phi*, $r$ match the shape of the inputs. For a matrix input $C$ the outputs will be column vectors with rows corresponding to the rows of the input matrix.

*theta* describes the azimuth angle relative to the positive x-axis measured in the xy-plane.

*phi* is the elevation angle measured relative to the xy-plane.

$r$ is the distance to the origin (0, 0, 0).

The coordinate transformation is computed using:

$$\theta = \arctan\left(\frac{y}{x}\right)$$

$$\phi = \arctan\left(\frac{z}{\sqrt{x^2 + y^2}}\right)$$

$$r = \sqrt{x^2 + y^2 + z^2}$$

**See also:** [sph2cart], page 642, [cart2pol], page 641, [pol2cart], page 641.

`[x, y, z] = sph2cart (theta, phi, r)`
`[x, y, z] = sph2cart (S)`

Transform spherical coordinates to Cartesian coordinates.

The inputs *theta*, *phi*, and $r$ must be the same shape, or scalar. If called with a single matrix argument then each row of $S$ must represent a spherical coordinate triplet (*theta*, *phi*, $r$).

The outputs *x*, *y*, *z* match the shape of the inputs. For a matrix input *S* the outputs are column vectors with rows corresponding to the rows of the input matrix.

*theta* describes the azimuth angle relative to the positive x-axis measured in the xy-plane.

*phi* is the elevation angle measured relative to the xy-plane.

*r* is the distance to the origin (0, 0, 0).

The coordinate transformation is computed using:

$$x = r \cos\phi \cos\theta$$

$$y = r \cos\phi \sin\theta$$

$$z = r \sin\phi$$

**See also:** [cart2sph], page 642, [pol2cart], page 641, [cart2pol], page 641.

## 17.9 Mathematical Constants

`x = e`
`x = e (n)`
`x = e (m, n, ...)`
`x = e ([m, n, ...])`
`x = e (..., class)`

Return a scalar, matrix, or N-dimensional array whose elements are all equal to the base of natural logarithms.

The constant *e* satisfies the equation $\log(e) = 1$.

If called with no arguments, return the scalar value *e*.

If invoked with a single scalar integer argument *n*, return a square NxN matrix.

If invoked with two or more scalar integer arguments, or a vector of integer values, return an array with the given dimensions.

The optional argument *class* specifies the class of the return array. The only valid options are `"double"` (default) or `"single"`.

**See also:** [log], page 603, [exp], page 603, [pi], page 643, [I], page 644.

`p = pi`
`p = pi (n)`
`p = pi (m, n, ...)`
`p = pi ([m, n, ...])`
`p = pi (..., class)`

Return a scalar, matrix, or N-dimensional array whose elements are all equal to the ratio of the circumference of a circle to its diameter($\pi$).

If called with no arguments, return the scalar value $\pi$.

If invoked with a single scalar integer argument *n*, return a square NxN matrix.

If invoked with two or more scalar integer arguments, or a vector of integer values, return an array with the given dimensions.

The optional argument *class* specifies the class of the return array. The only valid options are `"double"` (default) or `"single"`.

**See also:** [e], page 643, [I], page 644.

`x = I`
`x = I (n)`
`x = I (m, n, ...)`
`x = I ([m, n, ...])`
`x = I (..., class)`

Return a scalar, matrix, or N-dimensional array whose elements are all equal to the pure imaginary unit, defined as $\sqrt{-1}$.

`I`, and its equivalents `i`, `j`, and `J`, are functions so any of the names may be reused for other purposes (such as `i` for a counter variable).

If called with no arguments, return the scalar value `complex (0, 1)`.

If invoked with a single scalar integer argument *n*, return a square NxN matrix.

If invoked with two or more scalar integer arguments, or a vector of integer values, return an array with the given dimensions.

The optional argument *class* specifies the class of the return array. The only valid options are `"double"` (default) or `"single"`.

**See also:** [e], page 643, [pi], page 643, [log], page 603, [exp], page 603.

`x = Inf`
`x = Inf (n)`
`x = Inf (m, n, ...)`
`x = Inf ([m, n, ...])`
`x = Inf (..., class)`
`x = Inf (..., "like", var)`

Return a scalar, matrix or N-dimensional array whose elements are all equal to the IEEE 754 representation for positive infinity.

Infinity is produced when results are too large to be represented using the IEEE 754 floating point format for numbers. Two common examples which produce infinity are division by zero and overflow.

```
[ 1/0 e^800 ]
⇒ Inf   Inf
```

If called with no arguments, return the scalar value `Inf`.

If invoked with a single scalar integer argument *n*, return a square NxN matrix.

If invoked with two or more scalar integer arguments, or a vector of integer values, return an array with the given dimensions.

The optional argument *class* specifies the class of the return array. The only valid options are `"double"` (default) or `"single"`.

If a variable *var* is specified after `"like"`, the output *x* will have the same data type, complexity, and sparsity as *var*.

**See also:** [isinf], page 567, [NaN], page 645.

`x = NaN`
`x = NaN (n)`
`x = NaN (m, n, ...)`
`x = NaN ([m, n, ...])`
`x = NaN (..., class)`
`x = NaN (..., "like", var)`

Return a scalar, matrix, or N-dimensional array whose elements are all equal to the IEEE 754 symbol NaN (Not a Number).

`NaN` is the result of operations which do not produce a well defined numerical result. Common operations which produce a `NaN` are arithmetic with infinity ($\infty - \infty$), zero divided by zero (0/0), and any operation involving another `NaN` value (`5 + NaN`).

Note that `NaN` always compares not equal to `NaN` (`NaN != NaN`). This behavior is specified by the IEEE 754 standard for floating point arithmetic. To find `NaN` values, use the `isnan` function.

If called with no arguments, return the scalar value `NaN`.

If invoked with a single scalar integer argument *n*, return a square NxN matrix.

If invoked with two or more scalar integer arguments, or a vector of integer values, return an array with the given dimensions.

The optional argument *class* specifies the class of the return array. The only valid options are `"double"` (default) or `"single"`.

If a variable *var* is specified after `"like"`, the output *x* will have the same data type, complexity, and sparsity as *var*.

**See also:** [isnan], page 567, [Inf], page 644.

`d = eps`
`d = eps (x)`
`d = eps (m, n, ...)`
`d = eps ([m, n, ...])`
`d = eps (..., class)`

Return a scalar, matrix or N-dimensional array whose elements are eps, the machine precision.

More precisely, `eps` is the relative spacing between any two adjacent numbers in the machine's floating point system. This number depends both on the system and where the number lies in the range representable by the floating point system. On machines that support IEEE 754 floating point arithmetic, `eps (1.0)` is approximately $2.2204 \times 10^{-16}$ for double precision and $1.1921 \times 10^{-7}$ for single precision.

If called with no arguments, return the scalar value `eps (1.0)`.

Given a floating point argument *x*, return an array *d* of the same size where each element is the distance between the element of *x* and the next largest value.

If invoked with two or more scalar integer arguments, or a vector of integer values, return an array with the given dimensions whose elements are all the scalar value `eps`.

The optional argument *class* specifies the class of the return array. The only valid options are `"double"` (default) or `"single"`.

**See also:** [realmax], page 646, [realmin], page 646, [intmax], page 60, [flintmax], page 61.

*Rmax* = realmax
*Rmax* = realmax (*n*)
*Rmax* = realmax (*m*, *n*, ...)
*Rmax* = realmax ([*m*, *n*, ...])
*Rmax* = realmax (..., *class*)
*Rmax* = realmax (..., "*like*", *var*)

Return a scalar, matrix, or N-dimensional array whose elements are all equal to the largest floating point number that is representable.

The actual value is system-dependent. On machines that support IEEE 754 floating point arithmetic, `realmax` is approximately $1.7977 \times 10^{308}$ for double precision and $3.4028 \times 10^{38}$ for single precision.

If called with no arguments, return the scalar value `realmax ("double")`.

If invoked with a single scalar integer argument *n*, return a square NxN matrix.

If invoked with two or more scalar integer arguments, or a vector of integer values, return an array with the given dimensions.

The optional argument *class* specifies the class of the return array. The only valid options are `"double"` (default) or `"single"`.

If a variable *var* is specified after `"like"`, the output *Rmax* will have the same data type, complexity, and sparsity as *var*.

**See also:** [realmin], page 646, [intmax], page 60, [flintmax], page 61, [eps], page 645.

*Rmin* = realmin
*Rmin* = realmin (*n*)
*Rmin* = realmin (*m*, *n*, ...)
*Rmin* = realmin ([*m*, *n*, ...])
*Rmin* = realmin (..., *class*)
*Rmin* = realmin (..., "*like*", *var*)

Return a scalar, matrix, or N-dimensional array whose elements are all equal to the smallest normalized floating point number that is representable.

The actual value is system-dependent. On machines that support IEEE 754 floating point arithmetic, `realmin` is approximately $2.2251 \times 10^{-308}$ for double precision and $1.1755 \times 10^{-38}$ for single precision.

If called with no arguments, return the scalar value `realmin ("double")`.

If invoked with a single scalar integer argument *n*, return a square NxN matrix.

If invoked with two or more scalar integer arguments, or a vector of integer values, return an array with the given dimensions.

The optional argument *class* specifies the class of the return array. The only valid options are `"double"` (default) or `"single"`.

If a variable *var* is specified after `"like"`, the output *Rmin* will have the same data type, complexity, and sparsity as *var*.

**See also:** [realmax], page 646, [intmin], page 61, [eps], page 645.

# 18 Linear Algebra

This chapter documents the linear algebra functions provided in Octave. Reference material for many of these functions may be found in Golub and Van Loan, *Matrix Computations, 2nd Ed.*, Johns Hopkins, 1989, and in the LAPACK *Users' Guide*, SIAM, 1992. The LAPACK *Users' Guide* is available at: *http://www.netlib.org/lapack/lug/*.

A common text for engineering courses is G. Strang, *Linear Algebra and Its Applications, 4th Edition*. It has become a widespread reference for linear algebra. An alternative is P. Lax, *Linear Algebra and Its Applications*, and also is a good choice. It claims to be suitable for high school students with substantial mathematical interests as well as first-year undergraduates.

## 18.1 Techniques Used for Linear Algebra

Octave includes a polymorphic solver that selects an appropriate matrix factorization depending on the properties of the matrix itself. Generally, the cost of determining the matrix type is small relative to the cost of factorizing the matrix itself. In any case the matrix type is cached once it is calculated so that it is not re-determined each time it is used in a linear equation.

The selection tree for how the linear equation is solved or a matrix inverse is formed is given by:

1. If the matrix is upper or lower triangular sparse use a forward or backward substitution using the LAPACK xTRTRS function, and goto 4.
2. If the matrix is square, Hermitian with a real positive diagonal, attempt Cholesky factorization using the LAPACK xPOTRF function.
3. If the Cholesky factorization failed or the matrix is not Hermitian with a real positive diagonal, and the matrix is square, factorize using the LAPACK xGETRF function.
4. If the matrix is not square, or any of the previous solvers flags a singular or near singular matrix, find a least squares solution using the LAPACK xGELSD function.

The user can force the type of the matrix with the `matrix_type` function. This overcomes the cost of discovering the type of the matrix. However, it should be noted that identifying the type of the matrix incorrectly will lead to unpredictable results, and so `matrix_type` should be used with care.

It should be noted that the test for whether a matrix is a candidate for Cholesky factorization, performed above, and by the `matrix_type` function, does not make certain that the matrix is Hermitian. However, the attempt to factorize the matrix will quickly detect a non-Hermitian matrix.

## 18.2 Basic Matrix Functions

```
AA = balance (A)
AA = balance (A, opt)
[DD, AA] = balance (A, opt)
[D, P, AA] = balance (A, opt)
```

`[CC, DD, AA, BB] = balance (A, B, opt)`

Balance the matrix $A$ to reduce numerical errors in future calculations.

Compute `AA = DD \ A * DD` in which $AA$ is a matrix whose row and column norms are roughly equal in magnitude, and `DD = P * D`, in which $P$ is a permutation matrix and $D$ is a diagonal matrix of powers of two. This allows the equilibration to be computed without round-off. Results of eigenvalue calculation are typically improved by balancing first.

If two output values are requested, `balance` returns the diagonal $D$ and the permutation $P$ separately as vectors. In this case, `DD = eye(n)(:,P) * diag (D)`, where $n$ is the matrix size.

If four output values are requested, compute `AA = CC*A*DD` and `BB = CC*B*DD`, in which $AA$ and $BB$ have nonzero elements of approximately the same magnitude and $CC$ and $DD$ are permuted diagonal matrices as in $DD$ for the algebraic eigenvalue problem.

The eigenvalue balancing option *opt* may be one of:

`"noperm"`, `"S"`
: Scale only; do not permute.

`"noscal"`, `"P"`
: Permute only; do not scale.

Algebraic eigenvalue balancing uses standard LAPACK routines.

Generalized eigenvalue problem balancing uses Ward's algorithm (SIAM Journal on Scientific and Statistical Computing, 1981).

`bw = bandwidth (A, type)`
`[lower, upper] = bandwidth (A)`

Compute the bandwidth of $A$.

The *type* argument is the string `"lower"` for the lower bandwidth and `"upper"` for the upper bandwidth. If no *type* is specified return both the lower and upper bandwidth of $A$.

The lower/upper bandwidth of a matrix is the number of subdiagonals/superdiagonals with nonzero entries.

**See also:** [isbanded], page 71, [isdiag], page 71, [istril], page 72, [istriu], page 72.

`c = cond (A)`
`c = cond (A, p)`

Compute the $p$-norm condition number of a matrix with respect to inversion.

`cond (A)` is defined as $\| A \|_p * \| A^{-1} \|_p$ .

By default, `p = 2` is used which implies a (relatively slow) singular value decomposition. Other possible selections are `p = 1, Inf, "fro"` which are generally faster. For a full discussion of possible $p$ values, see [norm], page 654.

The condition number of a matrix quantifies the sensitivity of the matrix inversion operation when small changes are made to matrix elements. Ideally the condition number will be close to 1. When the number is large this indicates small changes (such as underflow or round-off error) will produce large changes in the resulting

output. In such cases the solution results from numerical computing are not likely to be accurate.

**See also:** [condest], page 751, [rcond], page 656, [condeig], page 649, [norm], page 654, [svd], page 669.

`c = condeig (a)`
`[v, lambda, c] = condeig (a)`

Compute condition numbers of a matrix with respect to eigenvalues.

The condition numbers are the reciprocals of the cosines of the angles between the left and right eigenvectors; Large values indicate that the matrix has multiple distinct eigenvalues.

The input *a* must be a square numeric matrix.

The outputs are:

- *c* is a vector of condition numbers for the eigenvalues of *a*.
- *v* is the matrix of right eigenvectors of *a*. The result is equivalent to calling `[v, lambda] = eig (a)`.
- *lambda* is the diagonal matrix of eigenvalues of *a*. The result is equivalent to calling `[v, lambda] = eig (a)`.

Example

```
a = [1, 2; 3, 4];
c = condeig (a)
  ⇒ c =
       1.0150
       1.0150
```

**See also:** [eig], page 649, [cond], page 648, [balance], page 647.

`d = det (A)`
`[d, rcond] = det (A)`

Compute the determinant of *A*.

Return an estimate of the reciprocal condition number if requested.

Programming Notes: Routines from LAPACK are used for full matrices and code from UMFPACK is used for sparse matrices.

The determinant should not be used to check a matrix for singularity. For that, use any of the condition number functions: `cond`, `condest`, `rcond`.

**See also:** [cond], page 648, [condest], page 751, [rcond], page 656.

`lambda = eig (A)`
`lambda = eig (A, B)`
`[V, lambda] = eig (A)`
`[V, lambda] = eig (A, B)`
`[V, lambda, W] = eig (A)`
`[V, lambda, W] = eig (A, B)`
`[...] = eig (A, balanceOption)`
`[...] = eig (A, B, algorithm)`

`[...] = eig (..., eigvalOption)`

Compute the eigenvalues (*lambda*) and optionally the right eigenvectors ($V$) and the left eigenvectors ($W$) of a matrix or pair of matrices.

The flag *balanceOption* can be one of:

`"balance"` (default)
: Preliminary balancing is on.

`"nobalance"`
: Disables preliminary balancing.

The flag *eigvalOption* can be one of:

`"matrix"` Return the eigenvalues in a diagonal matrix. (default if 2 or 3 outputs are requested)

`"vector"` Return the eigenvalues in a column vector. (default if only 1 output is requested, e.g., `lambda = eig (A)`)

The flag *algorithm* can be one of:

`"chol"` Use the Cholesky factorization of B. (default if $A$ is symmetric (Hermitian) and $B$ is symmetric (Hermitian) positive definite)

`"qz"` Use the QZ algorithm. (used whenever $A$ or $B$ are not symmetric)

| **A and B** | **no flag** | **chol** | **qz** |
|---|---|---|---|
| both are symmetric | `"chol"` | `"chol"` | `"qz"` |
| at least one is not symmetric | `"qz"` | `"qz"` | `"qz"` |

The eigenvalues returned by `eig` are not ordered.

**See also:** [eigs], page 754, [svd], page 669.

`G = givens (x, y)`
`[c, s] = givens (x, y)`

Compute the Givens rotation matrix $G$. The Givens matrix is a $2 \times 2$ orthogonal matrix

$$G = \begin{bmatrix} c & s \\ -s' & c \end{bmatrix}$$

such that

$$G \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} * \\ 0 \end{bmatrix}$$

with $x$ and $y$ scalars.

If two output arguments are requested, return the factors $c$ and $s$ rather than the Givens rotation matrix.

For example:

```
givens (1, 1)
   ⇒   0.70711   0.70711
      -0.70711   0.70711
```

Note: The Givens matrix represents a counterclockwise rotation of a 2-D plane and can be used to introduce zeros into a matrix prior to complete factorization.

**See also:** [planerot], page 651, [qr], page 661.

`S = gsvd (A, B)`
`[U, V, X, C, S] = gsvd (A, B)`
`[U, V, X, C, S] = gsvd (A, B, 0)`

Compute the generalized singular value decomposition of $(A, B)$.

The generalized singular value decomposition is defined by the following relations:

$$A = UCX^\dagger$$

$$B = VSX^\dagger$$

$$C^\dagger C + S^\dagger S = eye(columns(A))$$

The function `gsvd` normally returns just the vector of generalized singular values

$$\sqrt{\frac{diag(C^\dagger C)}{diag(S^\dagger S)}}$$

If asked for five return values, it also computes $U$, $V$, $X$, and $C$.

If the optional third input is present, `gsvd` constructs the "economy-sized" decomposition where the number of columns of $U$, $V$ and the number of rows of $C$, $S$ is less than or equal to the number of columns of $A$. This option is not yet implemented.

Programming Note: the code is a wrapper to the corresponding LAPACK dggsvd and zggsvd routines. If matrices $A$ and $B$ are *both* rank deficient then LAPACK will return an incorrect factorization. Programmers should avoid this combination.

**See also:** [svd], page 669.

`[G, y] = planerot (x)`

Compute the Givens rotation matrix for the two-element column vector $x$. The Givens matrix is a $2 \times 2$ orthogonal matrix

$$G = \begin{bmatrix} c & s \\ -s' & c \end{bmatrix}$$

such that

$$G \begin{bmatrix} x(1) \\ x(2) \end{bmatrix} = \begin{bmatrix} * \\ 0 \end{bmatrix}$$

Note: The Givens matrix represents a counterclockwise rotation of a 2-D plane and can be used to introduce zeros into a matrix prior to complete factorization.

**See also:** [givens], page 650, [qr], page 661.

`x = inv (A)`
`[x, rcond] = inv (A)`
`[...] = inverse (...)`

Compute the inverse of the square matrix $A$.

Return an estimate of the reciprocal condition number if requested, otherwise warn of an ill-conditioned matrix if the reciprocal condition number is small.

In general it is best to avoid calculating the inverse of a matrix directly. For example, it is both faster and more accurate to solve systems of equations ($A*x = b$) with `y = A \ b`, rather than `y = inv (A) * b`.

If called with a sparse matrix, then in general $x$ will be a full matrix requiring significantly more storage. Avoid forming the inverse of a sparse matrix if possible.

Programming Note: `inverse` is an alias for `inv` and can be used interchangeably.

**See also:** [ldivide], page 172, [rdivide], page 173, [pinv], page 655.

`x = linsolve (A, b)`
`x = linsolve (A, b, opts)`
`[x, R] = linsolve (...)`

Solve the linear system `A*x = b`.

With no options, this function is equivalent to the left division operator (`x = A \ b`) or the matrix-left-divide function (`x = mldivide (A, b)`).

Octave ordinarily examines the properties of the matrix $A$ and chooses a solver that best matches the matrix. By passing a structure *opts* to `linsolve` you can inform Octave directly about the matrix $A$. In this case Octave will skip the matrix examination and proceed directly to solving the linear system producing a small speedup.

**Warning:** If the matrix $A$ does not have the properties listed in the *opts* structure then the result will not be accurate AND no warning will be given. When in doubt, let Octave examine the matrix and choose the appropriate solver as this step takes little time and the result is cached so that it is done just once per linear system. Benchmarking with Octave 12 and square matrices of 10,000 rows shows just 300 milliseconds for this step.

Possible *opts* fields (set value to `true`/`false`):

| | |
|---|---|
| `LT` | $A$ is lower triangular |
| `UT` | $A$ is upper triangular |
| `UHESS` | $A$ is upper Hessenberg (currently makes no difference) |
| `SYM` | $A$ is symmetric or complex Hermitian (currently makes no difference) |
| `POSDEF` | $A$ is positive definite |
| `RECT` | $A$ is general rectangular (currently makes no difference) |
| `TRANSA` | Solve `A'*x = b` if true rather than `A*x = b` |

The optional second output $R$ is the reciprocal condition number of $A$ for square matrices or `0` for rectangular matrices.

**See also:** [mldivide], page 172, [matrix_type], page 652, [rcond], page 656.

`type = matrix_type (A)`
`type = matrix_type (A, "nocompute")`
`A = matrix_type (A, type)`
`A = matrix_type (A, "upper", perm)`
`A = matrix_type (A, "lower", perm)`

*A* = `matrix_type` (*A*, "*banded*", *nl*, *nu*)

Identify the matrix type or mark a matrix as a particular type.

This allows more rapid solutions of linear equations involving *A* to be performed.

Called with a single argument, `matrix_type` returns the type of the matrix and caches it for future use.

Called with more than one argument, `matrix_type` allows the type of the matrix to be defined.

If the option `"nocompute"` is given, the function will not attempt to guess the type if it is still unknown. This is useful for debugging purposes.

The possible matrix types depend on whether the matrix is full or sparse, and can be one of the following

`"unknown"`
: Remove any previously cached matrix type, and mark type as unknown.

`"full"`
: Mark the matrix as full.

`"positive definite"`
: Probable full positive definite matrix.

`"diagonal"`
: Diagonal matrix. (Sparse matrices only)

`"permuted diagonal"`
: Permuted Diagonal matrix. The permutation does not need to be specifically indicated, as the structure of the matrix explicitly gives this. (Sparse matrices only)

`"upper"`
: Upper triangular. If the optional third argument *perm* is given, the matrix is assumed to be a permuted upper triangular with the permutations defined by the vector *perm*.

`"lower"`
: Lower triangular. If the optional third argument *perm* is given, the matrix is assumed to be a permuted lower triangular with the permutations defined by the vector *perm*.

`"banded"`
`"banded positive definite"`
: Banded matrix with the band size of *nl* below the diagonal and *nu* above it. If *nl* and *nu* are 1, then the matrix is tridiagonal and treated with specialized code. In addition the matrix can be marked as probably a positive definite. (Sparse matrices only)

`"singular"`
: The matrix is assumed to be singular and will be treated with a minimum norm solution.

Note that the matrix type will be discovered automatically on the first attempt to solve a linear equation involving *A*. Therefore `matrix_type` is only useful to give Octave hints of the matrix type. Incorrectly defining the matrix type will result in incorrect results from solutions of linear equations; it is entirely **the responsibility of the user** to correctly identify the matrix type.

Also, the test for positive definiteness is a low-cost test for a Hermitian matrix with a real positive diagonal. This does not guarantee that the matrix is positive definite, but only that it is a probable candidate. When such a matrix is factorized, a Cholesky factorization is first attempted, and if that fails the matrix is then treated with an LU factorization. Once the matrix has been factorized, `matrix_type` will return the correct classification of the matrix.

`n = norm (A)`
`n = norm (A, p)`
`n = norm (A, p, opt)`

Compute the p-norm of the matrix $A$.

If the second argument is not given, `p = 2` is used.

If $A$ is a matrix (or sparse matrix):

$p = 1$ 1-norm, the largest column sum of the absolute values of $A$.

$p = 2$ Largest singular value of $A$.

$p =$ `Inf` or `"inf"`
Infinity norm, the largest row sum of the absolute values of $A$.

$p =$ `"fro"`
Frobenius norm of $A$, `sqrt (sum (diag (A' * A)))`.

other $p$, $p >$ `1`
maximum `norm (A*x, p)` such that `norm (x, p) == 1`

If $A$ is a vector or a scalar:

$p =$ `Inf` or `"inf"`
`max (abs (A))`.

$p =$ `-Inf` `min (abs (A))`.

$p =$ `"fro"`
Frobenius norm of $A$, `sqrt (sumsq (abs (A)))`.

$p = 0$ Hamming norm—the number of nonzero elements.

other $p$, $p >$ `1`
p-norm of $A$, `(sum (abs (A) .^ p)) ^ (1/p)`.

other $p$ $p <$ `1`
the p-pseudonorm defined as above.

If *opt* is the value `"rows"`, treat each row as a vector and compute its norm. The result is returned as a column vector. Similarly, if *opt* is `"columns"` or `"cols"` then compute the norms of each column and return a row vector.

**See also:** [normest], page 750, [normest1], page 750, [vecnorm], page 656, [cond], page 648, [svd], page 669.

`Z = null (A)`
`Z = null (A, tol)`

Return an orthonormal basis $Z$ of the null space of $A$.

The dimension of the null space $Z$ is taken as the number of singular values of $A$ not greater than *tol*. If the argument *tol* is missing, it is computed as

```
max (size (A)) * max (svd (A, 0)) * eps
```

**See also:** [orth], page 655, [svd], page 669.

`B = orth (A)`
`B = orth (A, tol)`

Return an orthonormal basis of the range space of $A$.

The dimension of the range space is taken as the number of singular values of $A$ greater than *tol*. If the argument *tol* is missing, it is computed as

```
max (size (A)) * max (svd (A)) * eps
```

**See also:** [null], page 654.

`[y, h] = mgorth (x, v)`

Orthogonalize a given column vector $x$ with respect to a set of orthonormal vectors comprising the columns of $v$ using the modified Gram-Schmidt method.

On exit, $y$ is a unit vector such that:

```
norm (y) = 1
v' * y = 0
x = [v, y]*h'
```

`B = pinv (A)`
`B = pinv (A, tol)`

Return the Moore-Penrose pseudoinverse of $A$.

Singular values less than *tol* are ignored.

If the second argument is omitted, it is taken to be

```
tol = max ([rows(x), columns(x)]) * norm (x) * eps
```

**See also:** [inv], page 651, [ldivide], page 172.

`k = rank (A)`
`k = rank (A, tol)`

Compute the rank of matrix $A$, using the singular value decomposition.

The rank is taken to be the number of singular values of $A$ that are greater than the specified tolerance *tol*. If the second argument is omitted, it is taken to be

```
tol = max (size (A)) * sigma(1) * eps;
```

where `eps` is machine precision and `sigma(1)` is the largest singular value of $A$.

The rank of a matrix is the number of linearly independent rows or columns and equals the dimension of the row and column space. The function `orth` may be used to compute an orthonormal basis of the column space.

For testing if a system `A*x = b` of linear equations is solvable, one can use

```
rank (A) == rank ([A b])
```

In this case, `x = A \ b` finds a particular solution $x$. The general solution is $x$ plus the null space of matrix $A$. The function `null` may be used to compute a basis of the null space.

Example:

```
A = [1 2 3
     4 5 6
     7 8 9];
rank (A)
  ⇒ 2
```

In this example, the number of linearly independent rows is only 2 because the final row is a linear combination of the first two rows:

```
A(3,:) == -A(1,:) + 2 * A(2,:)
```

**See also:** [null], page 654, [orth], page 655, [sprank], page 753, [svd], page 669, [eps], page 645.

`c = rcond (A)`

Compute the 1-norm estimate of the reciprocal condition number as returned by LAPACK.

If the matrix is well-conditioned then $c$ will be near 1 and if the matrix is poorly conditioned it will be close to 0.

The matrix $A$ must not be sparse. If the matrix is sparse then `condest (A)` or `rcond (full (A))` should be used instead.

**See also:** [cond], page 648, [condest], page 751.

`t = trace (A)`

Compute the trace of $A$, the sum of the elements along the main diagonal.

The implementation is straightforward: `sum (diag (A))`.

**See also:** [eig], page 649.

`r = rref (A)`
`r = rref (A, tol)`
`[r, k] = rref (...)`

Return the reduced row echelon form of $A$.

*tol* defaults to `eps * max (size (A)) * norm (A, inf)`.

The optional return argument $k$ contains the vector of "bound variables", which are those columns on which elimination has been performed.

`n = vecnorm (A)`
`n = vecnorm (A, p)`
`n = vecnorm (A, p, dim)`

Return the vector p-norm of the elements of array $A$ along dimension *dim*.

The p-norm of a vector is defined as

$$\|A\|_p = \left[\sum_{i=1}^{N} |A_i|^p\right]^{1/p}$$

The input $p$ must be a positive scalar. If omitted it defaults to 2 (Euclidean norm or distance). Other special values of $p$ are 1 (Manhattan norm, sum of absolute values) and `Inf` (absolute value of largest element).

The input *dim* specifies the dimension of the array on which the function operates and must be a positive integer. If omitted the first non-singleton dimension is used.

**See also:** [norm], page 654.

## 18.3 Matrix Factorizations

`R = chol (A)`
`[R, p] = chol (A)`
`[R, p, Q] = chol (A)`
`[R, p, Q] = chol (A, "vector")`
`[L, ...] = chol (..., "lower")`
`[R, ...] = chol (..., "upper")`

Compute the upper Cholesky factor, $R$, of the real symmetric or complex Hermitian positive definite matrix $A$.

The upper Cholesky factor $R$ is computed by using the upper triangular part of matrix $A$ and is defined by $R^T R = A$.

Calling `chol` using the optional `"upper"` flag has the same behavior. In contrast, using the optional `"lower"` flag, `chol` returns the lower triangular factorization, computed by using the lower triangular part of matrix $A$, such that $LL^T = A$.

Called with one output argument `chol` fails if matrix $A$ is not positive definite. Note that if matrix $A$ is not real symmetric or complex Hermitian then the lower triangular part is considered to be the (complex conjugate) transpose of the upper triangular part, or vice versa, given the `"lower"` flag.

Called with two or more output arguments $p$ flags whether the matrix $A$ was positive definite and `chol` does not fail. A zero value of $p$ indicates that matrix $A$ is positive definite and $R$ gives the factorization. Otherwise, $p$ will have a positive value.

If called with three output arguments matrix $A$ must be sparse and a sparsity preserving row/column permutation is applied to matrix $A$ prior to the factorization. That is $R$ is the factorization of `A(Q,Q)` such that $R^T R = Q^T A Q$.

The sparsity preserving permutation is generally returned as a matrix. However, given the optional flag `"vector"`, $Q$ will be returned as a vector such that $R^T R = A(Q, Q)$.

In general the lower triangular factorization is significantly faster for sparse matrices.

**See also:** [hess], page 659, [lu], page 659, [qr], page 661, [qz], page 664, [schur], page 666, [svd], page 669, [ichol], page 762, [cholinv], page 657, [chol2inv], page 657, [cholupdate], page 658, [cholinsert], page 658, [choldelete], page 658, [cholshift], page 659.

`Ainv = cholinv (A)`

Compute the inverse of the symmetric positive definite matrix $A$ using the Cholesky factorization.

**See also:** [chol], page 657, [chol2inv], page 657, [inv], page 651.

`Ainv = chol2inv (R)`

Invert a symmetric, positive definite square matrix from its Cholesky decomposition, $R$.

Note that $R$ should be an upper-triangular matrix with positive diagonal elements. `chol2inv (U)` provides `inv (R'*R)` but is much faster than using `inv`.

**See also:** [chol], page 657, [cholinv], page 657, [inv], page 651.

`[R1, info] = cholupdate (R, u, op)`

Update or downdate a Cholesky factorization.

Given an upper triangular matrix $R$ and a column vector $u$, attempt to determine another upper triangular matrix $R1$ such that

- $R1$'*$R1$ = $R$'*$R$ + $u$*$u$' if $op$ is `"+"`
- $R1$'*$R1$ = $R$'*$R$ - $u$*$u$' if $op$ is `"-"`

If $op$ is `"-"`, $info$ is set to

- 0 if the downdate was successful,
- 1 if $R$'*$R$ - $u$*$u$' is not positive definite,
- 2 if $R$ is singular.

If $info$ is not present, an error message is printed in cases 1 and 2.

**See also:** [chol], page 657, [cholinsert], page 658, [choldelete], page 658, [cholshift], page 659.

`R1 = cholinsert (R, j, u)`
`[R1, info] = cholinsert (R, j, u)`

Update a Cholesky factorization given a row or column to insert in the original factored matrix.

Given a Cholesky factorization of a real symmetric or complex Hermitian positive definite matrix $A$ = $R$'*$R$, $R$ upper triangular, return the Cholesky factorization of $A1$, where A1(p,p) = A, A1(:,j) = A1(j,:)' = u and p = [1:j-1,j+1:n+1]. u(j) should be positive.

On return, $info$ is set to

- 0 if the insertion was successful,
- 1 if $A1$ is not positive definite,
- 2 if $R$ is singular.

If $info$ is not present, an error message is printed in cases 1 and 2.

**See also:** [chol], page 657, [cholupdate], page 658, [choldelete], page 658, [cholshift], page 659.

`R1 = choldelete (R, j)`

Update a Cholesky factorization given a row or column to delete from the original factored matrix.

Given a Cholesky factorization of a real symmetric or complex Hermitian positive definite matrix $A$ = $R$'*$R$, $R$ upper triangular, return the Cholesky factorization of A(p,p), where p = [1:j-1,j+1:n+1].

**See also:** [chol], page 657, [cholupdate], page 658, [cholinsert], page 658, [cholshift], page 659.

`R1 = cholshift (R, i, j)`
Update a Cholesky factorization given a range of columns to shift in the original factored matrix.

Given a Cholesky factorization of a real symmetric or complex Hermitian positive definite matrix $A = R'*R$, $R$ upper triangular, return the Cholesky factorization of $A$(p,p), where p is the permutation
`p = [1:i-1, circshift(i:j, 1), j+1:n]` if $i$ < $j$
or
`p = [1:j-1, circshift(j:i,-1), i+1:n]` if $j$ < $i$.

**See also:** [chol], page 657, [cholupdate], page 658, [cholinsert], page 658, [choldelete], page 658.

`H = hess (A)`
`[P, H] = hess (A)`
Compute the Hessenberg decomposition of the matrix $A$.

The Hessenberg decomposition is

$$A = PHP^T$$

where $P$ is a square unitary matrix ($P^TP = I$), and $H$ is upper Hessenberg ($H_{i,j} = 0, \forall i > j + 1$).

The Hessenberg decomposition is usually used as the first step in an eigenvalue computation, but has other applications as well (see Golub, Nash, and Van Loan, IEEE Transactions on Automatic Control, 1979).

**See also:** [eig], page 649, [chol], page 657, [lu], page 659, [qr], page 661, [qz], page 664, [schur], page 666, [svd], page 669.

`[L, U] = lu (A)`
`[L, U, P] = lu (A)`
`[L, U, P, Q] = lu (S)`
`[L, U, P, Q, R] = lu (S)`
`[...] = lu (S, thresh)`
`y = lu (...)`
`[...] = lu (..., "vector")`
Compute the LU decomposition of $A$.

If $A$ is full then subroutines from LAPACK are used, and if $A$ is sparse then UMFPACK is used.

The result is returned in a permuted form, according to the optional return value $P$. For example, given the matrix `A = [1, 2; 3, 4]`,

```
[L, U, P] = lu (A)
```

returns

```
L =

  1.00000  0.00000
  0.33333  1.00000

U =

  3.00000  4.00000
  0.00000  0.66667

P =

  0  1
  1  0
```

The matrix is not required to be square.

When called with two or three output arguments and a sparse input matrix, `lu` does not attempt to perform sparsity preserving column permutations. Called with a fourth output argument, the sparsity preserving column transformation $Q$ is returned, such that `P * A * Q = L * U`. This is the **preferred** way to call `lu` with sparse input matrices.

Called with a fifth output argument and a sparse input matrix, `lu` attempts to use a scaling factor $R$ on the input matrix such that `P * (R \ A) * Q = L * U`. This typically leads to a sparser and more stable factorization.

An additional input argument *thresh* that defines the pivoting threshold can be given. *thresh* can be a scalar, in which case it defines the UMFPACK pivoting tolerance for both symmetric and unsymmetric cases. If *thresh* is a 2-element vector, then the first element defines the pivoting tolerance for the unsymmetric UMFPACK pivoting strategy and the second for the symmetric strategy. By default, the values defined by `spparms` are used ([0.1, 0.001]).

Given the string argument `"vector"`, `lu` returns the values of $P$ and $Q$ as vector values, such that for full matrix, `A(P,:) = L * U`, and `R(P,:) * A(:,Q) = L * U`.

With two output arguments, returns the permuted forms of the upper and lower triangular matrices, such that `A = L * U`. With one output argument $y$, then the matrix returned by the LAPACK routines is returned. If the input matrix is sparse then the matrix $L$ is embedded into $U$ to give a return value similar to the full case. For both full and sparse matrices, `lu` loses the permutation information.

**See also:** [luupdate], page 660, [ilu], page 764, [chol], page 657, [hess], page 659, [qr], page 661, [qz], page 664, [schur], page 666, [svd], page 669.

`[L, U] = luupdate (L, U, x, y)`
`[L, U, P] = luupdate (L, U, P, x, y)`

Given an LU factorization of a real or complex matrix $A = L*U$, $L$ lower unit trapezoidal and $U$ upper trapezoidal, return the LU factorization of $A + x*y.'$, where $x$ and $y$ are column vectors (rank-1 update) or matrices with equal number of columns (rank-k update).

Optionally, row-pivoted updating can be used by supplying a row permutation (pivoting) matrix $P$; in that case, an updated permutation matrix is returned. Note that if $L$, $U$, $P$ is a pivoted LU factorization as obtained by `lu`:

```
[L, U, P] = lu (A);
```

then a factorization of `A+x*y.'` can be obtained either as

```
[L1, U1] = lu (L, U, P*x, y)
```

or

```
[L1, U1, P1] = lu (L, U, P, x, y)
```

The first form uses the unpivoted algorithm, which is faster, but less stable. The second form uses a slower pivoted algorithm, which is more stable.

The matrix case is done as a sequence of rank-1 updates; thus, for large enough k, it will be both faster and more accurate to recompute the factorization from scratch.

**See also:** [lu], page 659, [cholupdate], page 658, [qrupdate], page 663.

```
[Q, R] = qr (A)
[Q, R, P] = qr (A)
R = qr (A)
[C, R] = qr (A, B)
[C, R, P] = qr (A, B)
X = qr (A, B) # sparse only
[...] = qr (..., "econ")
[...] = qr (..., "vector")
[...] = qr (..., "matrix")
[...] = qr (..., 0)
```

Compute the QR factorization of $A$.

The QR factorization is $QR = A$ where $Q$ is an orthogonal matrix and $R$ is upper triangular.

For example, given the matrix `A = [1, 2; 3, 4]`,

```
[Q, R] = qr (A)
```

returns

```
Q =

  -0.31623  -0.94868
  -0.94868   0.31623

R =

  -3.16228  -4.42719
   0.00000  -0.63246
```

which multiplied together return the original matrix

```
Q * R
  ⇒
     1.0000   2.0000
     3.0000   4.0000
```

If just **one return value** is requested then it is $R$. (Note: unlike most commands, the single return value is not the first return value when multiple values are requested.)

If a third output $P$ is requested, then `qr` calculates the permuted QR factorization $QR = AP$ where $Q$ is an orthogonal matrix, $R$ is upper triangular, and $P$ is a permutation matrix.

If $A$ is dense, the QR factorization is computed using standard LAPACK subroutines. In this case, the permuted factorization has the additional property that the diagonal entries of $R$ are ordered by decreasing magnitude. In other words, `abs (diag (R))` will be ordered from largest to smallest.

If $A$ is sparse, $P$ is a fill-reducing ordering of the columns of $A$. In that case, the diagonal entries of $R$ are not ordered by decreasing magnitude.

For example, given the matrix `A = [1, 2; 3, 4]`,

```
[Q, R, P] = qr (A)
```

returns

```
Q =

  -0.44721  -0.89443
  -0.89443   0.44721

R =

  -4.47214  -3.13050
   0.00000   0.44721

P =

   0  1
   1  0
```

If the input matrix $A$ is sparse, the sparse QR factorization is computed by using SPQR or CXSPARSE (e.g., if SPQR is not available). Because the matrix $Q$ is, in general, a full matrix, it is recommended to request only one return value $R$. In that case, the computation avoids the construction of $Q$ and returns a sparse $R$ such that `R = chol (A' * A)`.

If $A$ is dense, and an additional input matrix $B$ is supplied, and two return values are requested, then `qr` returns $C$, where `C = Q' * B`. This allows the least squares approximation of `A \ B` to be calculated as

```
[C, R] = qr (A, B)
X = R \ C
```

If $A$ is a sparse MxN matrix and an additional matrix $B$ is supplied, one or two return values are possible. If one return value $X$ is requested and M < N, then $X$ is the minimum 2-norm solution of `A \ B`. If M >= N, $X$ is the least squares approximation of `A \ B`. If two return values are requested, $C$ and $R$ have the same meaning as in the dense case ($C$ is dense and $R$ is sparse). The version with one return parameter should be preferred because it uses less memory and can handle rank-deficient matrices better.

If the final argument is the string `"vector"` then $P$ is a permutation vector (of the columns of $A$) instead of a permutation matrix. In this case, the defining relationship is:

```
Q * R = A(:, P)
```

The default, however, is to return a permutation matrix and this may be explicitly specified by using a final argument of `"matrix"`.

When the optional argument is the string `"econ"`, an economy factorization is returned. If the original matrix $A$ has size MxN and M > N, then the economy factorization will calculate just N rows in $R$ and N columns in $Q$ and omit the zeros in $R$. If M $\leq$ N, there is no difference between the economy and standard factorizations.

If the optional argument is the numeric value `0` then `qr` acts as if the `"econ"` and `"vector"` arguments were both given. **Warning:** This syntax is accepted, but no longer recommended and may be removed in the future. Use `"econ"` instead.

Background: The QR factorization has applications in the solution of least squares problems

$$\min_x \|Ax - b\|_2$$

for overdetermined systems of equations (i.e., $A$ is a tall, thin matrix).

The permuted QR factorization `[Q, R, P] = qr (A)` allows the construction of an orthogonal basis of `span (A)`.

**See also:** [chol], page 657, [hess], page 659, [lu], page 659, [qz], page 664, [schur], page 666, [svd], page 669, [qrupdate], page 663, [qrinsert], page 663, [qrdelete], page 664, [qrshift], page 664.

`[Q1, R1] = qrupdate (Q, R, u, v)`

Update a QR factorization given update vectors or matrices.

Given a QR factorization of a real or complex matrix $A = Q^*R$, $Q$ unitary and $R$ upper trapezoidal, return the QR factorization of $A$ + $u^*v$', where $u$ and $v$ are column vectors (rank-1 update) or matrices with equal number of columns (rank-k update). Notice that the latter case is done as a sequence of rank-1 updates; thus, for k large enough, it will be both faster and more accurate to recompute the factorization from scratch.

The QR factorization supplied may be either full (Q is square) or economized (R is square).

**See also:** [qr], page 661, [qrinsert], page 663, [qrdelete], page 664, [qrshift], page 664.

`[Q1, R1] = qrinsert (Q, R, j, x, orient)`

Update a QR factorization given a row or column to insert in the original factored matrix.

Given a QR factorization of a real or complex matrix $A = Q^*R$, $Q$ unitary and $R$ upper trapezoidal, return the QR factorization of [A(:,1:j-1) x A(:,j:n)], where $u$ is a column vector to be inserted into $A$ (if *orient* is `"col"`), or the QR factorization of [A(1:j-1,:);x;A(:,j:n)], where $x$ is a row vector to be inserted into $A$ (if *orient* is `"row"`).

The default value of *orient* is `"col"`. If *orient* is `"col"`, *u* may be a matrix and *j* an index vector resulting in the QR factorization of a matrix *B* such that B(:,*j*) gives *u* and B(:,*j*) = [] gives *A*. Notice that the latter case is done as a sequence of k insertions; thus, for k large enough, it will be both faster and more accurate to recompute the factorization from scratch.

If *orient* is `"col"`, the QR factorization supplied may be either full (Q is square) or economized (R is square).

If *orient* is `"row"`, full factorization is needed.

**See also:** [qr], page 661, [qrupdate], page 663, [qrdelete], page 664, [qrshift], page 664.

`[`*Q1*`, `*R1*`] = qrdelete (`*Q*`, `*R*`, `*j*`, `*orient*`)`

Update a QR factorization given a row or column to delete from the original factored matrix.

Given a QR factorization of a real or complex matrix $A = Q^*R$, $Q$ unitary and $R$ upper trapezoidal, return the QR factorization of [A(:,1:j-1), U, A(:,j:n)], where *u* is a column vector to be inserted into *A* (if *orient* is `"col"`), or the QR factorization of [A(1:j-1,:);X;A(:,j:n)], where *x* is a row *orient* is `"row"`). The default value of *orient* is `"col"`.

If *orient* is `"col"`, *j* may be an index vector resulting in the QR factorization of a matrix *B* such that A(:,*j*) = [] gives *B*. Notice that the latter case is done as a sequence of k deletions; thus, for k large enough, it will be both faster and more accurate to recompute the factorization from scratch.

If *orient* is `"col"`, the QR factorization supplied may be either full (Q is square) or economized (R is square).

If *orient* is `"row"`, full factorization is needed.

**See also:** [qr], page 661, [qrupdate], page 663, [qrinsert], page 663, [qrshift], page 664.

`[`*Q1*`, `*R1*`] = qrshift (`*Q*`, `*R*`, `*i*`, `*j*`)`

Update a QR factorization given a range of columns to shift in the original factored matrix.

Given a QR factorization of a real or complex matrix $A = Q^*R$, $Q$ unitary and $R$ upper trapezoidal, return the QR factorization of $A$(:,p), where p is the permutation
`p = [1:i-1, circshift(i:j, 1), j+1:n]` if $i < j$
or
`p = [1:j-1, circshift(j:i,-1), i+1:n]` if $j < i$.

**See also:** [qr], page 661, [qrupdate], page 663, [qrinsert], page 663, [qrdelete], page 664.

`[`*AA*`, `*BB*`, `*Q*`, `*Z*`, `*V*`, `*W*`] = qz (`*A*`, `*B*`)`
`[`*AA*`, `*BB*`, `*Q*`, `*Z*`, `*V*`, `*W*`] = qz (`*A*`, `*B*`, `*opt*`)`

Compute the QZ decomposition of a generalized eigenvalue problem.

The generalized eigenvalue problem is defined as

$$Ax = \lambda Bx$$

There are two calling forms of the function:

1. `[AA, BB, Q, Z, V, W] = qz (A, B)`

   Compute the complex QZ decomposition, generalized eigenvectors, and generalized eigenvalues.

$$AA = Q \cdot A \cdot Z, BB = Q \cdot B \cdot Z$$

$$A \cdot V \cdot \mathrm{diag}(BB) = B \cdot V \cdot \mathrm{diag}(AA)$$

$$\mathrm{diag}(BB) \cdot W^T \cdot A = \mathrm{diag}(AA) \cdot W^T \cdot B$$

   with $AA$ and $BB$ upper triangular, and $Q$ and $Z$ unitary. The matrices $V$ and $W$ respectively contain the right and left generalized eigenvectors.

2. `[AA, BB, Q, Z, V, W] = qz (A, B, opt)`

   The *opt* argument must be equal to either `"real"` or `"complex"`. If it is equal to `"complex"`, then this calling form is equivalent to the first one with only two input arguments.

   If *opt* is equal to `"real"`, then the real QZ decomposition is computed. In particular, $AA$ is only guaranteed to be quasi-upper triangular with 1-by-1 and 2-by-2 blocks on the diagonal, and $Q$ and $Z$ are orthogonal. The identities mentioned above for right and left generalized eigenvectors are only verified if $AA$ is upper triangular (i.e., when all the generalized eigenvalues are real, in which case the real and complex QZ coincide).

Note: `qz` performs permutation balancing, but not scaling (see [`balance`], page 647), which may lead to less accurate results than `eig`. The order of output arguments was selected for compatibility with MATLAB.

For eigenvalues, use the `ordeig` function to compute them based on the resulting $AA$ and $BB$ matrices.

**See also:** [eig], page 649, [gsvd], page 651, [balance], page 647, [chol], page 657, [hess], page 659, [lu], page 659, [qr], page 661, [qzhess], page 665, [schur], page 666, [ordeig], page 669.

`[aa, bb, q, z] = qzhess (A, B)`

Compute the Hessenberg-triangular decomposition of the matrix pencil `(A, B)`, returning `aa = q * A * z`, `bb = q * B * z`, with $q$ and $z$ orthogonal.

For example:

```
[aa, bb, q, z] = qzhess ([1, 2; 3, 4], [5, 6; 7, 8])
  ⇒ aa =
      -3.02244  -4.41741
       0.92998   0.69749
  ⇒ bb =
      -8.60233  -9.99730
       0.00000  -0.23250
  ⇒ q =
      -0.58124  -0.81373
      -0.81373   0.58124
  ⇒ z =
     Diagonal Matrix
       1   0
       0   1
```

The Hessenberg-triangular decomposition is the first step in Moler and Stewart's QZ decomposition algorithm.

Algorithm taken from Golub and Van Loan, *Matrix Computations, 2nd edition.*

**See also:** [lu], page 659, [chol], page 657, [hess], page 659, [qr], page 661, [qz], page 664, [schur], page 666, [svd], page 669.

`S = schur (A)`
`S = schur (A, "real")`
`S = schur (A, "complex")`
`S = schur (A, opt)`
`[U, S] = schur (...)`

Compute the Schur decomposition of $A$.

The Schur decomposition of a square matrix $A$ is defined as

$$S = U^T A U$$

where $U$ is a unitary matrix ($U^T U$ is identity) and $S$ is upper triangular. The eigenvalues of $A$ (and $S$) are the diagonal elements of $S$. If the matrix $A$ is real, then the real Schur decomposition is computed, in which the matrix $U$ is orthogonal and $S$ is block upper triangular with blocks of size at most $2 \times 2$ along the diagonal.

The default for real matrices is a real Schur decomposition. A complex decomposition may be forced by passing the flag `"complex"`.

The eigenvalues are optionally ordered along the diagonal according to the value of *opt*:

`opt = "a"` Move eigenvalues with negative real parts to the leading block of $S$. Mnemonic: `"a"` for Algebraic Riccati Equations, where this ordering is useful.

`opt = "d"` Move eigenvalues with magnitude less than one to the leading block of $S$. Mnemonic: `"d"` for Discrete Algebraic Riccati Equations, where this ordering is useful.

`opt = "u"` Unordered. No particular ordering of eigenvalues (default).

The leading $k$ columns of $U$ always span the $A$-invariant subspace corresponding to the $k$ leading eigenvalues of $S$.

**See also:** [rsf2csf], page 667, [ordschur], page 667, [trexc], page 667, [ordeig], page 669, [lu], page 659, [chol], page 657, [hess], page 659, [qr], page 661, [qz], page 664, [svd], page 669, [eig], page 649.

`[U, T] = rsf2csf (UR, TR)`

Convert a real, upper quasi-triangular Schur form $TR$ to a complex, upper triangular Schur form $T$.

Note that the following relations hold: $UR \cdot TR \cdot UR^T = UTU^\dagger$ and $U^\dagger U$ is the identity matrix I.

Note also that $U$ and $T$ are not unique.

**See also:** [schur], page 666, [ordschur], page 667, [trexc], page 667.

`[UR, SR] = ordschur (U, S, select)`

Reorder the real Schur factorization ($U$,$S$) obtained with the `schur` function, so that selected eigenvalues appear in the upper left diagonal blocks of the quasi triangular Schur matrix.

The logical vector *select* specifies the selected eigenvalues as they appear along $S$'s diagonal.

For example, given the matrix `A = [1, 2; 3, 4]`, and its Schur decomposition

```
[U, S] = schur (A)
```

which returns

```
U =

  -0.82456  -0.56577
   0.56577  -0.82456

S =

  -0.37228  -1.00000
   0.00000   5.37228
```

It is possible to reorder the decomposition so that the positive eigenvalue is in the upper left corner, by doing:

```
[U, S] = ordschur (U, S, [0,1])
```

**See also:** [schur], page 666, [trexc], page 667, [ordeig], page 669, [ordqz], page 668.

`[U, T] = trexc (U, T, M)`

Reorder the Schur factorization using LAPACK function `ZTREXC`.

Given a unitary matrix $U$ and an upper triangular Schur form $T$ from the Schur decomposition `A = U*T*U'`, this function reorders the diagonal elements of $T$ according to the swap specifications in $M$.

Each row of $M$ contains two indices `[IFST, ILST]` specifying that the diagonal element at position `IFST` should be moved to position `ILST`. The swaps are performed sequentially.

**See also:** [schur], page 666, [ordschur], page 667.

`[AR, BR, QR, ZR] = ordqz (AA, BB, Q, Z, keyword)`
`[AR, BR, QR, ZR] = ordqz (AA, BB, Q, Z, select)`

Reorder the QZ decomposition of a generalized eigenvalue problem.

The generalized eigenvalue problem is defined as

$$Ax = \lambda Bx$$

Its generalized Schur decomposition is computed using the `qz` algorithm:

`[AA, BB, Q, Z] = qz (A, B)`

where $AA$, $BB$, $Q$, and $Z$ fulfill

$$AA = Q \cdot A \cdot Z, BB = Q \cdot B \cdot Z$$

The `ordqz` function computes a unitary transformation $QR$ and $ZR$ such that the order of the eigenvalue on the diagonal of $AA$ and $BB$ is changed. The resulting reordered matrices $AR$ and $BR$ fulfill:

$$A_R = Q_R \cdot A \cdot Z_R, B_R = Q_R \cdot B \cdot Z_R$$

The function can either be called with the *keyword* argument which selects the eigenvalues in the top left block of $AR$ and $BR$ in the following way:

`"S"`, `"udi"`
: small: leading block has all $|\lambda| < 1$

`"B"`, `"udo"`
: big: leading block has all $|\lambda| \geq 1$

`"-"`, `"lhp"`
: negative real part: leading block has all eigenvalues in the open left half-plane

`"+"`, `"rhp"`
: non-negative real part: leading block has all eigenvalues in the closed right half-plane

If a logical vector *select* is given instead of a keyword the `ordqz` function reorders all eigenvalues `k` to the left block for which `select(k)` is true.

Note: The keywords are compatible with the ones from `qr`.

**See also:** [eig], page 649, [ordeig], page 669, [qz], page 664, [schur], page 666, [ordschur], page 667.

*lambda* = ordeig (*A*)
*lambda* = ordeig (*A*, *B*)

Return the eigenvalues of quasi-triangular matrices in their order of appearance in the matrix $A$.

The quasi-triangular matrix $A$ is usually the result of a Schur factorization. If called with a second input $B$ then the generalized eigenvalues of the pair $A$, $B$ are returned in the order of appearance of the matrix `A-lambda*B`. The pair $A$, $B$ is usually the result of a QZ decomposition.

**See also:** [ordschur], page 667, [ordqz], page 668, [eig], page 649, [schur], page 666, [qz], page 664.

*angle* = subspace (*A*, *B*)

Determine the largest principal angle between two subspaces spanned by the columns of matrices $A$ and $B$.

Reference: Andrew V. Knyazev, Merico E. Argentati, "Principal Angles between Subspaces in an A-Based Scalar Product: Algorithms and Perturbation Estimates", *SIAM Journal on Scientific Computing*, Vol. 23, No. 6, pp. 2008–2040.

*s* = svd (*A*)
[*U*, *S*, *V*] = svd (*A*)
[*U*, *S*, *V*] = svd (*A*, "*econ*")
[*U*, *S*, *V*] = svd (*A*, *0*)

Compute the singular value decomposition of $A$.

The singular value decomposition is defined by the relation

$$A = USV^{\dagger}$$

The function `svd` normally returns only the vector of singular values. When called with three return values, it computes $U$, $S$, and $V$. For example,

```
svd (hilb (3))
```

returns

```
ans =

  1.4083189
  0.1223271
  0.0026873
```

and

```
[u, s, v] = svd (hilb (3))
```

returns

```
u =

  -0.82704   0.54745   0.12766
  -0.45986  -0.52829  -0.71375
  -0.32330  -0.64901   0.68867

s =

  1.40832  0.00000  0.00000
  0.00000  0.12233  0.00000
  0.00000  0.00000  0.00269

v =

  -0.82704   0.54745   0.12766
  -0.45986  -0.52829  -0.71375
  -0.32330  -0.64901   0.68867
```

When given a second argument that is not 0, `svd` returns an economy-sized decomposition, eliminating the unnecessary rows or columns of $U$ or $V$.

If the second argument is exactly 0, then the choice of decomposition is based on the matrix $A$. If $A$ has more rows than columns then an economy-sized decomposition is returned, otherwise a regular decomposition is calculated.

Algorithm Notes: When calculating the full decomposition (left and right singular matrices in addition to singular values) there is a choice of two routines in LAPACK. The default routine used by Octave is `gesvd`. The alternative is `gesdd` which is 5X faster, but may use more memory and may be inaccurate for some input matrices. There is a third routine `gejsv`, suitable for better accuracy at extreme scale. See the documentation for `svd_driver` for more information on choosing a driver.

**See also:** [svd_driver], page 670, [svds], page 757, [eig], page 649, [lu], page 659, [chol], page 657, [hess], page 659, [qr], page 661, [qz], page 664.

*val* = svd_driver ()
*old_val* = svd_driver (*new_val*)
*old_val* = svd_driver (*new_val*, "*local*")

Query or set the underlying LAPACK driver used by `svd`.

Currently recognized values are `"gesdd"`, `"gesvd"`, and `"gejsv"`. The default is `"gesvd"`.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

Algorithm Notes: The LAPACK library routines `gesvd` and `gesdd` are different only when calculating the full singular value decomposition (left and right singular matrices as well as singular values). When calculating just the singular values the following discussion is not relevant.

The newer `gesdd` routine is based on a Divide-and-Conquer algorithm that is 5X faster than the alternative `gesvd`, which is based on QR factorization. However, the new

algorithm can use significantly more memory. For an MxN input matrix the memory usage is of order O(min(M,N) ^ 2), whereas the alternative is of order O(max(M,N)).

The routine `gejsv` uses a preconditioned Jacobi SVD algorithm. Unlike `gesvd` and `gesdd`, in `gejsv`, there is no bidiagonalization step that could contaminate accuracy in some extreme cases. Also, `gejsv` is known to be optimally accurate in some sense. However, the speed is slower (single threaded at its core) and uses more memory (O(min(M,N) ^ 2 + M + N)).

Beyond speed and memory issues, there have been instances where some input matrices were not accurately decomposed by `gesdd`. See currently active bug https://savannah.gnu.org/bugs/?55564. Until these accuracy issues are resolved in a new version of the LAPACK library, the default driver in Octave has been set to `"gesvd"`.

**See also:** [svd], page 669.

`[`*housv*`, `*beta*`, `*zer*`] = housh (`*x*`, `*j*`, `*z*`)`

Compute Householder reflection vector *housv* to reflect *x* to be the j-th column of identity, i.e.,

```
(I - beta*housv*housv')x =  norm (x)*e(j) if x(j) < 0,
(I - beta*housv*housv')x = -norm (x)*e(j) if x(j) >= 0
```

Inputs

| | |
|---|---|
| *x* | vector |
| *j* | index into vector |
| *z* | threshold for zero (usually should be the number 0) |

Outputs (see Golub and Van Loan):

| | |
|---|---|
| *beta* | If beta = 0, then no reflection need be applied (zer set to 0) |
| *housv* | householder vector |

`[`*u*`, `*h*`, `*nu*`] = krylov (`*A*`, `*V*`, `*k*`, `*eps1*`, `*pflg*`)`

Construct an orthogonal basis *u* of a block Krylov subspace.

The block Krylov subspace has the following form:

```
[v a*v a^2*v ... a^(k+1)*v]
```

The construction is made with Householder reflections to guard against loss of orthogonality.

If *V* is a vector, then *h* contains the Hessenberg matrix such that `a*u == u*h+rk*ek'`, in which `rk = a*u(:,k)-u*h(:,k)`, and `ek'` is the vector `[0, 0, ..., 1]` of length *k*. Otherwise, *h* is meaningless.

If *V* is a vector and *k* is greater than `length (A) - 1`, then *h* contains the Hessenberg matrix such that `a*u == u*h`.

The value of *nu* is the dimension of the span of the Krylov subspace (based on *eps1*).

If *b* is a vector and *k* is greater than *m-1*, then *h* contains the Hessenberg decomposition of *A*.

The optional parameter *eps1* is the threshold for zero. The default value is 1e-12.

If the optional parameter *pflg* is nonzero, row pivoting is used to improve numerical behavior. The default value is 0.

Reference: A. Hodel, P. Misra, "Partial Pivoting in the Computation of Krylov Subspaces of Large Sparse Systems", *Proceedings of the 42nd IEEE Conference on Decision and Control*, December 2003.

## 18.4 Functions of a Matrix

`r = expm (A)`

Return the exponential of a matrix.

The matrix exponential is defined as the infinite Taylor series

$$\exp(A) = I + A + \frac{A^2}{2!} + \frac{A^3}{3!} + \cdots$$

However, the Taylor series is *not* the way to compute the matrix exponential; see Moler and Van Loan, "Nineteen Dubious Ways to Compute the Exponential of a Matrix", *SIAM Review*, 1978. This routine uses Ward's diagonal Padé approximation method with three step preconditioning (*SIAM Journal on Numerical Analysis*, 1977). Diagonal Padé approximations are rational polynomials of matrices $D_q(A)^{-1}N_q(A)$ whose Taylor series matches the first $2q+1$ terms of the Taylor series above; direct evaluation of the Taylor series (with the same preconditioning steps) may be desirable in lieu of the Padé approximation when $D_q(A)$ is ill-conditioned.

**See also:** [logm], page 672, [sqrtm], page 672.

`s = logm (A)`
`s = logm (A, opt_iters)`
`[s, iters] = logm (...)`

Compute the matrix logarithm of the square matrix $A$.

The implementation utilizes a Padé approximant and the identity

```
logm (A) = 2^k * logm (A^(1 / 2^k))
```

The optional input *opt_iters* is the maximum number of square roots to compute and defaults to 100.

The optional output *iters* is the number of square roots actually computed.

**See also:** [expm], page 672, [sqrtm], page 672.

`s = sqrtm (A)`
`[s, error_estimate] = sqrtm (A)`

Compute the matrix square root of the square matrix $A$.

Ref: N.J. Higham, *A New sqrtm for* MATLAB, Numerical Analysis Report No. 336, Manchester Centre for Computational Mathematics, Manchester, England, January 1999.

**See also:** [expm], page 672, [logm], page 672.

`F = funm (A, fun)`
`F = funm (A, fun, options)`
`F = funm (A, fun, options, p1, ...)`
`[F, exitflag] = funm (...)`
`[F, exitflag, output] = funm (...)`

Evaluate a general matrix function.

`funm (A, fun)` evaluates the function *fun* at the square matrix $A$. The input `fun (x, k)` must return the $k$'th derivative of the function represented by *fun* evaluated at the vector $x$. The function *fun* must have a Taylor series representation with an infinite radius of convergence.

The special functions `exp`, `log`, `sin`, `cos`, `sinh`, and `cosh` can be passed by function handle; for example `funm (A, @cos)`.

For matrix square roots, use `sqrtm` instead. For matrix exponentials, either `expm` or `funm (A, @exp)` may be faster or more accurate, depending on $A$.

An optional third input in the form of an options struct *options* can be used to specify verbosity of the function and to influence certain of the algorithms. See the references below for more details on the latter.

*options* can have the following fields:

`Display` Specify what information will be printed to the screen during the course of calculations. It can be either a string value of `"off"` (no info, the default), `"on"` (some info), `"verbose"` (maximum info); or a scalar value between `0` (no info, the default) and `5` (maximum info). When `Display` is `"verbose"` or a scalar value $\geq$ `3`, a plot of the eigenvalues and groups will also be shown.

`TolBlk` Tolerance used in determining the blocking (positive scalar, default: `0.1`).

`TolTay` Tolerance used in the convergence test for evaluating the Taylor series (positive scalar, default: `eps`).

`MaxTerms` The maximum number of Taylor series terms (positive integer, default: `250`).

`MaxSqrt` The maximum number of square roots evaluated in the course of inverse scaling and squaring (positive integer, default: `100`). This option is only used when computing a logarithm where it functions similarly to `MaxTerms`.

`Ord` Define a custom blocking pattern in the form of a vector whose length equals the order of the matrix $A$.

Octave accepts any case for these fieldnames.

All inputs beyond *options* will be passed as positional arguments to the function *fun*.

Optional outputs:

*exitflag* Scalar exit flag that describes the exit condition:

`0` — The algorithm was successful.

`1` — One or more Taylor series evaluations did not converge, but the computed value of $F$ might still be accurate.

*output* Structure with the following fields:

`terms` Vector for which `output.terms(i)` is the number of Taylor series terms used when evaluating the i'th block, or, in the case of the logarithm, the number of square roots of matrices of dimension greater than 2.

`ind` Cell array for which the (i,j) block of the reordered Schur factor $T$ is `T(output.ind{i}, output.ind{j})`.

`ord` Ordering of the Schur form, as passed to `ordschur`.

`T` Reordered Schur form.

If the Schur form is diagonal then `output = struct ("terms", ones (n, 1), "ind", {1:n}, "ord", [], "T", T)`.

Example:

```
F = funm (magic (3), @sin);
⇒ F =
   -0.3850    1.0191    0.0162
    0.6179    0.2168   -0.1844
    0.4173   -0.5856    0.8185
```

The code

```
S = funm (X, @sin);
C = funm (X, @cos);
```

will produce the same results (within possible rounding error) as

```
E = expm (i*X);
C = real (E);
S = imag (E);
```

References:

Philip I. Davies and Nicholas J. Higham, "A Schur-Parlett algorithm for computing matrix functions", *SIAM Journal on Matrix Analysis and Applications*, Vol. 25(2), pp. 464–485, 2003.

Nicholas J. Higham, *Functions of Matrices: Theory and Computation*, SIAM, pp. 425, 2008, ISBN 978-0-898716-46-7.

**See also:** [expm], page 672, [logm], page 672, [sqrtm], page 672.

`C = kron (A, B)`
`C = kron (A1, A2, ...)`

Form the Kronecker product of two or more matrices.

This is defined block by block as

```
c = [ a(i,j)*b ]
```

For example:

```
kron (1:4, ones (3, 1))
     ⇒  1  2  3  4
        1  2  3  4
        1  2  3  4
```

If there are more than two input arguments *A1*, *A2*, ..., *An* the Kronecker product is computed as

```
kron (kron (A1, A2), ..., An)
```

Since the Kronecker product is associative, this is well-defined.

**See also:** [tensorprod], page 675.

`C = tensorprod (A, B, dimA, dimB)`
`C = tensorprod (A, B, dim)`
`C = tensorprod (A, B)`
`C = tensorprod (A, B, "all")`
`C = tensorprod (A, B, ..., "NumDimensionsA", value)`

Compute the tensor product between numeric tensors *A* and *B*.

The dimensions of *A* and *B* that are contracted are defined by *dimA* and *dimB*, respectively. *dimA* and *dimB* are scalars or equal length vectors that define the dimensions to match up. The matched dimensions of *A* and *B* must have the same number of elements.

When only *dim* is used, it is equivalent to `dimA = dimB = dim`.

When no dimensions are specified, `dimA = dimB = []`. This computes the outer product between *A* and *B*.

Using the `"all"` option results in the inner product between *A* and *B*. This requires `size (A) == size (B)`.

Use the property-value pair with the property name `"NumDimensionsA"` when *A* has trailing singleton dimensions that should be transferred to *C*. The specified *value* should be the total number of dimensions of *A*.

MATLAB Compatibility: Octave does not currently support the `"property_name=value"` syntax for the `"NumDimensionsA"` parameter.

**See also:** [kron], page 674, [dot], page 623, [mtimes], page 173.

`C = blkmm (A, B)`

Compute products of matrix blocks.

The blocks are given as 2-dimensional subarrays of the arrays *A*, *B*. The size of *A* must have the form `[m,k,...]` and size of *B* must be `[k,n,...]`. The result is then of size `[m,n,...]` and is computed as follows:

```
for i = 1:prod (size (A)(3:end))
  C(:,:,i) = A(:,:,i) * B(:,:,i)
endfor
```

`X = sylvester (A, B, C)`

Solve the Sylvester equation.

The Sylvester equation is defined as:

$$AX + XB = C$$

The solution is computed using standard LAPACK subroutines.

For example:

```
sylvester ([1, 2; 3, 4], [5, 6; 7, 8], [9, 10; 11, 12])
   ⇒ [ 0.50000, 0.66667; 0.66667, 0.50000 ]
```

## 18.5 Specialized Solvers

`x = bicg (A, b)`
`x = bicg (A, b, tol)`
`x = bicg (A, b, tol, maxit)`
`x = bicg (A, b, tol, maxit, M)`
`x = bicg (A, b, tol, maxit, M1, M2)`
`x = bicg (A, b, tol, maxit, M, [], x0)`
`x = bicg (A, b, tol, maxit, M1, M2, x0)`
`x = bicg (A, b, tol, maxit, M, [], x0, ...)`
`x = bicg (A, b, tol, maxit, M1, M2, x0, ...)`
`[x, flag, relres, iter, resvec] = bicg (A, b, ...)`

Solve the linear system of equations `A * x = b` by means of the Bi-Conjugate Gradient iterative method.

The input arguments are:

- *A* is the matrix of the linear system and it must be square. *A* can be passed as a matrix, function handle, or inline function `Afcn` such that `Afcn (x, "notransp") = A * x` and `Afcn (x, "transp") = A' * x`. Additional parameters to `Afcn` may be passed after *x0*.
- *b* is the right-hand side vector. It must be a column vector with the same number of rows as *A*.
- *tol* is the required relative tolerance for the residual error, `b - A * x`. The iteration stops if `norm (b - A * x)` $\leq$ `tol * norm (b)`. If *tol* is omitted or empty, then a tolerance of 1e-6 is used.
- *maxit* is the maximum allowed number of iterations; if *maxit* is omitted or empty then a value of 20 is used.
- *M1*, *M2* are the preconditioners. The preconditioner *M* is given as `M = M1 * M2`. Both *M1* and *M2* can be passed as a matrix or as a function handle or inline function `g` such that `g (x, "notransp") = M1 \ x` or `g (x, "notransp") = M2 \ x` and `g (x, "transp") = M1' \ x` or `g (x, "transp") = M2' \ x`. If *M1* is omitted or empty, then preconditioning is not applied. The preconditioned system is theoretically equivalent to applying the `bicg` method to the linear system `inv (M1) * A * inv (M2) * y = inv (M1) * b` and `inv (M2') * A' * inv (M1') * z = inv (M2') * b` and then setting `x = inv (M2) * y`.
- *x0* is the initial guess. If *x0* is omitted or empty then the function sets *x0* to a zero vector by default.

Any arguments which follow *x0* are treated as parameters, and passed in an appropriate manner to any of the functions (*Afcn* or *Mfcn*) or that have been given to `bicg`.

The output parameters are:

- *x* is the computed approximation to the solution of `A * x = b`. If the algorithm did not converge, then *x* is the iteration which has the minimum residual.
- *flag* indicates the exit status:
  - 0: The algorithm converged to within the prescribed tolerance.

  - 1: The algorithm did not converge and it reached the maximum number of iterations.
  - 2: The preconditioner matrix is singular.
  - 3: The algorithm stagnated, i.e., the absolute value of the difference between the current iteration *x* and the previous is less than `eps * norm (x,2)`.
  - 4: The algorithm could not continue because intermediate values became too small or too large for reliable computation.
- *relres* is the ratio of the final residual to its initial value, measured in the Euclidean norm.
- *iter* is the iteration which *x* is computed.
- *resvec* is a vector containing the residual at each iteration. The total number of iterations performed is given by `length (resvec) - 1`.

Consider a trivial problem with a tridiagonal matrix

```
n = 20;
A = toeplitz (sparse ([1, 1], [1, 2], [2, 1] * n ^ 2, 1, n)) + ...
    toeplitz (sparse (1, 2, -1, 1, n) * n / 2, ...
              sparse (1, 2, 1, 1, n) * n / 2);
b = A * ones (n, 1);
restart = 5;
[M1, M2] = ilu (A);  # in this tridiag case, it corresponds to lu (A)
M = M1 * M2;
Afcn = @(x, string) strcmp (string, "notransp") * (A * x) + ...
                    strcmp (string, "transp") * (A' * x);
Mfcn = @(x, string) strcmp (string, "notransp") * (M \ x) + ...
                    strcmp (string, "transp") * (M' \ x);
M1fcn = @(x, string) strcmp (string, "notransp") * (M1 \ x) + ...
                     strcmp (string, "transp") * (M1' \ x);
M2fcn = @(x, string) strcmp (string, "notransp") * (M2 \ x) + ...
                     strcmp (string, "transp") * (M2' \ x);
```

Example 1: simplest usage of `bicg`

```
x = bicg (A, b)
```

Example 2: `bicg` with a function that computes `A*x` and `A'*x`

```
x = bicg (Afcn, b, [], n)
```

Example 3: `bicg` with a preconditioner matrix *M*

```
x = bicg (A, b, 1e-6, n, M)
```

Example 4: `bicg` with a function as preconditioner

```
x = bicg (Afcn, b, 1e-6, n, Mfcn)
```

Example 5: `bicg` with preconditioner matrices *M1* and *M2*

```
x = bicg (A, b, 1e-6, n, M1, M2)
```

Example 6: `bicg` with functions as preconditioners

```
x = bicg (Afcn, b, 1e-6, n, M1fcn, M2fcn)
```

Example 7: `bicg` with as input a function requiring an argument

```
function y = Ap (A, x, string, z)
  ## compute A^z * x or (A^z)' * x
  y = x;
  if (strcmp (string, "notransp"))
    for i = 1:z
      y = A * y;
    endfor
  elseif (strcmp (string, "transp"))
    for i = 1:z
      y = A' * y;
    endfor
  endif
endfunction

Apfcn = @(x, string, p) Ap (A, x, string, p);
x = bicg (Apfcn, b, [], [], [], [], [], 2);
```

Reference:

Y. Saad, *Iterative Methods for Sparse Linear Systems*, 2nd edition, SIAM, 2003.

**See also:** [bicgstab], page 678, [cgs], page 680, [gmres], page 682, [pcg], page 758, [qmr], page 685, [tfqmr], page 686.

`x = bicgstab (A, b, tol, maxit, M1, M2, x0, ...)`
`x = bicgstab (A, b, tol, maxit, M, [], x0, ...)`
`[x, flag, relres, iter, resvec] = bicgstab (A, b, ...)`

Solve `A x = b` using the stabilized Bi-conjugate gradient iterative method.

The input parameters are:

- *A* is the matrix of the linear system and it must be square. *A* can be passed as a matrix, function handle, or inline function `Afcn` such that `Afcn(x) = A * x`. Additional parameters to `Afcn` are passed after *x0*.
- *b* is the right hand side vector. It must be a column vector with the same number of rows as *A*.
- *tol* is the required relative tolerance for the residual error, `b - A * x`. The iteration stops if `norm (b - A * x)` $\leq$ `tol * norm (b)`. If *tol* is omitted or empty, then a tolerance of 1e-6 is used.
- *maxit* the maximum number of outer iterations, if not given or set to [] the default value `min (20, numel (b))` is used.
- *M1*, *M2* are the preconditioners. The preconditioner *M* is given as `M = M1 * M2`. Both *M1* and *M2* can be passed as a matrix or as a function handle or inline function `g` such that `g(x) = M1 \ x` or `g(x) = M2 \ x`. The technique used is the right preconditioning, i.e., it is solved `A * inv (M) * y = b` and then `x = inv (M) * y`.
- *x0* the initial guess, if not given or set to [] the default value `zeros (size (b))` is used.

The arguments which follow *x0* are treated as parameters, and passed in a proper way to any of the functions (*A* or *M*) which are passed to `bicstab`.

The output parameters are:

- *x* is the approximation computed. If the method doesn't converge then it is the iterated with the minimum residual.
- *flag* indicates the exit status:
  - 0: iteration converged to the within the chosen tolerance
  - 1: the maximum number of iterations was reached before convergence
  - 2: the preconditioner matrix is singular
  - 3: the algorithm reached stagnation
  - 4: the algorithm can't continue due to a division by zero
- *relres* is the relative residual obtained with as `(A*x-b) / norm(b)`.
- *iter* is the (possibly half) iteration which *x* is computed. If it is an half iteration then it is `iter + 0.5`
- *resvec* is a vector containing the residual of each half and total iteration (There are also the half iterations since *x* is computed in two steps at each iteration). Doing `(length(resvec) - 1) / 2` is possible to see the total number of (total) iterations performed.

Let us consider a trivial problem with a tridiagonal matrix

```
n = 20;
A = toeplitz (sparse ([1, 1], [1, 2], [2, 1] * n ^ 2, 1, n))  + ...
    toeplitz (sparse (1, 2, -1, 1, n) * n / 2, ...
    sparse (1, 2, 1, 1, n) * n / 2);
b = A * ones (n, 1);
restart = 5;
[M1, M2] = ilu (A); # in this tridiag case, it corresponds to lu (A)
M = M1 * M2;
Afcn = @(x) A * x;
Mfcn = @(x) M \ x;
M1fcn = @(x) M1 \ x;
M2fcn = @(x) M2 \ x;
```

Example 1: simplest usage of `bicgstab`

```
x = bicgstab (A, b, [], n)
```

Example 2: `bicgstab` with a function which computes `A * x`

```
x = bicgstab (Afcn, b, [], n)
```

Example 3: `bicgstab` with a preconditioner matrix *M*

```
x = bicgstab (A, b, [], 1e-06, n, M)
```

Example 4: `bicgstab` with a function as preconditioner

```
x = bicgstab (Afcn, b, 1e-6, n, Mfcn)
```

Example 5: `bicgstab` with preconditioner matrices *M1* and *M2*

```
x = bicgstab (A, b, [], 1e-6, n, M1, M2)
```

Example 6: `bicgstab` with functions as preconditioners

```
x = bicgstab (Afcn, b, 1e-6, n, M1fcn, M2fcn)
```

EXAMPLE 7: `bicgstab` with as input a function requiring an argument

```
function y = Ap (A, x, z) # compute A^z * x
   y = x;
   for i = 1:z
     y = A * y;
   endfor
 endfunction
Apfcn = @(x, string, p) Ap (A, x, string, p);
x = bicgstab (Apfcn, b, [], [], [], [], [], 2);
```

EXAMPLE 8: explicit example to show that `bicgstab` uses a right preconditioner

```
[M1, M2] = ilu (A + 0.1 * eye (n)); # factorization of A perturbed
M = M1 * M2;

## reference solution computed by bicgstab after one iteration
[x_ref, fl] = bicgstab (A, b, [], 1, M)

## right preconditioning
[y, fl] = bicgstab (A / M, b, [], 1)
x = M \ y # compare x and x_ref
```

Reference:

Y. Saad, *Iterative Methods for Sparse Linear Systems*, 2nd edition, SIAM, 2003.

**See also:** [bicg], page 676, [cgs], page 680, [gmres], page 682, [pcg], page 758, [qmr], page 685, [tfqmr], page 686.

`x = cgs (A, b, tol, maxit, M1, M2, x0, ...)`
`x = cgs (A, b, tol, maxit, M, [], x0, ...)`
`[x, flag, relres, iter, resvec] = cgs (A, b, ...)`

Solve `A x = b`, where $A$ is a square matrix, using the Conjugate Gradients Squared method.

The input arguments are:

- $A$ is the matrix of the linear system and it must be square. $A$ can be passed as a matrix, function handle, or inline function `Afcn` such that `Afcn(x) = A * x`. Additional parameters to `Afcn` are passed after *x0*.
- $b$ is the right hand side vector. It must be a column vector with same number of rows of $A$.
- *tol* is the relative tolerance, if not given or set to [] the default value 1e-6 is used.
- *maxit* the maximum number of outer iterations, if not given or set to [] the default value `min (20, numel (b))` is used.
- *M1*, *M2* are the preconditioners. The preconditioner matrix is given as `M = M1 * M2`. Both *M1* and *M2* can be passed as a matrix or as a function handle or inline function `g` such that `g(x) = M1 \ x` or `g(x) = M2 \ x`. If M1 is empty or not passed then no preconditioners are applied. The technique used is the right preconditioning, i.e., it is solved `A*inv(M)*y = b` and then `x = inv(M)*y`.

– *x0* the initial guess, if not given or set to [] the default value `zeros (size (b))` is used.

The arguments which follow *x0* are treated as parameters, and passed in a proper way to any of the functions ($A$ or $P$) which are passed to `cgs`.

The output parameters are:

– $x$ is the approximation computed. If the method doesn't converge then it is the iterated with the minimum residual.
– *flag* indicates the exit status:
  – 0: iteration converged to the within the chosen tolerance
  – 1: the maximum number of iterations was reached before convergence
  – 2: the preconditioner matrix is singular
  – 3: the algorithm reached stagnation
  – 4: the algorithm can't continue due to a division by zero
– *relres* is the relative residual obtained with as `(A*x-b) / norm(b)`.
– *iter* is the iteration which $x$ is computed.
– *resvec* is a vector containing the residual at each iteration. Doing `length(resvec) - 1` is possible to see the total number of iterations performed.

Let us consider a trivial problem with a tridiagonal matrix

```
n = 20;
A = toeplitz (sparse ([1, 1], [1, 2], [2, 1] * n ^ 2, 1, n))  + ...
    toeplitz (sparse (1, 2, -1, 1, n) * n / 2, ...
    sparse (1, 2, 1, 1, n) * n / 2);
b = A * ones (n, 1);
restart = 5;
[M1, M2] = ilu (A); # in this tridiag case it corresponds to chol (A)'
M = M1 * M2;
Afcn = @(x) A * x;
Mfcn = @(x) M \ x;
M1fcn = @(x) M1 \ x;
M2fcn = @(x) M2 \ x;
```

Example 1: simplest usage of `cgs`

```
x = cgs (A, b, [], n)
```

Example 2: `cgs` with a function which computes `A * x`

```
x = cgs (Afcn, b, [], n)
```

Example 3: `cgs` with a preconditioner matrix $M$

```
x = cgs (A, b, [], 1e-06, n, M)
```

Example 4: `cgs` with a function as preconditioner

```
x = cgs (Afcn, b, 1e-6, n, Mfcn)
```

Example 5: `cgs` with preconditioner matrices *M1* and *M2*

```
x = cgs (A, b, [], 1e-6, n, M1, M2)
```

Example 6: `cgs` with functions as preconditioners

```
x = cgs (Afcn, b, 1e-6, n, M1fcn, M2fcn)
```

Example 7: `cgs` with as input a function requiring an argument

```
function y = Ap (A, x, z) # compute A^z * x
   y = x;
   for i = 1:z
     y = A * y;
   endfor
 endfunction
Apfcn = @(x, string, p) Ap (A, x, string, p);
x = cgs (Apfcn, b, [], [], [], [], [], 2);
```

Example 8: explicit example to show that `cgs` uses a right preconditioner

```
[M1, M2] = ilu (A + 0.3 * eye (n)); # factorization of A perturbed
M = M1 * M2;

## reference solution computed by cgs after one iteration
[x_ref, fl] = cgs (A, b, [], 1, M)

## right preconditioning
[y, fl] = cgs (A / M, b, [], 1)
x = M \ y # compare x and x_ref
```

References:

Y. Saad, *Iterative Methods for Sparse Linear Systems*, 2nd edition, SIAM, 2003.

**See also:** [pcg], page 758, [bicgstab], page 678, [bicg], page 676, [gmres], page 682, [qmr], page 685, [tfqmr], page 686.

`x = gmres (A, b, restart, tol, maxit, M1, M2, x0, ...)`
`x = gmres (A, b, restart, tol, maxit, M, [], x0, ...)`
`[x, flag, relres, iter, resvec] = gmres (A, b, ...)`

Solve `A x = b` using the Preconditioned GMRES iterative method with restart, a.k.a. PGMRES(restart).

The input arguments are:

- *A* is the matrix of the linear system and it must be square. *A* can be passed as a matrix, function handle, or inline function `Afcn` such that `Afcn(x) = A * x`. Additional parameters to `Afcn` are passed after *x0*.
- *b* is the right hand side vector. It must be a column vector with the same numbers of rows as *A*.
- *restart* is the number of iterations before that the method restarts. If it is [] or N = numel (b), then the restart is not applied.
- *tol* is the required relative tolerance for the preconditioned residual error, `inv (M) * (b - a * x)`. The iteration stops if `norm (inv (M) * (b - a * x))` $\leq$ *tol* * `norm (inv (M) * B)`. If *tol* is omitted or empty, then a tolerance of 1e-6 is used.
- *maxit* is the maximum number of outer iterations, if not given or set to [], then the default value `min (10, N / restart)` is used. Note that, if *restart* is empty, then *maxit* is the maximum number of iterations. If *restart* and *maxit* are not

empty, then the maximum number of iterations is `restart * maxit`. If both *restart* and *maxit* are empty, then the maximum number of iterations is set to `min (10, N)`.

- *M1*, *M2* are the preconditioners. The preconditioner $M$ is given as `M = M1 * M2`. Both *M1* and *M2* can be passed as a matrix, function handle, or inline function `g` such that `g(x) = M1 \ x` or `g(x) = M2 \ x`. If *M1* is [] or not given, then the preconditioner is not applied. The technique used is the left-preconditioning, i.e., it is solved `inv(M) * A * x = inv(M) * b` instead of `A * x = b`.
- *x0* is the initial guess, if not given or set to [], then the default value `zeros (size (b))` is used.

The arguments which follow *x0* are treated as parameters, and passed in a proper way to any of the functions ($A$ or $M$ or *M1* or *M2*) which are passed to `gmres`.

The outputs are:

- $x$ the computed approximation. If the method does not converge, then it is the iterated with minimum residual.
- *flag* indicates the exit status:

  | | |
  |---|---|
  | 0 | iteration converged to within the specified tolerance |
  | 1 | maximum number of iterations exceeded |
  | 2 | the preconditioner matrix is singular |
  | 3 | algorithm reached stagnation (the relative difference between two consecutive iterations is less than eps) |

- *relres* is the value of the relative preconditioned residual of the approximation $x$.
- *iter* is a vector containing the number of outer iterations and inner iterations performed to compute $x$. That is:
  - *iter(1)*: number of outer iterations, i.e., how many times the method restarted. (if *restart* is empty or $N$, then it is 1, if not $1 \leq$ *iter(1)* $\leq$ *maxit*).
  - *iter(2)*: the number of iterations performed before the restart, i.e., the method restarts when `iter(2) = restart`. If *restart* is empty or $N$, then $1 \leq$ *iter(2)* $\leq$ *maxit*.

  To be more clear, the approximation $x$ is computed at the iteration `(iter(1) - 1) * restart + iter(2)`. Since the output $x$ corresponds to the minimal preconditioned residual solution, the total number of iterations that the method performed is given by `length (resvec) - 1`.
- *resvec* is a vector containing the preconditioned relative residual at each iteration, including the 0-th iteration `norm (A * x0 - b)`.

Let us consider a trivial problem with a tridiagonal matrix

```
n = 20;
A = toeplitz (sparse ([1, 1], [1, 2], [2, 1] * n ^ 2, 1, n))  + ...
    toeplitz (sparse (1, 2, -1, 1, n) * n / 2, ...
    sparse (1, 2, 1, 1, n) * n / 2);
b = A * ones (n, 1);
restart = 5;
[M1, M2] = ilu (A); # in this tridiag case, it corresponds to lu (A)
M = M1 * M2;
Afcn = @(x) A * x;
Mfcn = @(x) M \ x;
M1fcn = @(x) M1 \ x;
M2fcn = @(x) M2 \ x;
```

Example 1: simplest usage of `gmres`

```
x = gmres (A, b, [], [], n)
```

Example 2: `gmres` with a function which computes `A * x`

```
x = gmres (Afcn, b, [], [], n)
```

Example 3: usage of `gmres` with the restart

```
x = gmres (A, b, restart);
```

Example 4: `gmres` with a preconditioner matrix $M$ with and without restart

```
x = gmres (A, b, [], 1e-06, n, M)
x = gmres (A, b, restart, 1e-06, n, M)
```

Example 5: `gmres` with a function as preconditioner

```
x = gmres (Afcn, b, [], 1e-6, n, Mfcn)
```

Example 6: `gmres` with preconditioner matrices *M1* and *M2*

```
x = gmres (A, b, [], 1e-6, n, M1, M2)
```

Example 7: `gmres` with functions as preconditioners

```
x = gmres (Afcn, b, 1e-6, n, M1fcn, M2fcn)
```

Example 8: `gmres` with as input a function requiring an argument

```
  function y = Ap (A, x, p) # compute A^p * x
     y = x;
     for i = 1:p
       y = A * y;
     endfor
  endfunction
Apfcn = @(x, p) Ap (A, x, p);
x = gmres (Apfcn, b, [], [], [], [], [], [], 2);
```

Example 9: explicit example to show that `gmres` uses a left preconditioner

```
[M1, M2] = ilu (A + 0.1 * eye (n)); # factorization of A perturbed
M = M1 * M2;

## reference solution computed by gmres after two iterations
[x_ref, fl] = gmres (A, b, [], [], 1, M)

## left preconditioning
[x, fl] = gmres (M \ A, M \ b, [], [], 1)
x # compare x and x_ref
```

Reference:

Y. Saad, *Iterative Methods for Sparse Linear Systems*, 2nd edition, SIAM, 2003.

**See also:** [bicg], page 676, [bicgstab], page 678, [cgs], page 680, [pcg], page 758, [pcr], page 761, [qmr], page 685, [tfqmr], page 686.

*x* = qmr (*A*, *b*, *rtol*, *maxit*, *M1*, *M2*, *x0*)
*x* = qmr (*A*, *b*, *rtol*, *maxit*, *P*)
[*x*, *flag*, *relres*, *iter*, *resvec*] = qmr (*A*, *b*, ...)

Solve `A x = b` using the Quasi-Minimal Residual iterative method (without look-ahead).

- *rtol* is the relative tolerance, if not given or set to [] the default value 1e-6 is used.
- *maxit* the maximum number of outer iterations, if not given or set to [] the default value `min (20, numel (b))` is used.
- *x0* the initial guess, if not given or set to [] the default value `zeros (size (b))` is used.

*A* can be passed as a matrix or as a function handle or inline function `f` such that `f(x, "notransp") = A*x` and `f(x, "transp") = A'*x`.

The preconditioner *P* is given as `P = M1 * M2`. Both *M1* and *M2* can be passed as a matrix or as a function handle or inline function `g` such that `g(x, "notransp") = M1 \ x` or `g(x, "notransp") = M2 \ x` and `g(x, "transp") = M1' \ x` or `g(x, "transp") = M2' \ x`.

If called with more than one output parameter

- *flag* indicates the exit status:
  - 0: iteration converged to the within the chosen tolerance
  - 1: the maximum number of iterations was reached before convergence
  - 3: the algorithm reached stagnation

  (the value 2 is unused but skipped for compatibility).
- *relres* is the final value of the relative residual.
- *iter* is the number of iterations performed.
- *resvec* is a vector containing the residual norms at each iteration.

References:

1. R. Freund and N. Nachtigal, "QMR: a quasi-minimal residual method for non-Hermitian linear systems", *Numerische Mathematik*, 60, pp. 315–339, 1991.

2. R. Barrett, M. Berry, T. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhour, R. Pozo, C. Romine, and H. van der Vorst, *Templates for the solution of linear systems: Building blocks for iterative methods*, SIAM, 2nd ed., 1994.

**See also:** [bicg], page 676, [bicgstab], page 678, [cgs], page 680, [gmres], page 682, [pcg], page 758.

`x = tfqmr (A, b, tol, maxit, M1, M2, x0, ...)`
`x = tfqmr (A, b, tol, maxit, M, [], x0, ...)`
`[x, flag, relres, iter, resvec] = tfqmr (A, b, ...)`

Solve `A x = b` using the Transpose-Tree qmr method, based on the cgs.

The input parameters are:

- *A* is the matrix of the linear system and it must be square. *A* can be passed as a matrix, function handle, or inline function `Afcn` such that `Afcn(x) = A * x`. Additional parameters to `Afcn` are passed after *x0*.
- *b* is the right hand side vector. It must be a column vector with the same number of rows as *A*.
- *tol* is the relative tolerance, if not given or set to [] the default value 1e-6 is used.
- *maxit* the maximum number of outer iterations, if not given or set to [] the default value `min (20, numel (b))` is used. To be compatible, since the method as different behaviors in the iteration number is odd or even, is considered as iteration in `tfqmr` the entire odd-even cycle. That is, to make an entire iteration, the algorithm performs two sub-iterations: the odd one and the even one.
- *M1*, *M2* are the preconditioners. The preconditioner *M* is given as `M = M1 * M2`. Both *M1* and *M2* can be passed as a matrix or as a function handle or inline function `g` such that `g(x) = M1 \ x` or `g(x) = M2 \ x`. The technique used is the right-preconditioning, i.e., it is solved `A*inv(M)*y = b` and then `x = inv(M)*y`, instead of `A x = b`.
- *x0* the initial guess, if not given or set to [] the default value `zeros (size (b))` is used.

The arguments which follow *x0* are treated as parameters, and passed in a proper way to any of the functions (*A* or *M*) which are passed to `tfqmr`.

The output parameters are:

- *x* is the approximation computed. If the method doesn't converge then it is the iterated with the minimum residual.
- *flag* indicates the exit status:
  - 0: iteration converged to the within the chosen tolerance
  - 1: the maximum number of iterations was reached before convergence
  - 2: the preconditioner matrix is singular
  - 3: the algorithm reached stagnation
  - 4: the algorithm can't continue due to a division by zero
- *relres* is the relative residual obtained as `(A*x-b) / norm (b)`.
- *iter* is the iteration which *x* is computed.

– *resvec* is a vector containing the residual at each iteration (including `norm (b - A x0)`). Doing `length (resvec) - 1` is possible to see the total number of iterations performed.

Let us consider a trivial problem with a tridiagonal matrix

```
n = 20;
A = toeplitz (sparse ([1, 1], [1, 2], [2, 1] * n ^ 2, 1, n))  + ...
    toeplitz (sparse (1, 2, -1, 1, n) * n / 2, ...
    sparse (1, 2, 1, 1, n) * n / 2);
b = A * ones (n, 1);
restart = 5;
[M1, M2] = ilu (A); # in this tridiag case it corresponds to chol (A)'
M = M1 * M2;
Afcn = @(x) A * x;
Mfcn = @(x) M \ x;
M1fcn = @(x) M1 \ x;
M2fcn = @(x) M2 \ x;
```

Example 1: simplest usage of `tfqmr`

```
x = tfqmr (A, b, [], n)
```

Example 2: `tfqmr` with a function which computes *A * x*

```
x = tfqmr (Afcn, b, [], n)
```

Example 3: `tfqmr` with a preconditioner matrix *M*

```
x = tfqmr (A, b, [], 1e-06, n, M)
```

Example 4: `tfqmr` with a function as preconditioner

```
x = tfqmr (Afcn, b, 1e-6, n, Mfcn)
```

Example 5: `tfqmr` with preconditioner matrices *M1* and *M2*

```
x = tfqmr (A, b, [], 1e-6, n, M1, M2)
```

Example 6: `tfmqr` with functions as preconditioners

```
x = tfqmr (Afcn, b, 1e-6, n, M1fcn, M2fcn)
```

Example 7: `tfqmr` with as input a function requiring an argument

```
function y = Ap (A, x, z) # compute A^z * x
   y = x;
   for i = 1:z
     y = A * y;
   endfor
 endfunction
Apfcn = @(x, string, p) Ap (A, x, string, p);
x = tfqmr (Apfcn, b, [], [], [], [], [], 2);
```

Example 8: explicit example to show that `tfqmr` uses a right preconditioner

```
[M1, M2] = ilu (A + 0.3 * eye (n)); # factorization of A perturbed
M = M1 * M2;

## reference solution computed by tfqmr after one iteration
[x_ref, fl] = tfqmr (A, b, [], 1, M)

## right preconditioning
[y, fl] = tfqmr (A / M, b, [], 1)
x = M \ y # compare x and x_ref
```

Reference:

Y. Saad, *Iterative Methods for Sparse Linear Systems*, 2nd edition, SIAM, 2003.

**See also:** [bicg], page 676, [bicgstab], page 678, [cgs], page 680, [gmres], page 682, [pcg], page 758, [qmr], page 685, [pcr], page 761.

# 19 Vectorization and Faster Code Execution

Vectorization is a programming technique that uses vector operations instead of element-by-element loop-based operations. Besides frequently producing more succinct Octave code, vectorization also allows for better optimization in the subsequent implementation. The optimizations may occur either in Octave's own Fortran, C, or C++ internal implementation, or even at a lower level depending on the compiler and external numerical libraries used to build Octave. The ultimate goal is to make use of your hardware's vector instructions if possible or to perform other optimizations in software.

Vectorization is not a concept unique to Octave, but it is particularly important because Octave is a matrix-oriented language. Vectorized Octave code will see a dramatic speed up (10X–100X) in most cases.

This chapter discusses vectorization and other techniques for writing faster code.

## 19.1 Basic Vectorization

To a very good first approximation, the goal in vectorization is to write code that avoids loops and uses whole-array operations. As a trivial example, consider

```
for i = 1:n
  for j = 1:m
    c(i,j) = a(i,j) + b(i,j);
  endfor
endfor
```

compared to the much simpler

```
c = a + b;
```

This isn't merely easier to write; it is also internally much easier to optimize. Octave delegates this operation to an underlying implementation which, among other optimizations, may use special vector hardware instructions or could conceivably even perform the additions in parallel. In general, if the code is vectorized, the underlying implementation has more freedom about the assumptions it can make in order to achieve faster execution.

This is especially important for loops with "cheap" bodies. Often it suffices to vectorize just the innermost loop to get acceptable performance. A general rule of thumb is that the "order" of the vectorized body should be greater or equal to the "order" of the enclosing loop.

As a less trivial example, instead of

```
for i = 1:n-1
  a(i) = b(i+1) - b(i);
endfor
```

write

```
a = b(2:n) - b(1:n-1);
```

This shows an important general concept about using arrays for indexing instead of looping over an index variable. See Section 8.1 [Index Expressions], page 159. Also use

boolean indexing generously. If a condition needs to be tested, this condition can also be written as a boolean index. For instance, instead of

```
for i = 1:n
  if (a(i) > 5)
    a(i) -= 20
  endif
endfor
```

write

```
a(a>5) -= 20;
```

which exploits the fact that `a > 5` produces a boolean index.

Use elementwise vector operators whenever possible to avoid looping (operators like `.*` and `.^`). See Section 8.3 [Arithmetic Ops], page 170.

Also exploit broadcasting in these elementwise operators both to avoid looping and unnecessary intermediate memory allocations. See Section 19.2 [Broadcasting], page 691.

Use built-in and library functions if possible. Built-in and compiled functions are very fast. Even with an m-file library function, chances are good that it is already optimized, or will be optimized more in a future release.

For instance, even better than

```
a = b(2:n) - b(1:n-1);
```

is

```
a = diff (b);
```

Most Octave functions are written with vector and array arguments in mind. If you find yourself writing a loop with a very simple operation, chances are that such a function already exists. The following functions occur frequently in vectorized code:

- Index manipulation
  - find
  - sub2ind
  - ind2sub
  - sort
  - unique
  - lookup
  - ifelse / merge
- Repetition
  - repmat
  - repelems
- Vectorized arithmetic
  - sum
  - prod
  - cumsum
  - cumprod

  - sumsq
  - diff
  - dot
  - cummax
  - cummin
- Shape of higher dimensional arrays
  - reshape
  - resize
  - permute
  - squeeze
  - deal

## 19.2 Broadcasting

Broadcasting refers to how Octave binary operators and functions behave when their matrix or array operands or arguments differ in size. Since version 3.6.0, Octave now automatically broadcasts vectors, matrices, and arrays when using elementwise binary operators and functions. Broadly speaking, smaller arrays are "broadcast" across the larger one, until they have a compatible shape. The rule is that corresponding array dimensions must either

1. be equal, or
2. one of them must be 1.

In case all dimensions are equal, no broadcasting occurs and ordinary element-by-element arithmetic takes place. For arrays of higher dimensions, if the number of dimensions isn't the same, then missing trailing dimensions are treated as 1. When one of the dimensions is 1, the array with that singleton dimension gets copied along that dimension until it matches the dimension of the other array. For example, consider

```
x = [1 2 3;
     4 5 6;
     7 8 9];

y = [10 20 30];

x + y
```

Without broadcasting, `x + y` would be an error because the dimensions do not agree. However, with broadcasting it is as if the following operation were performed:

```
x = [1 2 3
     4 5 6
     7 8 9];

y = [10 20 30
     10 20 30
     10 20 30];

x + y
⇒    11   22   33
     14   25   36
     17   28   39
```

That is, the smaller array of size `[1 3]` gets copied along the singleton dimension (the number of rows) until it is `[3 3]`. No actual copying takes place, however. The internal implementation reuses elements along the necessary dimension in order to achieve the desired effect without copying in memory.

Both arrays can be broadcast across each other, for example, all pairwise differences of the elements of a vector with itself:

```
y - y'
⇒    0   10   20
   -10    0   10
   -20  -10    0
```

Here the vectors of size `[1 3]` and `[3 1]` both get broadcast into matrices of size `[3 3]` before ordinary matrix subtraction takes place.

A special case of broadcasting that may be familiar is when all dimensions of the array being broadcast are 1, i.e., the array is a scalar. Thus, for example, operations like `x - 42` and `max (x, 2)` are basic examples of broadcasting.

For a higher-dimensional example, suppose `img` is an RGB image of size `[m n 3]` and we wish to multiply each color by a different scalar. The following code accomplishes this with broadcasting,

```
img .*= permute ([0.8, 0.9, 1.2], [1, 3, 2]);
```

Note the usage of permute to match the dimensions of the `[0.8, 0.9, 1.2]` vector with `img`.

For functions that are not written with broadcasting semantics, `bsxfun` can be useful for coercing them to broadcast.

*C* = **bsxfun** (*f*, *A*, *B*)

Apply a binary function *f* element-by-element to two array arguments *A* and *B*, expanding singleton dimensions in either input argument as necessary.

*f* is a function handle, inline function, or string containing the name of the function to evaluate. The function *f* must be capable of accepting two column-vector arguments of equal length, or one column vector argument and a scalar.

The dimensions of *A* and *B* must be equal or singleton. The singleton dimensions of the arrays will be expanded to the same dimensionality as the other array.

**See also:** [arrayfun], page 695, [cellfun], page 697.

Broadcasting is only applied if either of the two broadcasting conditions hold. As usual, however, broadcasting does not apply when two dimensions differ and neither is 1:

```
x = [1 2 3
     4 5 6];
y = [10 20
     30 40];
x + y
```

This will produce an error about nonconformant arguments.

Besides common arithmetic operations, several functions of two arguments also broadcast. The full list of functions and operators that broadcast is

```
plus      +
minus     -
times     .*
rdivide   ./
ldivide   .\
power     .^
lt        <
le        <=
eq        ==
gt        >
ge        >=
ne        !=  ~=
and       &
or        |
atan2
hypot
max
min
mod
rem
xor

+=  -=  .*=  ./=  .\=  .^=  &=  |=
```

Here is a real example of the power of broadcasting. The Floyd-Warshall algorithm is used to calculate the shortest path lengths between every pair of vertices in a graph. A naive implementation for a graph adjacency matrix of order $n$ might look like this:

```
for k = 1:n
  for i = 1:n
    for j = 1:n
      dist(i,j) = min (dist(i,j), dist(i,k) + dist(k,j));
    endfor
  endfor
endfor
```

Upon vectorizing the innermost loop, it might look like this:

```
for k = 1:n
  for i = 1:n
    dist(i,:) = min (dist(i,:), dist(i,k) + dist(k,:));
  endfor
endfor
```

Using broadcasting in both directions, it looks like this:

```
for k = 1:n
  dist = min (dist, dist(:,k) + dist(k,:));
endfor
```

The relative time performance of the three techniques for a given graph with 100 vertices is 7.3 seconds for the naive code, 87 milliseconds for the singly vectorized code, and 1.3 milliseconds for the fully broadcast code. For a graph with 1000 vertices, vectorization takes 11.7 seconds while broadcasting takes only 1.15 seconds. Therefore in general it is worth writing code with broadcasting semantics for performance.

However, beware of resorting to broadcasting if a simpler operation will suffice. For matrices *a* and *b*, consider the following:

```
c = sum (permute (a, [1, 3, 2]) .* permute (b, [3, 2, 1]), 3);
```

This operation broadcasts the two matrices with permuted dimensions across each other during elementwise multiplication in order to obtain a larger 3-D array, and this array is then summed along the third dimension. A moment of thought will prove that this operation is simply the much faster ordinary matrix multiplication, `c = a*b;`.

A note on terminology: "broadcasting" is the term popularized by the Numpy numerical environment in the Python programming language. In other programming languages and environments, broadcasting may also be known as *binary singleton expansion* (BSX, in MATLAB, and the origin of the name of the `bsxfun` function), *recycling* (R programming language), *single-instruction multiple data* (SIMD), or *replication*.

### 19.2.1 Broadcasting and Legacy Code

The new broadcasting semantics almost never affect code that worked in previous versions of Octave. Consequently, all code inherited from MATLAB that worked in previous versions of Octave should still work without change in Octave. The only exception is code such as

```
try
  c = a.*b;
catch
  c = a.*a;
end_try_catch
```

that may have relied on matrices of different size producing an error. Because such operation is now valid Octave syntax, this will no longer produce an error. Instead, the following code should be used:

```
if (isequal (size (a), size (b)))
  c = a .* b;
else
  c = a .* a;
endif
```

## 19.3 Function Application

As a general rule, functions should already be written with matrix arguments in mind and should consider whole matrix operations in a vectorized manner. Sometimes, writing functions in this way appears difficult or impossible for various reasons. For those situations, Octave provides facilities for applying a function to each element of an array, cell, or struct.

*B* = arrayfun (*fcn*, *A*)
*B* = arrayfun (*fcn*, *A1*, *A2*, ...)
[*B1*, *B2*, ...] = arrayfun (*fcn*, *A*, ...)
*B* = arrayfun (..., "*UniformOutput*", *val*)
*B* = arrayfun (..., "*ErrorHandler*", *errfcn*)

Execute a function on each element of an array.

This is useful for functions that do not accept array arguments. If the function does accept array arguments it is *better* to call the function directly.

The first input argument *fcn* can be a string, a function handle, an inline function, or an anonymous function. The input argument *A* can be a logical array, a numeric array, a string array, a structure array, or a cell array. `arrayfun` passes all elements of *A* individually to the function *fcn* and collects the results. The equivalent pseudo-code is

```
cls = class (fcn (A(1));
B = zeros (size (A), cls);
for i = 1:numel (A)
  B(i) = fcn (A(i))
endfor
```

The named function can also take more than two input arguments, with the input arguments given as third input argument *A2*, fourth input argument *A2*, ... If given more than one array input argument then all input arguments must have the same sizes. For example:

```
arrayfun (@atan2, [1, 0], [0, 1])
     ⇒ [ 1.57080   0.00000 ]
```

If the parameter *val* after a further string input argument `"UniformOutput"` is set `true` (the default), then the named function *fcn* must return a single element which then will be concatenated into the return value and is of type matrix. Otherwise, if that parameter is set to `false`, then the outputs are concatenated in a cell array. For example:

```
arrayfun (@(x,y) x:y, "abc", "def", "UniformOutput", false)
⇒
   {
     [1,1] = abcd
     [1,2] = bcde
     [1,3] = cdef
   }
```

If more than one output arguments are given then the named function must return the number of return values that also are expected, for example:

```
[A, B, C] = arrayfun (@find, [10; 0], "UniformOutput", false)
⇒
A =
{
   [1,1] =  1
   [2,1] = [](0x0)
}
B =
{
   [1,1] =  1
   [2,1] = [](0x0)
}
C =
{
   [1,1] =  10
   [2,1] = [](0x0)
}
```

If the parameter *errfcn* after a further string input argument `"ErrorHandler"` is another string, a function handle, an inline function, or an anonymous function, then *errfcn* defines a function to call in the case that *fcn* generates an error. The definition of the function must be of the form

```
function [...] = errfcn (s, ...)
```

where there is an additional input argument to *errfcn* relative to *fcn*, given by *s*. This is a structure with the elements `"identifier"`, `"message"`, and `"index"` giving, respectively, the error identifier, the error message, and the index of the array elements that caused the error. The size of the output argument of *errfcn* must have the same size as the output argument of *fcn*, otherwise a real error is thrown. For example:

```
function y = ferr (s, x), y = "MyString"; endfunction
arrayfun (@str2num, [1234],
          "UniformOutput", false, "ErrorHandler", @ferr)
⇒
   {
     [1,1] = MyString
   }
```

**See also:** [spfun], page 696, [cellfun], page 697, [structfun], page 699.

`y = spfun (f, S)`

Compute `f (S)` for the nonzero elements of *S*.

The input function *f* is applied only to the nonzero elements of the input matrix *S* which is typically sparse. The function *f* can be passed as a string, function handle, or inline function.

The output *y* is a sparse matrix with the same sparsity structure as the input *S*. `spfun` preserves sparsity structure which is different than simply applying the function *f* to the sparse matrix *S* when `f (0) != 0`.

Example

Sparsity preserving `spfun` versus normal function application

```
S = pi * speye (2,2)
S =

Compressed Column Sparse (rows = 2, cols = 2, nnz = 2 [50%])

  (1, 1) -> 3.1416
  (2, 2) -> 3.1416

y = spfun (@cos, S)
y =

Compressed Column Sparse (rows = 2, cols = 2, nnz = 2 [50%])

  (1, 1) -> -1
  (2, 2) -> -1

y = cos (S)
y =

Compressed Column Sparse (rows = 2, cols = 2, nnz = 4 [100%])

  (1, 1) -> -1
  (2, 1) -> 1
  (1, 2) -> 1
  (2, 2) -> -1
```

**See also:** [arrayfun], page 695, [cellfun], page 697, [structfun], page 699.

`A = cellfun (@fcn, C)`
`A = cellfun ("fcn", C)`
`A = cellfun (fcn, C)`
`A = cellfun (fcn, C1, C2, ...)`
`[A1, A2, ...] = cellfun (...)`
`A = cellfun (..., "UniformOutput", val)`
`A = cellfun (..., "ErrorHandler", errfcn)`
`A = cellfun ('isempty', C)`
`A = cellfun ('islogical', C)`
`A = cellfun ('isnumeric', C)`
`A = cellfun ('isreal', C)`
`A = cellfun ('length', C)`
`A = cellfun ('numel', C)`
`A = cellfun ('prodofsize', C)`
`A = cellfun ('size', C, dim)`
`A = cellfun ('isclass', C, class)`

Evaluate the function *fcn* on the elements of the cell array *C*.

`cellfun` accepts an arbitrary function *fcn* given as a name in a character string, as a function handle, or as an inline function. Specifying *fcn* with a character string is preferred as the performance is ~3X better for builtin functions and equivalent for m-files.

`cellfun` has a limited number of functions which have been specially-coded for high-performance (~8X faster than a function handle). These functions are only used if the function is specified by name as a string, and only the simplest calling form—`cellfun ('fcn', C)`—without options is supported. If you need access to an overloaded version of a function, such as `numel` for a particular `classdef` file, then you cannot use the accelerated function name and must use a function handle instead, e.g., `@numel`.

The high-performance functions are

`'isempty'`
: Return true for empty elements.

`'islogical'`
: Return true for logical elements.

`'isnumeric'`
: Return true for numeric elements.

`'isreal'` Return true for real elements.

`'length'` Return a vector of the lengths of cell elements.

`'ndims'` Return the number of dimensions of each element.

`'numel'`
`'prodofsize'`
: Return the number of elements contained within each cell element. The number is the product of the dimensions of the object of each cell element.

`'size'` Return the size along dimension *dim*.

`'isclass'`
: Return true for elements which are of type *class*.

Elements in *C* are passed to the function individually and the result of each function invocation is collected in the output. The function can take more than one argument with the inputs arguments given by *C1*, *C2*, etc. Input arguments that are singleton (1x1) cells will be automatically expanded to the size of the other arguments. For example:

```
cellfun ("atan2", {1, 0}, {0, 1})
     ⇒ [ 1.57080   0.00000 ]
```

The number of output arguments of `cellfun` matches the number of output arguments of the function and can be greater than one. When there are multiple outputs of the function they will be collected into the output arguments of `cellfun` like this:

```
function [a, b] = twoouts (x)
  a = x;
  b = x*x;
endfunction
[aa, bb] = cellfun (@twoouts, {1, 2, 3})
     ⇒
        aa =
           1 2 3
        bb =
           1 4 9
```

Note that, by default, the output argument(s) are arrays of the same size as the input arguments.

If the parameter `"UniformOutput"` is set to true (the default), then the function must return scalars which will be concatenated into the return array(s). If `"UniformOutput"` is false, the outputs are concatenated into a cell array (or cell arrays). For example:

```
cellfun ("lower", {"Foo", "Bar", "FooBar"},
         "UniformOutput", false)
⇒ {"foo", "bar", "foobar"}
```

The parameter `"ErrorHandler"` specifies a function *errfcn* to call if *fcn* generates an error. The form of the function is

```
function [...] = errfcn (s, ...)
```

where there is an additional input argument to *errfcn* relative to *fcn*, given by *s*. This is a structure with the elements `"identifier"`, `"message"`, and `"index"` giving respectively the error identifier, the error message, and the index into the input arguments of the element that caused the error. For example:

```
function y = errfcn (s, x), y = NaN; endfunction
cellfun ("factorial", {-1, 2}, "ErrorHandler", @errfcn)
⇒ [NaN 2]
```

Programming Note: Use `cellfun` intelligently. The `cellfun` function is a useful tool for avoiding loops. It is often used with anonymous function handles; however, calling an anonymous function involves an overhead quite comparable to the overhead of an m-file function. Passing a handle to a built-in function is faster, because the interpreter is not involved in the internal loop. For example:

```
C = {...}
v = cellfun (@(x) det (x), C); # compute determinants
v = cellfun (@det, C);         # 40% faster
```

**See also:** [arrayfun], page 695, [structfun], page 699, [spfun], page 696.

`A = structfun (`*`fcn`*`, `*`S`*`)`
`A = structfun (...,` "*ErrorHandler*"`,` *`errfcn`*`)`
`A = structfun (...,` "*UniformOutput*"`,` *`val`*`)`
`[`*`A, B, ...`*`] = structfun (...)`

Evaluate the function named *name* on the fields of the structure *S*. The fields of *S* are passed to the function *fcn* individually.

`structfun` accepts an arbitrary function *fcn* in the form of an inline function, function handle, or the name of a function (in a character string). In the case of a character string argument, the function must accept a single argument named *x*, and it must return a string value. If the function returns more than one argument, they are returned as separate output variables.

If the parameter `"UniformOutput"` is set to true (the default), then the function must return a single element which will be concatenated into the return value. If `"UniformOutput"` is false, the outputs are placed into a structure with the same fieldnames as the input structure.

```
s.name1 = "John Smith";
s.name2 = "Jill Jones";
structfun (@(x) regexp (x, '(\w+)$', "matches"){1}, s,
           "UniformOutput", false)
  ⇒ scalar structure containing the fields:
       name1 = Smith
       name2 = Jones
```

Given the parameter `"ErrorHandler"`, *errfcn* defines a function to call in case *fcn* generates an error. The form of the function is

```
function [...] = errfcn (se, ...)
```

where there is an additional input argument to *errfcn* relative to *fcn*, given by *se*. This is a structure with the elements `"identifier"`, `"message"` and `"index"`, giving respectively the error identifier, the error message, and the index into the input arguments of the element that caused the error. For an example on how to use an error handler, see [cellfun], page 697.

**See also:** [cellfun], page 697, [arrayfun], page 695, [spfun], page 696.

Consistent with earlier advice, seek to use Octave built-in functions whenever possible for the best performance. This advice applies especially to the four functions above. For example, when adding two arrays together element-by-element one could use a handle to the built-in addition function `@plus` or define an anonymous function `@(x,y) x + y`. But, the anonymous function is 60% slower than the first method. See Section 34.4.2 [Operator Overloading], page 984, for a list of basic functions which might be used in place of anonymous ones.

## 19.4 Accumulation

Whenever it's possible to categorize according to indices the elements of an array when performing a computation, accumulation functions can be useful.

`A = accumarray (subs, vals)`
`A = accumarray (subs, vals, sz)`
`A = accumarray (subs, vals, sz, fcn)`
`A = accumarray (subs, vals, sz, fcn, fillval)`
`A = accumarray (subs, vals, sz, fcn, fillval, issparse)`

Create an array by accumulating the elements of a vector into the positions defined by their subscripts.

The subscripts are defined by the rows of the matrix *subs* and the values by *vals*. Each row of *subs* corresponds to one of the values in *vals*. If *vals* is a scalar, it will be used for each of the row of *subs*. If *subs* is a cell array of vectors, all vectors must be of the same length, and the subscripts in the *k*th vector must correspond to the *k*th dimension of the result.

The size of the matrix will be determined by the subscripts themselves. However, if *sz* is defined it determines the matrix size. The length of *sz* must correspond to the number of columns in *subs*. An exception is if *subs* has only one column, in which case *sz* may be the dimensions of a vector and the subscripts of *subs* are taken as the indices into it.

The default action of `accumarray` is to sum the elements with the same subscripts. This behavior can be modified by defining the *fcn* function. This should be a function or function handle that accepts a column vector and returns a scalar. The result of the function should not depend on the order of the subscripts.

The elements of the returned array that have no subscripts associated with them are set to zero. Defining *fillval* to some other value allows these values to be defined. This behavior changes, however, for certain values of *fcn*. If *fcn* is `@min` (respectively, `@max`) then the result will be filled with the minimum (respectively, maximum) integer if *vals* is of integral type, logical false (respectively, logical true) if *vals* is of logical type, zero if *fillval* is zero and all values are non-positive (respectively, non-negative), and NaN otherwise.

By default `accumarray` returns a full matrix. If *issparse* is logically true, then a sparse matrix is returned instead.

The following `accumarray` example constructs a frequency table that in the first column counts how many occurrences each number in the second column has, taken from the vector *x*. Note the usage of `unique` for assigning to all repeated elements of *x* the same index (see [unique], page 871).

```
x = [91, 92, 90, 92, 90, 89, 91, 89, 90, 100, 100, 100];
[u, ~, j] = unique (x);
[accumarray(j', 1), u']
  ⇒  2    89
     3    90
     2    91
     2    92
     3   100
```

Another example, where the result is a multi-dimensional 3-D array and the default value (zero) appears in the output:

```
accumarray ([1, 1, 1;
             2, 1, 2;
             2, 3, 2;
             2, 1, 2;
             2, 3, 2], 101:105)
⇒ ans(:,:,1) = [101, 0, 0; 0, 0, 0]
⇒ ans(:,:,2) = [0, 0, 0; 206, 0, 208]
```

The sparse option can be used as an alternative to the `sparse` constructor (see [sparse], page 732). Thus

```
sparse (i, j, sv)
```

can be written with `accumarray` as

```
accumarray ([i, j], sv', [], [], 0, true)
```

For repeated indices, `sparse` adds the corresponding value. To take the minimum instead, use `min` as an accumulator function:

```
accumarray ([i, j], sv', [], @min, 0, true)
```

The complexity of accumarray in general for the non-sparse case is generally O(M+N), where N is the number of subscripts and M is the maximum subscript (linearized in multi-dimensional case). If *fcn* is one of `@sum` (default), `@max`, `@min` or `@(x) {x}`, an optimized code path is used. Note that for general reduction function the interpreter overhead can play a major part and it may be more efficient to do multiple accumarray calls and compute the results in a vectorized manner.

**See also:** [accumdim], page 702, [unique], page 871, [sparse], page 732.

```
A = accumdim (subs, vals)
A = accumdim (subs, vals, dim)
A = accumdim (subs, vals, dim, n)
A = accumdim (subs, vals, dim, n, fcn)
A = accumdim (subs, vals, dim, n, fcn, fillval)
```

Create an array by accumulating the slices of an array into the positions defined by their subscripts along a specified dimension.

The subscripts are defined by the index vector *subs*. The dimension is specified by *dim*. If not given, it defaults to the first non-singleton dimension. The length of *subs* must be equal to `size (vals, dim)`.

The extent of the result matrix in the working dimension will be determined by the subscripts themselves. However, if *n* is defined it determines this extent.

The default action of `accumdim` is to sum the subarrays with the same subscripts. This behavior can be modified by defining the *fcn* function. This should be a function or function handle that accepts an array and a dimension, and reduces the array along this dimension. As a special exception, the built-in `min` and `max` functions can be used directly, and `accumdim` accounts for the middle empty argument that is used in their calling.

The slices of the returned array that have no subscripts associated with them are set to zero. Defining *fillval* to some other value allows these values to be defined.

An example of the use of `accumdim` is:

```
accumdim ([1, 2, 1, 2, 1], [ 7, -10,   4;
                            -5, -12,   8;
                           -12,   2,   8;
                           -10,   9,  -3;
                            -5,  -3, -13])
⇒ [-10,-11,-1;-15,-3,5]
```

**See also:** [accumarray], page 700.

## 19.5 Memoization

Memoization is a technique to cache the results of slow function calls and return the cached value when the function is called with the same inputs again, instead of reevaluating it. It is very common to replace function calls with lookup tables if the same inputs are happening over and over again in a known, predictable way. Memoization is, at its core, an extension of this practice where the lookup table is extended even during runtime for new arguments not seen previously. A basic theoretical background can be found on Wikipedia or any undergraduate-level computer science textbook.

Octave's `memoize` function provides drop-in memoization functionality for any user function or Octave function, including compiled functions.

*mem_fcn_handle* = memoize (*fcn_handle*)

Create a memoized version *mem_fcn_handle* of function *fcn_handle*.

Each call to the memoized version *mem_fcn_handle* checks the inputs against an internally maintained table, and if the inputs have occurred previously, then the result of the function call is returned from the table itself instead of evaluating the full function again. This speeds up the execution of functions that are called with the same inputs multiple times.

For example, here we take a slow user-written function named `slow_fcn` and memoize it to a new handle `cyc`. The first executions of both versions take the same time, but the subsequent executions of the memoized version returns the previously computed value, thus reducing 2.4 seconds of runtime to only 2.4 milliseconds. The final check verifies that the same result was returned from both versions.

```
>> tic; p = slow_fcn (5040); toc
Elapsed time is 2.41244 seconds.
>> tic; p = slow_fcn (5040); toc
Elapsed time is 2.41542 seconds.

>> cyc = memoize (@slow_fcn);
>> tic; r = cyc (5040); toc
Elapsed time is 2.42609 seconds.
>> tic; r = cyc (5040); toc
Elapsed time is 0.00236511 seconds.

>> all (p == r)
ans = 1
```

**See also:** [clearAllMemoizedCaches], page 704.

To memoize a function `z = foo(x, y)`, use this general pattern:

```
foo2 = memoize (@(x, y) foo(x, y));
z = foo2 (x, y);
```

In the above example, the first line creates a memoized version `foo2` of the function `foo`. For simple functions with only trivial wrapping, this line can also be shortened to:

```
foo2 = memoize (@foo);
```

The second line `z = foo2 (x, y);` calls that memoized version `foo2` instead of the original function, allowing `memoize` to intercept the call and replace it with a looked-up value from a table if the inputs have occurred before, instead of evaluating the original function again.

Note that this will not accelerate the *first* call to the function but only subsequent calls.

Note that due to the overhead incurred by `memoize` to create and manage the lookup tables for each function, this technique is useful only for functions that take at least a couple of seconds to execute. Such functions can be replaced by table lookups taking only a millisecond or so, but if the original function itself was taking only milliseconds, memoizing it will not speed it up.

Recursive functions can be memoized as well, using a pattern like:

```
function z = foo (x, y)
  persistent foo2 = memoize (@foo);
  foo2.CacheSize = 1e6;

  ## Call the memoized version when recursing
  z = foo2 (x, y);
endfunction
```

The `CacheSize` can be optionally increased in anticipation of a large number of function calls, such as from inside a recursive function. If `CacheSize` is exceeded, the memoization tables are resized, causing a slowdown. Increasing the `CacheSize` thus works like preallocation to speed up execution.

The function `clearAllMemoizedCaches` clears the memoization tables when they are no longer needed.

`clearAllMemoizedCaches ()`

Clear all memoized caches.

Memoization maintains internal tables of which functions have been called with which inputs. This function clears those tables to free memory, or for a fresh start.

**See also:** [memoize], page 703.

## 19.6 Miscellaneous Techniques

Here are some other ways of improving the execution speed of Octave programs.

- Avoid computing costly intermediate results multiple times. Octave currently does not eliminate common subexpressions. Also, certain internal computation results are cached for variables. For instance, if a matrix variable is used multiple times as an index, checking the indices (and internal conversion to integers) is only done once.
- Be aware of lazy copies (copy-on-write). When a copy of an object is created, the data is not immediately copied, but rather shared. The actual copying is postponed until the copied data needs to be modified. For example:

```
a = zeros (1000); # create a 1000x1000 matrix
b = a; # no copying done here
b(1) = 1; # copying done here
```

Lazy copying applies to whole Octave objects such as matrices, cells, struct, and also individual cell or struct elements (not array elements).

Additionally, index expressions also use lazy copying when Octave can determine that the indexed portion is contiguous in memory. For example:

```
a = zeros (1000); # create a 1000x1000 matrix
b = a(:,10:100);  # no copying done here
b = a(10:100,:);  # copying done here
```

This applies to arrays (matrices), cell arrays, and structs indexed using '`()`'. Index expressions generating comma-separated lists can also benefit from shallow copying in some cases. In particular, when *a* is a struct array, expressions like `{a.x}, {a(:,2).x}` will use lazy copying, so that data can be shared between a struct array and a cell array.

Most indexing expressions do not live longer than their parent objects. In rare cases, however, a lazily copied slice outlasts its parent, in which case it becomes orphaned, still occupying unnecessarily more memory than needed. To provide a remedy working in most real cases, Octave checks for orphaned lazy slices at certain situations, when a value is stored into a "permanent" location, such as a named variable or cell or struct element, and possibly economizes them. For example:

```
a = zeros (1000); # create a 1000x1000 matrix
b = a(:,10:100);  # lazy slice
a = []; # the original "a" array is still allocated
c{1} = b; # b is reallocated at this point
```

- Avoid deep recursion. Function calls to m-file functions carry a relatively significant overhead, so rewriting a recursion as a loop often helps. Also, note that the maximum level of recursion is limited.
- Avoid resizing matrices unnecessarily. When building a single result matrix from a series of calculations, set the size of the result matrix first, then insert values into it. Write

```
result = zeros (big_n, big_m)
for i = over:and_over
  ridx = ...
  cidx = ...
  result(ridx, cidx) = new_value ();
endfor
```

instead of

```
result = [];
for i = ever:and_ever
  result = [ result, new_value() ];
endfor
```

Sometimes the number of items can not be computed in advance, and stack-like operations are needed. When elements are being repeatedly inserted or removed from the end of an array, Octave detects it as stack usage and attempts to use a smarter memory management strategy by pre-allocating the array in bigger chunks. This strategy is also applied to cell and struct arrays.

```
a = [];
while (condition)
  ...
  a(end+1) = value; # "push" operation
  ...
  a(end) = []; # "pop" operation
  ...
endwhile
```

- Avoid calling `eval` or `feval` excessively. Parsing input or looking up the name of a function in the symbol table are relatively expensive operations.

  If you are using `eval` merely as an exception handling mechanism, and not because you need to execute some arbitrary text, use the `try` statement instead. See Section 10.9 [The try Statement], page 202.

- Use `ignore_function_time_stamp` when appropriate. If you are calling lots of functions, and none of them will need to change during your run, set the variable `ignore_function_time_stamp` to `"all"`. This will stop Octave from checking the time stamp of a function file to see if it has been updated while the program is being run.

## 19.7 Examples

The following are examples of vectorization questions asked by actual users of Octave and their solutions.

- For a vector `A`, the following loop

  ```
  n = length (A) - 1;
  B = zeros (n, 2);
  for i = 1:n
    ## this will be two columns, the first is the difference and
    ## the second the mean of the two elements used for the diff.
    B(i,:) = [A(i+1)-A(i), (A(i+1) + A(i))/2];
  endfor
  ```

  can be turned into the following one-liner:

  ```
  B = [diff(A)(:), 0.5*(A(1:end-1)+A(2:end))(:)]
  ```

  Note the usage of colon indexing to flatten an intermediate result into a column vector. This is a common vectorization trick.

# 20 Nonlinear Equations

## 20.1 Solvers

Octave can solve sets of nonlinear equations of the form

$$f(x) = 0$$

using the function `fsolve`, which is based on the MINPACK subroutine `hybrd`. This is an iterative technique so a starting point must be provided. This also has the consequence that convergence is not guaranteed even if a solution exists.

`x = fsolve (fcn, x0)`
`x = fsolve (fcn, x0, options)`
`[x, fval] = fsolve (...)`
`[x, fval, info] = fsolve (...)`
`[x, fval, info, output] = fsolve (...)`
`[x, fval, info, output, fjac] = fsolve (...)`

Solve a system of nonlinear equations defined by the function *fcn*.

*fcn* is a function handle, inline function, or string containing the name of the function to evaluate. *fcn* should accept a vector (array) defining the unknown variables, and return a vector of left-hand sides of the equations. Right-hand sides are defined to be zeros. In other words, this function attempts to determine a vector $x$ such that `fcn (x)` gives (approximately) all zeros.

*x0* is an initial guess for the solution. The shape of *x0* is preserved in all calls to *fcn*, but otherwise is treated as a column vector.

*options* is a structure specifying additional parameters which control the algorithm. Currently, `fsolve` recognizes these options: `"AutoScaling"`, `"ComplexEqn"`, `"FinDiffType"`, `"FunValCheck"`, `"Jacobian"`, `"MaxFunEvals"`, `"MaxIter"`, `"OutputFcn"`, `"TolFun"`, `"TolX"`, `"TypicalX"`, and `"Updating"`.

If `"AutoScaling"` is `"on"`, the variables will be automatically scaled according to the column norms of the (estimated) Jacobian. As a result, `"TolFun"` becomes scaling-independent. By default, this option is `"off"` because it may sometimes deliver unexpected (though mathematically correct) results.

If `"ComplexEqn"` is `"on"`, `fsolve` will attempt to solve complex equations in complex variables, assuming that the equations possess a complex derivative (i.e., are holomorphic). If this is not what you want, you should unpack the real and imaginary parts of the system to get a real system.

If `"Jacobian"` is `"on"`, it specifies that *fcn*—when called with 2 output arguments—also returns the Jacobian matrix of right-hand sides at the requested point.

`"MaxFunEvals"` proscribes the maximum number of function evaluations before optimization is halted. The default value is `100 * number_of_variables`, i.e., `100 * length (x0)`. The value must be a positive integer.

If `"Updating"` is `"on"`, the function will attempt to use Broyden updates to update the Jacobian, in order to reduce the number of Jacobian calculations. If your user

function always calculates the Jacobian (regardless of number of output arguments) then this option provides no advantage and should be disabled.

`"TolX"` specifies the termination tolerance in the unknown variables, while `"TolFun"` is a tolerance for equations. Default is `1e-6` for both `"TolX"` and `"TolFun"`.

For a description of the other options, see [optimset], page 825. To initialize an options structure with default values for `fsolve` use `options = optimset ("fsolve")`.

The first output *x* is the solution while the second output *fval* contains the value of the function *fcn* evaluated at *x* (ideally a vector of all zeros).

The third output *info* reports whether the algorithm succeeded and may take one of the following values:

1 Converged to a solution point. Relative residual error is less than specified by `TolFun`.

2 Last relative step size was less than `TolX`.

3 Last relative decrease in residual was less than `TolFun`.

0 Iteration limit (either `MaxIter` or `MaxFunEvals`) exceeded.

-1 Stopped by `OutputFcn`.

-2 The Jacobian became excessively small and the search stalled.

-3 The trust region radius became excessively small.

*output* is a structure containing runtime information about the `fsolve` algorithm. Fields in the structure are:

`iterations`
: Number of iterations through loop.

`successful`
: Number of successful iterations.

`funcCount`
: Number of function evaluations.

The final output *fjac* contains the value of the Jacobian evaluated at *x*.

Note: If you only have a single nonlinear equation of one variable, using `fzero` is usually a much better idea.

Note about user-supplied Jacobians: As an inherent property of the algorithm, a Jacobian is always requested for a solution vector whose residual vector is already known, and it is the last accepted successful step. Often this will be one of the last two calls, but not always. If the savings by reusing intermediate results from residual calculation in Jacobian calculation are significant, the best strategy is to employ `OutputFcn`: After a vector is evaluated for residuals, if `OutputFcn` is called with that vector, then the intermediate results should be saved for future Jacobian evaluation, and should be kept until a Jacobian evaluation is requested or until `OutputFcn` is called with a different vector, in which case they should be dropped in favor of this most recent vector. A short example how this can be achieved follows:

```
function [fval, fjac] = user_fcn (x, optimvalues, state)
```

```
persistent sav = [], sav0 = [];
if (nargin == 1)
  ## evaluation call
  if (nargout == 1)
    sav0.x = x;  # mark saved vector
    ## calculate fval, save results to sav0.
  elseif (nargout == 2)
    ## calculate fjac using sav.
  endif
else
  ## outputfcn call.
  if (all (x == sav0.x))
    sav = sav0;
  endif
  ## maybe output iteration status, etc.
endif
endfunction

## ...

fsolve (@user_fcn, x0, optimset ("OutputFcn", @user_fcn, ...))
```

**See also:** [fzero], page 710, [optimset], page 825.

The following is a complete example. To solve the set of equations

$$-2x^2 + 3xy + 4\sin(y) - 6 = 0$$
$$3x^2 - 2xy^2 + 3\cos(x) + 4 = 0$$

you first need to write a function to compute the value of the given function. For example:

```
function y = f (x)
  y = zeros (2, 1);
  y(1) = -2*x(1)^2 + 3*x(1)*x(2)   + 4*sin(x(2)) - 6;
  y(2) =  3*x(1)^2 - 2*x(1)*x(2)^2 + 3*cos(x(1)) + 4;
endfunction
```

Then, call `fsolve` with a specified initial condition to find the roots of the system of equations. For example, given the function `f` defined above,

```
[x, fval, info] = fsolve (@f, [1; 2])
```

results in the solution

```
x =

  0.57983
  2.54621

fval =

  -5.7184e-10
   5.5460e-10

info = 1
```

A value of `info = 1` indicates that the solution has converged.

When no Jacobian is supplied (as in the example above) it is approximated numerically. This requires more function evaluations, and hence is less efficient. In the example above we could compute the Jacobian analytically as

$$\begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \frac{\partial f_1}{\partial x_2} \\ \frac{\partial f_2}{\partial x_1} & \frac{\partial f_2}{\partial x_2} \end{bmatrix} = \begin{bmatrix} 3x_2 - 4x_1 & 4\cos(x_2) + 3x_1 \\ -2x_2^2 - 3\sin(x_1) + 6x_1 & -4x_1x_2 \end{bmatrix}$$

and compute it with the following Octave function

```
function [y, jac] = f (x)
  y = zeros (2, 1);
  y(1) = -2*x(1)^2 + 3*x(1)*x(2)   + 4*sin(x(2)) - 6;
  y(2) =  3*x(1)^2 - 2*x(1)*x(2)^2 + 3*cos(x(1)) + 4;
  if (nargout == 2)
    jac = zeros (2, 2);
    jac(1,1) =  3*x(2) - 4*x(1);
    jac(1,2) =  4*cos(x(2)) + 3*x(1);
    jac(2,1) = -2*x(2)^2 - 3*sin(x(1)) + 6*x(1);
    jac(2,2) = -4*x(1)*x(2);
  endif
endfunction
```

The Jacobian can then be used with the following call to `fsolve`:

```
[x, fval, info] = fsolve (@f, [1; 2], optimset ("jacobian", "on"));
```

which gives the same solution as before.

`x = fzero (`*`fcn`*`, `*`x0`*`)`
`x = fzero (`*`fcn`*`, `*`x0`*`, `*`options`*`)`
`[`*`x`*`, `*`fval`*`] = fzero (...)`
`[`*`x`*`, `*`fval`*`, `*`info`*`] = fzero (...)`
`[`*`x`*`, `*`fval`*`, `*`info`*`, `*`output`*`] = fzero (...)`

Find a zero of a univariate function.

*fcn* is a function handle, inline function, or string containing the name of the function to evaluate.

*x0* should be a two-element vector specifying two points which bracket a zero. In other words, there must be a change in sign of the function between *x0*(1) and *x0*(2). More mathematically, the following must hold

```
sign (fcn (x0(1))) * sign (fcn (x0(2))) <= 0
```

If *x0* is a single scalar then several nearby and distant values are probed in an attempt to obtain a valid bracketing. If this is not successful, the function fails.

*options* is a structure specifying additional options. Currently, `fzero` recognizes these options: `"Display"`, `"FunValCheck"`, `"MaxFunEvals"`, `"MaxIter"`, `"OutputFcn"`, and `"TolX"`.

`"MaxFunEvals"` proscribes the maximum number of function evaluations before the search is halted. The default value is `Inf`. The value must be a positive integer.

`"MaxIter"` proscribes the maximum number of algorithm iterations before the search is halted. The default value is `Inf`. The value must be a positive integer.

`"TolX"` specifies the termination tolerance for the solution *x*. The default value is `eps`.

For a description of the other options, see [optimset], page 825. To initialize an options structure with default values for `fzero` use `options = optimset ("fzero")`.

On exit, the function returns *x*, the approximate zero point, and *fval*, the function evaluated at *x*.

The third output *info* reports whether the algorithm succeeded and may take one of the following values:

- 1 The algorithm converged to a solution.
- 0 Maximum number of iterations or function evaluations has been reached.
- -1 The algorithm has been terminated by a user `OutputFcn`.
- -5 The algorithm may have converged to a singular point.
- -6 No valid initial bracketing with a sign change found.

*output* is a structure containing runtime information about the `fzero` algorithm. Fields in the structure are:

- intervaliterations Number of iterations searching for a bracketing interval.
- iterations Number of iterations searching for a zero.
- funcCount Number of function evaluations.
- algorithm The string `"bisection, interpolation"`.
- message A string with the final status of the algorithm.
- bracketx A two-element vector with the final bracketing of the zero along the x-axis.
- brackety A two-element vector with the final bracketing of the zero along the y-axis.

Programming Notes: The algorithm relies on the function crossing the x-axis and having both positive and negative values. It will not generally work for functions which only touch the x-axis, e.g., $x^2$. The function will generally be faster if a 2-element bracketing interval *x0* is supplied since no interval search will be conducted.

**See also:** [fsolve], page 707, [fminbnd], page 712, [optimset], page 825.

## 20.2 Minimizers

Often it is useful to find the minimum value of a function rather than just the zeroes where it crosses the x-axis. `fminbnd` is designed for the simpler, but very common, case of a univariate function where the interval to search is bounded. For unbounded minimization of a function with potentially many variables use `fminunc` or `fminsearch`. The two functions use different internal algorithms and some knowledge of the objective function is required. For functions which can be differentiated, `fminunc` is appropriate. For functions with discontinuities, or for which a gradient search would fail, use `fminsearch`. See Chapter 25 [Optimization], page 813, for minimization with the presence of constraint functions. Note that searches can be made for maxima by simply inverting the objective function ($F_{max} = -F_{min}$).

`x = fminbnd (fcn, a, b)`
`x = fminbnd (fcn, a, b, options)`
`[x, fval, info, output] = fminbnd (...)`

Find a minimum point of a univariate function.

*fcn* is a function handle, inline function, or string containing the name of the function to evaluate.

The starting interval is specified by *a* (left boundary) and *b* (right boundary). The endpoints must be finite.

*options* is a structure specifying additional parameters which control the algorithm. Currently, `fminbnd` recognizes these options: `"Display"`, `"FunValCheck"`, `"MaxFunEvals"`, `"MaxIter"`, `"OutputFcn"`, `"TolX"`.

`"MaxFunEvals"` proscribes the maximum number of function evaluations before optimization is halted. The default value is 500. The value must be a positive integer.

`"MaxIter"` proscribes the maximum number of algorithm iterations before optimization is halted. The default value is 500. The value must be a positive integer.

`"TolX"` specifies the termination tolerance for the solution *x*. The default is `1e-4`.

For a description of the other options, see [optimset], page 825. To initialize an options structure with default values for `fminbnd` use `options = optimset ("fminbnd")`.

On exit, the function returns *x*, the approximate minimum point, and *fval*, the function evaluated *x*.

The third output *info* reports whether the algorithm succeeded and may take one of the following values:

- 1 The algorithm converged to a solution.
- 0 Iteration limit (either `MaxIter` or `MaxFunEvals`) exceeded.
- -1 The algorithm was terminated by a user `OutputFcn`.

Application Notes:

1. The search for a minimum is restricted to be in the finite interval bound by *a* and *b*. If you have only one initial point to begin searching from then you will need to use an unconstrained minimization algorithm such as `fminunc` or `fminsearch`. `fminbnd` internally uses a Golden Section search strategy.

2. Use Section 11.12.2 [Anonymous Functions], page 250, to pass additional parameters to *fcn*. For specific examples of doing so for `fminbnd` and other minimization functions see the Section 20.2 [Minimizers], page 712, section of the GNU Octave manual.

**See also:** [fzero], page 710, [fminunc], page 713, [fminsearch], page 714, [optimset], page 825.

`x = fminunc (fcn, x0)`
`x = fminunc (fcn, x0, options)`
`[x, fval] = fminunc (fcn, ...)`
`[x, fval, info] = fminunc (fcn, ...)`
`[x, fval, info, output] = fminunc (fcn, ...)`
`[x, fval, info, output, grad] = fminunc (fcn, ...)`
`[x, fval, info, output, grad, hess] = fminunc (fcn, ...)`

Solve an unconstrained optimization problem defined by the function *fcn*.

`fminunc` attempts to determine a vector *x* such that `fcn (x)` is a local minimum.

*fcn* is a function handle, inline function, or string containing the name of the function to evaluate. *fcn* should accept a vector (array) defining the unknown variables, and return the objective function value, optionally with gradient.

*x0* determines a starting guess. The shape of *x0* is preserved in all calls to *fcn*, but otherwise is treated as a column vector.

*options* is a structure specifying additional parameters which control the algorithm. Currently, `fminunc` recognizes these options: `"AutoScaling"`, `"FinDiffType"`, `"FunValCheck"`, `"GradObj"`, `"MaxFunEvals"`, `"MaxIter"`, `"OutputFcn"`, `"TolFun"`, `"TolX"`, `"TypicalX"`.

If `"AutoScaling"` is `"on"`, the variables will be automatically scaled according to the column norms of the (estimated) Jacobian. As a result, `"TolFun"` becomes scaling-independent. By default, this option is `"off"` because it may sometimes deliver unexpected (though mathematically correct) results.

If `"GradObj"` is `"on"`, it specifies that *fcn*—when called with two output arguments—also returns the Jacobian matrix of partial first derivatives at the requested point.

`"MaxFunEvals"` proscribes the maximum number of function evaluations before optimization is halted. The default value is `100 * number_of_variables`, i.e., `100 * length (x0)`. The value must be a positive integer.

`"MaxIter"` proscribes the maximum number of algorithm iterations before optimization is halted. The default value is 400. The value must be a positive integer.

`"TolX"` specifies the termination tolerance for the unknown variables *x*, while `"TolFun"` is a tolerance for the objective function value *fval*. The default is `1e-6` for both options.

For a description of the other options, see [optimset], page 825.

On return, *x* is the location of the minimum and *fval* contains the value of the objective function at *x*.

*info* may be one of the following values:

| | |
|---|---|
| 1 | Converged to a solution point. Relative gradient error is less than specified by `TolFun`. |
| 2 | Last relative step size was less than `TolX`. |
| 3 | Last relative change in function value was less than `TolFun`. |
| 0 | Iteration limit exceeded—either maximum number of algorithm iterations `MaxIter` or maximum number of function evaluations `MaxFunEvals`. |
| -1 | Algorithm terminated by `OutputFcn`. |
| -3 | The trust region radius became excessively small. |

Optionally, `fminunc` can return a structure with convergence statistics (*output*), the output gradient (*grad*) at the solution *x*, and approximate Hessian (*hess*) at the solution *x*.

Application Notes:

1. If the objective function is a single nonlinear equation of one variable then using `fminbnd` is usually a better choice.
2. The algorithm used by `fminunc` is a gradient search which depends on the objective function being differentiable. If the function has discontinuities it may be better to use a derivative-free algorithm such as `fminsearch`.
3. Use Section 11.12.2 [Anonymous Functions], page 250, to pass additional parameters to *fcn*. For specific examples of doing so for `fminunc` and other minimization functions see the Section 20.2 [Minimizers], page 712, section of the GNU Octave manual.

**See also:** [fminbnd], page 712, [fminsearch], page 714, [optimset], page 825.

`x = fminsearch (fcn, x0)`
`x = fminsearch (fcn, x0, options)`
`x = fminsearch (problem)`
`[x, fval, exitflag, output] = fminsearch (...)`

Find a value of *x* which minimizes the multi-variable function *fcn*.

*fcn* is a function handle, inline function, or string containing the name of the function to evaluate.

The search begins at the point *x0* and iterates using the Nelder & Mead Simplex algorithm (a derivative-free method). This algorithm is better-suited to functions which have discontinuities or for which a gradient-based search such as `fminunc` fails.

Options for the search are provided in the parameter *options* using the function `optimset`. Currently, `fminsearch` accepts the options: `"Display"`, `"FunValCheck"`,`"MaxFunEvals"`, `"MaxIter"`, `"OutputFcn"`, `"TolFun"`, `"TolX"`.

`"MaxFunEvals"` proscribes the maximum number of function evaluations before optimization is halted. The default value is `200 * number_of_variables`, i.e., `200 * length (x0)`. The value must be a positive integer.

`"MaxIter"` proscribes the maximum number of algorithm iterations before optimization is halted. The default value is `200 * number_of_variables`, i.e., `200 * length (x0)`. The value must be a positive integer.

For a description of the other options, see [optimset], page 825. To initialize an options structure with default values for `fminsearch` use `options = optimset ("fminsearch")`.

`fminsearch` may also be called with a single structure argument with the following fields:

`objective`
: The objective function.

`x0`
: The initial point.

`solver`
: Must be set to `"fminsearch"`.

`options`
: A structure returned from `optimset` or an empty matrix to indicate that defaults should be used.

The field `options` is optional. All others are required.

On exit, the function returns *x*, the minimum point, and *fval*, the function value at the minimum.

The third output *exitflag* reports whether the algorithm succeeded and may take one of the following values:

1
: if the algorithm converged (size of the simplex is smaller than `TolX` **AND** the step in function value between iterations is smaller than `TolFun`).

0
: if the maximum number of iterations or the maximum number of function evaluations are exceeded.

-1
: if the iteration is stopped by the `"OutputFcn"`.

The fourth output is a structure *output* containing runtime about the algorithm. Fields in the structure are `funcCount` containing the number of function calls to *fcn*, `iterations` containing the number of iteration steps, `algorithm` with the name of the search algorithm (always: `"Nelder-Mead simplex direct search"`), and `message` with the exit message.

Example:

```
fminsearch (@(x) (x(1)-5).^2+(x(2)-8).^4, [0;0])
```

Application Notes:

1. If you need to find the minimum of a single variable function it is probably better to use `fminbnd`.
2. The legacy, undocumented syntax for passing parameters to *fcn* by appending them to the input argument list after *options* has been removed in Octave 10. The cross-platform compatible method of passing parameters to *fcn* is through the use of Section 11.12.2 [Anonymous Functions], page 250. For specific examples of doing so for `fminsearch` and other minimization functions see the Section 20.2 [Minimizers], page 712, section of the GNU Octave manual.

**See also:** [fminbnd], page 712, [fminunc], page 713, [optimset], page 825.

Certain minimization operations require additional parameters be passed to the function, `F`, being minimized. E.g., `F = F(x, C)`. Octave's minimizer functions are designed to only pass one optimization variable to `F`, but parameter passing can be accomplished by defining an Section 11.12.2 [Anonymous Function], page 250, that contains the necessary parameter(s). See the example below:

```
A = 2; B = 3;
f = @(x) sin (A*x + B);
fminbnd (f, 0, 2)
   ⇒    0.8562
```

Note that anonymous functions retain the value of parameters at the time they are defined. Changing a parameter value after function definition will not affect output of the function until it is redefined:

```
B = 4;
fminbnd (f, 0, 2)
   ⇒    0.8562
f = @(x) sin (A*x + B);
fminbnd (f, 0, 2)
   ⇒    0.3562
```

The function `humps` is a useful function for testing zero and extrema finding functions.

`y = humps (x)`
`[x, y] = humps (x)`

Evaluate a function with multiple minima, maxima, and zero crossings.

The output $y$ is the evaluation of the rational function:

$$y = -\frac{1200x^4 - 2880x^3 + 2036x^2 - 348x - 88}{200x^4 - 480x^3 + 406x^2 - 138x + 17}$$

$x$ may be a scalar, vector or array. If $x$ is omitted, the default range [0:0.05:1] is used.

When called with two output arguments, [$x$, $y$], $x$ will contain the input values, and $y$ will contain the output from `humps`.

Programming Notes: `humps` has two local maxima located near $x = 0.300$ and 0.893, a local minimum near $x = 0.637$, and zeros near $x =$ -0.132 and 1.300. `humps` is a useful function for testing algorithms which find zeros or local minima and maxima.

Try `demo humps` to see a plot of the `humps` function.

**See also:** [fzero], page 710, [fminbnd], page 712, [fminunc], page 713, [fminsearch], page 714.

# 21 Diagonal and Permutation Matrices

## 21.1 Creating and Manipulating Diagonal/Permutation Matrices

A diagonal matrix is defined as a matrix that has zero entries outside the main diagonal; that is, $D_{ij} = 0$ if $i \neq j$ Most often, square diagonal matrices are considered; however, the definition can equally be applied to non-square matrices, in which case we usually speak of a rectangular diagonal matrix.

A permutation matrix is defined as a square matrix that has a single element equal to unity in each row and each column; all other elements are zero. That is, there exists a permutation (vector) $p$ such that $P_{ij} = 1$ if $j = p_i$ and $P_{ij} = 0$ otherwise.

Octave provides special treatment of real and complex rectangular diagonal matrices, as well as permutation matrices. They are stored as special objects, using efficient storage and algorithms, facilitating writing both readable and efficient matrix algebra expressions in the Octave language. The special treatment may be disabled by using the functions *optimize_diagonal_matrix* and *optimize_permutation_matrix*.

```
val = optimize_diagonal_matrix ()
old_val = optimize_diagonal_matrix (new_val)
old_val = optimize_diagonal_matrix (new_val, "local")
```

Query or set whether a special space-efficient format is used for storing diagonal matrices.

The default value is true. If this option is set to false, Octave will store diagonal matrices as full matrices.

When called from inside a function with the `"local"` option, the setting is changed locally for the function and any subroutines it calls. The original setting is restored when exiting the function.

**See also:** [optimize_range], page 57, [optimize_permutation_matrix], page 717.

```
val = optimize_permutation_matrix ()
old_val = optimize_permutation_matrix (new_val)
old_val = optimize_permutation_matrix (new_val, "local")
```

Query or set whether a special space-efficient format is used for storing permutation matrices.

The default value is true. If this option is set to false, Octave will store permutation matrices as full matrices.

When called from inside a function with the `"local"` option, the setting is changed locally for the function and any subroutines it calls. The original setting is restored when exiting the function.

**See also:** [optimize_range], page 57, [optimize_diagonal_matrix], page 717.

The space savings are significant as demonstrated by the following code.

```
x = diag (rand (10, 1));
xf = full (x);
sizeof (x)
⇒ 80
sizeof (xf)
⇒ 800
```

### 21.1.1 Creating Diagonal Matrices

The most common and easiest way to create a diagonal matrix is using the built-in function *diag*. The expression `diag (v)`, with *v* a vector, will create a square diagonal matrix with elements on the main diagonal given by the elements of *v*, and size equal to the length of *v*. `diag (v, m, n)` can be used to construct a rectangular diagonal matrix. The result of these expressions will be a special diagonal matrix object, rather than a general matrix object.

Diagonal matrix with unit elements can be created using *eye*. Some other built-in functions can also return diagonal matrices. Examples include *balance* or *inv*.

Example:

```
  diag (1:4)
⇒
Diagonal Matrix

   1   0   0   0
   0   2   0   0
   0   0   3   0
   0   0   0   4

  diag (1:3,5,3)

⇒
Diagonal Matrix

   1   0   0
   0   2   0
   0   0   3
   0   0   0
   0   0   0
```

### 21.1.2 Creating Permutation Matrices

For creating permutation matrices, Octave does not introduce a new function, but rather overrides an existing syntax: permutation matrices can be conveniently created by indexing an identity matrix by permutation vectors. That is, if $q$ is a permutation vector of length $n$, the expression

```
P = eye (n) (:, q);
```

will create a permutation matrix - a special matrix object.

```
eye (n) (q, :)
```

will also work (and create a row permutation matrix), as well as

```
eye (n) (q1, q2).
```

For example:

```
  eye (4) ([1,3,2,4],:)
⇒
Permutation Matrix

   1   0   0   0
   0   0   1   0
   0   1   0   0
   0   0   0   1

  eye (4) (:,[1,3,2,4])
⇒
Permutation Matrix

   1   0   0   0
   0   0   1   0
   0   1   0   0
   0   0   0   1
```

Mathematically, an identity matrix is both diagonal and permutation matrix. In Octave, `eye (n)` returns a diagonal matrix, because a matrix can only have one class. You can convert this diagonal matrix to a permutation matrix by indexing it by an identity permutation, as shown below. This is a special property of the identity matrix; indexing other diagonal matrices generally produces a full matrix.

```
  eye (3)
⇒
Diagonal Matrix

   1   0   0
   0   1   0
   0   0   1

  eye(3)(1:3,:)
⇒
Permutation Matrix

   1   0   0
   0   1   0
   0   0   1
```

Some other built-in functions can also return permutation matrices. Examples include *inv* or *lu*.

### 21.1.3 Explicit and Implicit Conversions

The diagonal and permutation matrices are special objects in their own right. A number of operations and built-in functions are defined for these matrices to use special, more efficient

code than would be used for a full matrix in the same place. Examples are given in further sections.

To facilitate smooth mixing with full matrices, backward compatibility, and compatibility with MATLAB, the diagonal and permutation matrices should allow any operation that works on full matrices, and will either treat it specially, or implicitly convert themselves to full matrices.

Instances include matrix indexing, except for extracting a single element or a leading submatrix, indexed assignment, or applying most mapper functions, such as *exp*.

An explicit conversion to a full matrix can be requested using the built-in function *full*. It should also be noted that the diagonal and permutation matrix objects will cache the result of the conversion after it is first requested (explicitly or implicitly), so that subsequent conversions will be very cheap.

## 21.2 Linear Algebra with Diagonal/Permutation Matrices

As has been already said, diagonal and permutation matrices make it possible to use efficient algorithms while preserving natural linear algebra syntax. This section describes in detail the operations that are treated specially when performed on these special matrix objects.

### 21.2.1 Expressions Involving Diagonal Matrices

Assume $D$ is a diagonal matrix. If $M$ is a full matrix, then `D*M` will scale the rows of $M$. That means, if `S = D*M`, then for each pair of indices i,j it holds

$$S_{ij} = D_{ii} M_{ij}$$

Similarly, `M*D` will do a column scaling.

The matrix $D$ may also be rectangular, m-by-n where `m != n`. If `m < n`, then the expression `D*M` is equivalent to

```
D(:,1:m) * M(1:m,:),
```

i.e., trailing `n-m` rows of $M$ are ignored. If `m > n`, then `D*M` is equivalent to

```
[D(1:n,:) * M; zeros(m-n, columns (M))],
```

i.e., null rows are appended to the result. The situation for right-multiplication `M*D` is analogous.

The expressions `D \ M` and `M / D` perform inverse scaling. They are equivalent to solving a diagonal (or rectangular diagonal) in a least-squares minimum-norm sense. In exact arithmetic, this is equivalent to multiplying by a pseudoinverse. The pseudoinverse of a rectangular diagonal matrix is again a rectangular diagonal matrix with swapped dimensions, where each nonzero diagonal element is replaced by its reciprocal. The matrix division algorithms do, in fact, use division rather than multiplication by reciprocals for better numerical accuracy; otherwise, they honor the above definition. Note that a diagonal matrix is never truncated due to ill-conditioning; otherwise, it would not be of much use for scaling. This is typically consistent with linear algebra needs. A full matrix that only happens to be diagonal (and is thus not a special object) is of course treated normally.

Multiplication and division by diagonal matrices work efficiently also when combined with sparse matrices, i.e., `D*S`, where $D$ is a diagonal matrix and $S$ is a sparse matrix scales

the rows of the sparse matrix and returns a sparse matrix. The expressions `S*D`, `D\S`, `S/D` work analogically.

If *D1* and *D2* are both diagonal matrices, then the expressions

```
D1 + D2
D1 - D2
D1 * D2
D1 / D2
D1 \ D2
```

again produce diagonal matrices, provided that normal dimension matching rules are obeyed. The relations used are same as described above.

Also, a diagonal matrix $D$ can be multiplied or divided by a scalar, or raised to a scalar power if it is square, producing diagonal matrix result in all cases.

A diagonal matrix can also be transposed or conjugate-transposed, giving the expected result. Extracting a leading submatrix of a diagonal matrix, i.e., `D(1:m,1:n)`, will produce a diagonal matrix, other indexing expressions will implicitly convert to full matrix.

Adding a diagonal matrix to a full matrix only operates on the diagonal elements. Thus,

```
A = A + eps * eye (n)
```

is an efficient method of augmenting the diagonal of a matrix. Subtraction works analogically.

When involved in expressions with other element-by-element operators, `.*`, `./`, `.\` or `.^`, an implicit conversion to full matrix will take place. This is not always strictly necessary but chosen to facilitate better consistency with MATLAB.

### 21.2.2 Expressions Involving Permutation Matrices

If $P$ is a permutation matrix and $M$ a matrix, the expression `P*M` will permute the rows of $M$. Similarly, `M*P` will yield a column permutation. Matrix division `P\M` and `M/P` can be used to do inverse permutation.

The previously described syntax for creating permutation matrices can actually help an user to understand the connection between a permutation matrix and a permuting vector. Namely, the following holds, where `I = eye (n)` is an identity matrix:

```
I(p,:) * M = (I*M) (p,:) = M(p,:)
```

Similarly,

```
M * I(:,p) = (M*I) (:,p) = M(:,p)
```

The expressions `I(p,:)` and `I(:,p)` are permutation matrices.

A permutation matrix can be transposed (or conjugate-transposed, which is the same, because a permutation matrix is never complex), inverting the permutation, or equivalently, turning a row-permutation matrix into a column-permutation one. For permutation matrices, transpose is equivalent to inversion, thus `P\M` is equivalent to `P'*M`. Transpose of a permutation matrix (or inverse) is a constant-time operation, flipping only a flag internally, and thus the choice between the two above equivalent expressions for inverse permuting is completely up to the user's taste.

Multiplication and division by permutation matrices works efficiently also when combined with sparse matrices, i.e., `P*S`, where $P$ is a permutation matrix and $S$ is a sparse

matrix permutes the rows of the sparse matrix and returns a sparse matrix. The expressions `S*P`, `P\S`, `S/P` work analogically.

Two permutation matrices can be multiplied or divided (if their sizes match), performing a composition of permutations. Also a permutation matrix can be indexed by a permutation vector (or two vectors), giving again a permutation matrix. Any other operations do not generally yield a permutation matrix and will thus trigger the implicit conversion.

## 21.3 Functions That Are Aware of These Matrices

This section lists the built-in functions that are aware of diagonal and permutation matrices on input, or can return them as output. Passed to other functions, these matrices will in general trigger an implicit conversion. (Of course, user-defined dynamically linked functions may also work with diagonal or permutation matrices).

### 21.3.1 Diagonal Matrix Functions

*inv* and *pinv* can be applied to a diagonal matrix, yielding again a diagonal matrix. *det* will use an efficient straightforward calculation when given a diagonal matrix, as well as *cond*. The following mapper functions can be applied to a diagonal matrix without converting it to a full one: *abs*, *real*, *imag*, *conj*, *sqrt*. A diagonal matrix can also be returned from the *balance* and *svd* functions. The *sparse* function will convert a diagonal matrix efficiently to a sparse matrix.

### 21.3.2 Permutation Matrix Functions

*inv* and *pinv* will invert a permutation matrix, preserving its specialness. *det* can be applied to a permutation matrix, efficiently calculating the sign of the permutation (which is equal to the determinant).

A permutation matrix can also be returned from the built-in functions *lu* and *qr*, if a pivoted factorization is requested.

The *sparse* function will convert a permutation matrix efficiently to a sparse matrix. The *find* function will also work efficiently with a permutation matrix, making it possible to conveniently obtain the permutation indices.

## 21.4 Examples of Usage

The following can be used to solve a linear system `A*x = b` using the pivoted LU factorization:

```
[L, U, P] = lu (A); ## now L*U = P*A
x = U \ (L \ P) * b;
```

This is one way to normalize columns of a matrix $X$ to unit norm:

```
s = norm (X, "columns");
X /= diag (s);
```

The same can also be accomplished with broadcasting (see Section 19.2 [Broadcasting], page 691):

```
s = norm (X, "columns");
X ./= s;
```

The following expression is a way to efficiently calculate the sign of a permutation, given by a permutation vector *p*. It will also work in earlier versions of Octave, but slowly.

```
det (eye (length (p))(p, :))
```

Finally, here's how to solve a linear system `A*x = b` with Tikhonov regularization (ridge regression) using SVD (a skeleton only):

```
m = rows (A); n = columns (A);
[U, S, V] = svd (A);
## determine the regularization factor alpha
## alpha = ...
## transform to orthogonal basis
b = U'*b;
## Use the standard formula, replacing A with S.
## S is diagonal, so the following will be very fast and accurate.
x = (S'*S + alpha^2 * eye (n)) \ (S' * b);
## transform to solution basis
x = V*x;
```

## 21.5 Differences in Treatment of Zero Elements

Making diagonal and permutation matrices special matrix objects in their own right and the consequent usage of smarter algorithms for certain operations implies, as a side effect, small differences in treating zeros. The contents of this section apply also to sparse matrices, discussed in the following chapter. (see Chapter 22 [Sparse Matrices], page 727)

The IEEE 754 floating point standard defines the result of the expressions `0*Inf` and `0*NaN` as `NaN`. This is widely agreed to be a good compromise and Octave strives to conform to this standard. However, certain algorithms designed for maximized efficiency with structured and sparse matrices may not conform to the IEEE 754 floating point standard regarding operations with zero elements that should result in `NaN` such as `0*Inf`, `0/0`, or `0*NaN`.

Numerical software dealing with structured and sparse matrices (including Octave) almost always makes a distinction between a "numerical zero" and an "assumed zero". A "numerical zero" is a zero value occurring in a place where any floating-point value could occur. It is normally stored somewhere in memory as an explicit value. An "assumed zero", on the contrary, is a zero matrix element implied by the matrix structure (diagonal, triangular) or a sparsity pattern; its value is usually not stored explicitly anywhere, but is implied by the underlying data structure.

The primary distinction is that an assumed zero, when multiplied or divided by any number, yields **always** a zero, even when, e.g., multiplied by `Inf` or divided by `NaN`. The reason for this behavior is that the numerical multiplication is not actually performed anywhere by the underlying algorithm; the result is just assumed to be zero.

Octave mitigates this behavior by branching the underlying algorithm into specialized implementations to handle operations that actually affect assumed zeros. This is primarily a trade-off between computational efficiency and conformity to the floating point standard by treating assumed zeros as numerical zeros. Besides the difference in computation speed, this behavior also affects sparsity patterns that should be taken into account when operating on very large sparse matrices which may result in out-of-memory errors.

Examples of treating assumed zeros as numerical zeros in sparse matrices:

```
speye (3)
⇒
Compressed Column Sparse (rows = 3, cols = 3, nnz = 3 [33%])

  (1, 1) -> 1
  (2, 2) -> 1
  (3, 3) -> 1

Inf * speye (3)
⇒
Compressed Column Sparse (rows = 3, cols = 3, nnz = 9 [100%])

  (1, 1) -> Inf
  (2, 1) -> NaN
  (3, 1) -> NaN
  (1, 2) -> NaN
  (2, 2) -> Inf
  (3, 2) -> NaN
  (1, 3) -> NaN
  (2, 3) -> NaN
  (3, 3) -> Inf

Inf * full (eye (3))
⇒
   Inf   NaN   NaN
   NaN   Inf   NaN
   NaN   NaN   Inf

speye (3) / 0
⇒
Compressed Column Sparse (rows = 3, cols = 3, nnz = 9 [100%])

  (1, 1) -> Inf
  (2, 1) -> NaN
  (3, 1) -> NaN
  (1, 2) -> NaN
  (2, 2) -> Inf
  (3, 2) -> NaN
  (1, 3) -> NaN
  (2, 3) -> NaN
  (3, 3) -> Inf

full (eye (3)) / 0
⇒
   Inf   NaN   NaN
```

```
   NaN   Inf   NaN
   NaN   NaN   Inf
```

Although in sparse matrices only the sparsity pattern may be affected by certain operations, when dealing with diagonal matrices, such operations may also affect the returned type. The following rules apply:

- finite scalar * diagonal matrix is a diagonal matrix
- nonfinite scalar * diagonal matrix is a matrix
- diagonal matrix / 0 is a matrix
- diagonal matrix / NaN is a matrix
- diagonal matrix / any other scalar is a diagonal matrix

For example:

```
3 * eye (3)
⇒

Diagonal Matrix

   3   0   0
   0   3   0
   0   0   3

Inf * eye (3)
⇒

      Inf      NaN      NaN
      NaN      Inf      NaN
      NaN      NaN      Inf

eye (3) / 0
⇒

      Inf      NaN      NaN
      NaN      Inf      NaN
      NaN      NaN      Inf

eye (3) / NaN
⇒

      NaN      NaN      NaN
      NaN      NaN      NaN
      NaN      NaN      NaN

eye (3) / Inf
⇒
```

```
Diagonal Matrix

   0   0   0
   0   0   0
   0   0   0

eye (3) / 2
⇒

Diagonal Matrix

   0.5000        0        0
        0   0.5000        0
        0        0   0.5000
```

Note that Octave (as well as MATLAB) does not strictly conform to the IEEE 754 floating point standard and there are certain cases like in the following examples where the underlying algorithms operate under assumed zeros without treating then as numerical zeros. As of Octave 11, the intention is to make Octave fully compliant with the IEEE 754 standard so that operations on dense, sparse, or any other special matrix types produce the same consistent results. Thus, it is expected that the output in the following examples may change in future Octave versions.

Examples of effects of assumed zeros vs. numerical zeros:

```
diag (1:3) * [NaN; 1; 1]
⇒
   NaN
     2
     3

sparse (1:3,1:3,1:3) * [NaN; 1; 1]
⇒
   NaN
     2
     3
[1,0,0;0,2,0;0,0,3] * [NaN; 1; 1]
⇒
   NaN
   NaN
   NaN
```

# 22 Sparse Matrices

## 22.1 Creation and Manipulation of Sparse Matrices

The size of mathematical problems that can be treated at any particular time is generally limited by the available computing resources. Both, the speed of the computer and its available memory place limitation on the problem size.

There are many classes of mathematical problems which give rise to matrices, where a large number of the elements are zero. In this case it makes sense to have a special matrix type to handle this class of problems where only the nonzero elements of the matrix are stored. Not only does this reduce the amount of memory to store the matrix, but it also means that operations on this type of matrix can take advantage of the a priori knowledge of the positions of the nonzero elements to accelerate their calculations.

A matrix type that stores only the nonzero elements is generally called sparse. It is the purpose of this document to discuss the basics of the storage and creation of sparse matrices and the fundamental operations on them.

### 22.1.1 Storage of Sparse Matrices

It is not strictly speaking necessary for the user to understand how sparse matrices are stored. However, such an understanding will help to get an understanding of the size of sparse matrices. Understanding the storage technique is also necessary for those users wishing to create their own oct-files.

There are many different means of storing sparse matrix data. What all of the methods have in common is that they attempt to reduce the complexity and storage given a priori knowledge of the particular class of problems that will be solved. A good summary of the available techniques for storing sparse matrix is given by Saad[1]. With full matrices, knowledge of the point of an element of the matrix within the matrix is implied by its position in the computers memory. However, this is not the case for sparse matrices, and so the positions of the nonzero elements of the matrix must equally be stored.

An obvious way to do this is by storing the elements of the matrix as triplets, with two elements being their position in the array (rows and column) and the third being the data itself. This is conceptually easy to grasp, but requires more storage than is strictly needed.

The storage technique used within Octave is the compressed column format. It is similar to the Yale format.[2] In this format the position of each element in a row and the data are stored as previously. However, if we assume that all elements in the same column are stored adjacent in the computers memory, then we only need to store information on the number of nonzero elements in each column, rather than their positions. Thus assuming that the matrix has more nonzero elements than there are columns in the matrix, we win in terms of the amount of memory used.

In fact, the column index contains one more element than the number of columns, with the first element always being zero. The advantage of this is a simplification in the code, in

[1] Y. Saad "SPARSKIT: A basic toolkit for sparse matrix computation", 1994, https://www-users.cs.umn.edu/~saad/software/SPARSKIT/paper.ps

[2] https://en.wikipedia.org/wiki/Sparse_matrix#Yale_format

that there is no special case for the first or last columns. A short example, demonstrating this in C is.

```
for (j = 0; j < nc; j++)
  for (i = cidx(j); i < cidx(j+1); i++)
     printf ("nonzero element (%i,%i) is %d\n",
         ridx(i), j, data(i));
```

A clear understanding might be had by considering an example of how the above applies to an example matrix. Consider the matrix

```
1   2   0  0
0   0   0  3
0   0   0  4
```

The nonzero elements of this matrix are

```
(1, 1)  ⇒ 1
(1, 2)  ⇒ 2
(2, 4)  ⇒ 3
(3, 4)  ⇒ 4
```

This will be stored as three vectors *cidx*, *ridx* and *data*, representing the column indexing, row indexing and data respectively. The contents of these three vectors for the above matrix will be

```
cidx = [0, 1, 2, 2, 4]
ridx = [0, 0, 1, 2]
data = [1, 2, 3, 4]
```

Note that this is the representation of these elements with the first row and column assumed to start at zero, while in Octave itself the row and column indexing starts at one. Thus, the number of elements in the *i*-th column is given by `cidx (i + 1) - cidx (i)`.

Although Octave uses a compressed column format, it should be noted that compressed row formats are equally possible. However, in the context of mixed operations between mixed sparse and dense matrices, it makes sense that the elements of the sparse matrices are in the same order as the dense matrices. Octave stores dense matrices in column major ordering, and so sparse matrices are equally stored in this manner.

A further constraint on the sparse matrix storage used by Octave is that all elements in the rows are stored in increasing order of their row index, which makes certain operations faster. However, it imposes the need to sort the elements on the creation of sparse matrices. Having disordered elements is potentially an advantage in that it makes operations such as concatenating two sparse matrices together easier and faster, however it adds complexity and speed problems elsewhere.

### 22.1.2 Creating Sparse Matrices

There are several means to create sparse matrix.

Returned from a function

> There are many functions that directly return sparse matrices. These include *speye*, *sprand*, *diag*, etc.

Constructed from matrices or vectors
: The function *sparse* allows a sparse matrix to be constructed from three vectors representing the row, column and data. Alternatively, the function *spconvert* uses a three column matrix format to allow easy importation of data from elsewhere.

Created and then filled
: The function *sparse* or *spalloc* can be used to create an empty matrix that is then filled by the user

From a user binary program
: The user can directly create the sparse matrix within an oct-file.

There are several basic functions to return specific sparse matrices. For example the sparse identity matrix, is a matrix that is often needed. It therefore has its own function to create it as `speye (n)` or `speye (r, c)`, which creates an $n$-by-$n$ or $r$-by-$c$ sparse identity matrix.

Another typical sparse matrix that is often needed is a random distribution of random elements. The functions *sprand* and *sprandn* perform this for uniform and normal random distributions of elements. They have exactly the same calling convention, where `sprand (r, c, d)`, creates an $r$-by-$c$ sparse matrix with a density of filled elements of $d$.

Other functions of interest that directly create sparse matrices, are *diag* or its generalization *spdiags*, that can take the definition of the diagonals of the matrix and create the sparse matrix that corresponds to this. For example,

```
s = diag (sparse (randn (1,n)), -1);
```

creates a sparse ($n$+1)-by-($n$+1) sparse matrix with a single diagonal defined.

`B = spdiags (A)`
`[B, d] = spdiags (A)`
`B = spdiags (A, d)`
`A = spdiags (v, d, A)`
`A = spdiags (v, d, m, n)`

A generalization of the function `diag`.

Called with a single input argument, the nonzero diagonals $d$ of $A$ are extracted.

With two arguments the diagonals to extract are given by the vector $d$.

The other two forms of `spdiags` modify the input matrix by replacing the diagonals. They use the columns of $v$ to replace the diagonals represented by the vector $d$. If the sparse matrix $A$ is defined then the diagonals of this matrix are replaced. Otherwise a matrix of $m$ by $n$ is created with the diagonals given by the columns of $v$.

Negative values of $d$ represent diagonals below the main diagonal, and positive values of $d$ diagonals above the main diagonal.

For example:

```
spdiags (reshape (1:12, 4, 3), [-1 0 1], 5, 4)
   ⇒ 5 10  0  0
     1  6 11  0
     0  2  7 12
     0  0  3  8
     0  0  0  4
```

**See also:** [diag], page 579.

`s = speye ()`
`s = speye (n)`
`s = speye (m, n)`
`s = speye ([m, n])`

Return a sparse identity matrix of size *m*x*n*.

If called with no arguments, return the sparse scalar value `1`.

If invoked with a single scalar argument *n*, return a sparse square NxN identity matrix.

If supplied two scalar arguments (*m*, *n*), or a 2-element vector `[m, n]`, return a sparse MxN identity matrix with *m* rows and *n* columns.

Programming Note: The implementation is significantly more efficient than `sparse (eye (...))` as the full matrix is not constructed.

**See also:** [sparse], page 732, [spdiags], page 729, [eye], page 580.

`r = spones (S)`

Replace the nonzero entries of *S* with ones.

This creates a sparse matrix with the same structure as *S*.

**See also:** [sparse], page 732, [sprand], page 730, [sprandn], page 730, [sprandsym], page 731, [spfun], page 696, [spy], page 737.

`s = sprand (m, n, d)`
`s = sprand (m, n, d, rc)`
`s = sprand (s)`

Generate a sparse matrix with uniformly distributed random values.

The size of the matrix is *m*x*n* with a density of values *d*. *d* must be between 0 and 1. Values will be uniformly distributed on the interval (0, 1).

If called with a single matrix argument, a sparse matrix is generated with random values wherever the matrix *s* is nonzero.

If called with a scalar fourth argument *rc*, a random sparse matrix with reciprocal condition number *rc* is generated. If *rc* is a vector, then it specifies the first singular values of the generated matrix (`length (rc) <= min (m, n)`).

**See also:** [sprandn], page 730, [sprandsym], page 731, [rand], page 585.

`s = sprandn (m, n, d)`
`s = sprandn (m, n, d, rc)`
`s = sprandn (s)`

Generate a sparse matrix with normally distributed random values.

The size of the matrix is *m*x*n* with a density of values *d*. *d* must be between 0 and 1. Values will be normally distributed with a mean of 0 and a variance of 1.

If called with a single matrix argument, a sparse matrix is generated with random values wherever the matrix *s* is nonzero.

If called with a scalar fourth argument *rc*, a random sparse matrix with reciprocal condition number *rc* is generated. If *rc* is a vector, then it specifies the first singular values of the generated matrix (`length (rc) <= min (m, n)`).

**See also:** [sprand], page 730, [sprandsym], page 731, [randn], page 587.

*S* = sprandsym (*n*, *d*)
*S* = sprandsym (*s*)

Generate a symmetric random sparse matrix.

The size of the matrix will be *n*x*n*, with a density of values given by *d*. *d* must be between 0 and 1 inclusive. Values will be normally distributed with a mean of zero and a variance of 1.

If called with a single matrix argument, a random sparse matrix is generated wherever the matrix *s* is nonzero in its lower triangular part.

**See also:** [sprand], page 730, [sprandn], page 730, [spones], page 730, [sparse], page 732.

The recommended way for the user to create a sparse matrix, is to create two vectors containing the row and column index of the data and a third vector of the same size containing the data to be stored. For example,

```
ri = ci = d = [];
for j = 1:c
  ri = [ri; randperm(r,n)'];
  ci = [ci; j*ones(n,1)];
  d = [d; rand(n,1)];
endfor
s = sparse (ri, ci, d, r, c);
```

creates an *r*-by-*c* sparse matrix with a random distribution of *n* (<*r*) elements per column. The elements of the vectors do not need to be sorted in any particular order as Octave will sort them prior to storing the data. However, pre-sorting the data will make the creation of the sparse matrix faster.

The function *spconvert* takes a three or four column real matrix. The first two columns represent the row and column index respectively and the third and four columns, the real and imaginary parts of the sparse matrix. The matrix can contain zero elements and the elements can be sorted in any order. Adding zero elements is a convenient way to define the size of the sparse matrix. For example:

```
s = spconvert ([1 2 3 4; 1 3 4 4; 1 2 3 0]')
⇒ Compressed Column Sparse (rows=4, cols=4, nnz=3)
      (1 , 1) -> 1
      (2 , 3) -> 2
      (3 , 4) -> 3
```

An example of creating and filling a matrix might be

```
k = 5;
nz = r * k;
s = spalloc (r, c, nz)
for j = 1:c
  idx = randperm (r);
  s (:, j) = [zeros(r - k, 1); ...
        rand(k, 1)] (idx);
endfor
```

It should be noted, that due to the way that the Octave assignment functions are written that the assignment will reallocate the memory used by the sparse matrix at each iteration of the above loop. Therefore the *spalloc* function ignores the *nz* argument and does not pre-assign the memory for the matrix. Therefore, it is vitally important that code using to above structure should be vectorized as much as possible to minimize the number of assignments and reduce the number of memory allocations.

*FM* = `full` (*SM*)

Return a full storage matrix from a sparse, diagonal, or permutation matrix, or from a range.

**See also:** [sparse], page 732, [issparse], page 734.

*s* = `spalloc` (*m*, *n*, *nz*)

Create an *m*-by-*n* sparse matrix with pre-allocated space for at most *nz* nonzero elements.

This is useful for building a matrix incrementally by a sequence of indexed assignments. Subsequent indexed assignments after `spalloc` will reuse the pre-allocated memory, provided they are of one of the simple forms

- *s*`(I:J) =` *x*
- *s*`(:,I:J) =` *x*
- *s*`(K:L,I:J) =` *x*

**and** that the following conditions are met:

- the assignment does not decrease `nnz (`*S*`)`.
- after the assignment, `nnz (`*S*`)` does not exceed *nz*.
- no index is out of bounds.

Partial movement of data may still occur, but in general the assignment will be more memory and time efficient under these circumstances. In particular, it is possible to efficiently build a pre-allocated sparse matrix from a contiguous block of columns.

The amount of pre-allocated memory for a given matrix may be queried using the function `nzmax`.

Programming Note: Octave always reserves memory for at least one value, even if *nz* is 0.

**See also:** [nzmax], page 734, [sparse], page 732.

*S* = `sparse` (*A*)
*S* = `sparse` (*m*, *n*)
*S* = `sparse` (*i*, *j*, *sv*)
*S* = `sparse` (*i*, *j*, *sv*, *m*, *n*)
*S* = `sparse` (*i*, *j*, *sv*, *m*, *n*, "*unique*")
*S* = `sparse` (*i*, *j*, *sv*, *m*, *n*, *nzmax*)

Create a sparse matrix from a full matrix *A* or row, column, value triplets.

If *A* is a full matrix, convert it to a sparse matrix representation, removing all zero values in the process. The matrix *A* should be of type logical or double.

If two inputs *m* (rows) and *n* (columns) are specified then create an empty sparse matrix with the specified dimensions.

Given the integer index vectors *i* and *j*, and a 1-by-`nnz` vector of real or complex values *sv*, construct the sparse matrix `S(i(k),j(k)) = sv(k)` with overall dimensions *m* and *n*. If any of *i*, *j*, or *sv* are scalars, they are expanded to have a common size.

If *m* or *n* are not specified then their values are derived from the maximum index in the vectors *i* and *j* as given by `m = max (i)`, `n = max (j)`.

**Note**: If multiple values are specified with the same *i*, *j* indices, the corresponding value in *S* will be the sum of the values at the repeated location. See [accumarray], page 700, for an example of how to produce different behavior such as taking the minimum instead.

If the option `"unique"` is given, and more than one value is specified at the same *i*, *j* indices, then only the last specified value will be used. For completeness, the option `"sum"` can be given and will be ignored as the default behavior is to sum values at repeated locations.

`sparse (m, n)` will create an empty *m*x*n* sparse matrix and is equivalent to `sparse ([], [], [], m, n)`

The optional final argument reserves space for *nzmax* values in the sparse array and is useful if the eventual number of nonzero values will be greater than the number of values in *sv* used during the initial construction of the array. See [spalloc], page 732, for more information and usage instructions.

Example 1 (convert full matrix to sparse to save memory):

```
x = full (diag (1:1000));
sizeof (x)
⇒  8000000
s = sparse (x);
sizeof (xs)
⇒  24008
```

Example 2 (sum at repeated indices):

```
i = [1 1 2]; j = [1 1 2]; sv = [3 4 5];
sparse (i, j, sv, 3, 4)
⇒
   Compressed Column Sparse (rows = 3, cols = 4, nnz = 2 [17%])

     (1, 1) ->  7
     (2, 2) ->  5
```

Example 3 ("unique" option):

```
i = [1 1 2]; j = [1 1 2]; sv = [3 4 5];
sparse (i, j, sv, 3, 4, "unique")
⇒
   Compressed Column Sparse (rows = 3, cols = 4, nnz = 2 [17%])

     (1, 1) ->  4
     (2, 2) ->  5
```

**See also:** [full], page 732, [accumarray], page 700, [spalloc], page 732, [spdiags], page 729, [speye], page 730, [spones], page 730, [sprand], page 730, [sprandn], page 730, [sprandsym], page 731, [spconvert], page 734, [spfun], page 696.

`x = spconvert (m)`

Convert a simple sparse matrix format easily generated by other programs into Octave's internal sparse format.

The input *m* is either a 3 or 4 column real matrix, containing the row, column, real, and imaginary parts of the elements of the sparse matrix. An element with a zero real and imaginary part can be used to force a particular matrix size.

**See also:** [sparse], page 732.

The above problem of memory reallocation can be avoided in oct-files. However, the construction of a sparse matrix from an oct-file is more complex than can be discussed here. See Appendix A [External Code Interface], page 1097, for a full description of the techniques involved.

### 22.1.3 Finding Information about Sparse Matrices

There are a number of functions that allow information concerning sparse matrices to be obtained. The most basic of these is *issparse* that identifies whether a particular Octave object is in fact a sparse matrix.

Another very basic function is *nnz* that returns the number of nonzero entries there are in a sparse matrix, while the function *nzmax* returns the amount of storage allocated to the sparse matrix. Note that Octave tends to crop unused memory at the first opportunity for sparse objects. There are some cases of user created sparse objects where the value returned by *nzmax* will not be the same as *nnz*, but in general they will give the same result. The function *spstats* returns some basic statistics on the columns of a sparse matrix including the number of elements, the mean and the variance of each column.

`tf = issparse (x)`

Return true if *x* is a sparse matrix.

**See also:** [ismatrix], page 69.

`n = nnz (A)`

Return the number of nonzero elements in *A*.

**See also:** [nzmax], page 734, [nonzeros], page 734, [find], page 568.

`v = nonzeros (A)`

Return a column vector of the nonzero values of the matrix *A*.

**See also:** [find], page 568, [nnz], page 734.

`n = nzmax (SM)`

Return the amount of storage allocated to the sparse matrix *SM*.

Programming Note: Octave tends to crop unused memory at the first opportunity for sparse objects. Thus, in general the value of `nzmax` will be the same as `nnz`, except for some cases of user-created sparse objects.

Also, note that Octave always reserves storage for at least one value. Thus, for empty matrices `nnz` will report 0, but `nzmax` will report 1.

**See also:** [nnz], page 734, [spalloc], page 732, [sparse], page 732.

`[`*count*`, `*mean*`, `*var*`] = spstats (`*S*`)`
`[`*count*`, `*mean*`, `*var*`] = spstats (`*S*`, `*j*`)`
Return the stats for the nonzero elements of the sparse matrix *S*.

*count* is the number of nonzeros in each column, *mean* is the mean of the nonzeros in each column, and *var* is the variance of the nonzeros in each column.

Called with two input arguments, if *S* is the data and *j* is the bin number for the data, compute the stats for each bin. In this case, bins can contain data values of zero, whereas with `spstats (`*S*`)` the zeros may disappear.

When solving linear equations involving sparse matrices Octave determines the means to solve the equation based on the type of the matrix (see Section 22.2 [Sparse Linear Algebra], page 748). Octave probes the matrix type when the div (/) or ldiv (\) operator is first used with the matrix and then caches the type. However the *matrix_type* function can be used to determine the type of the sparse matrix prior to use of the div or ldiv operators. For example,

```
a = tril (sprandn (1024, 1024, 0.02), -1) ...
    + speye (1024);
matrix_type (a);
ans = Lower
```

shows that Octave correctly determines the matrix type for lower triangular matrices. *matrix_type* can also be used to force the type of a matrix to be a particular type. For example:

```
a = matrix_type (tril (sprandn (1024, ...
   1024, 0.02), -1) + speye (1024), "Lower");
```

This allows the cost of determining the matrix type to be avoided. However, incorrectly defining the matrix type will result in incorrect results from solutions of linear equations, and so it is entirely the responsibility of the user to correctly identify the matrix type

There are several graphical means of finding out information about sparse matrices. The first is the *spy* command, which displays the structure of the nonzero elements of the matrix. See Figure 22.1, for an example of the use of *spy*. More advanced graphical information can be obtained with the *treeplot*, *etreeplot* and *gplot* commands.

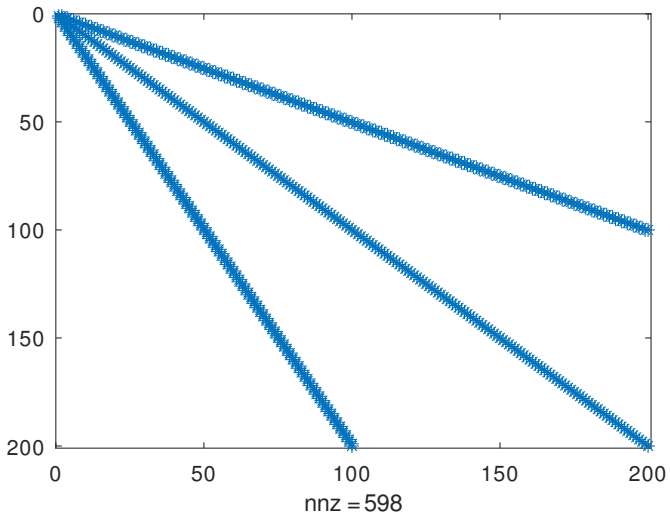



Figure 22.1: Structure of simple sparse matrix.

One use of sparse matrices is in graph theory, where the interconnections between nodes are represented as an adjacency matrix. That is, if the i-th node in a graph is connected to the j-th node. Then the ij-th node (and in the case of undirected graphs the ji-th node) of the sparse adjacency matrix is nonzero. If each node is then associated with a set of coordinates, then the *gplot* command can be used to graphically display the interconnections between nodes.

As a trivial example of the use of *gplot* consider the example,

```
A = sparse ([2,6,1,3,2,4,3,5,4,6,1,5],
    [1,1,2,2,3,3,4,4,5,5,6,6],1,6,6);
xy = [0,4,8,6,4,2;5,0,5,7,5,7]';
gplot (A,xy)
```

which creates an adjacency matrix `A` where node 1 is connected to nodes 2 and 6, node 2 with nodes 1 and 3, etc. The coordinates of the nodes are given in the n-by-2 matrix `xy`. See Figure 22.2.

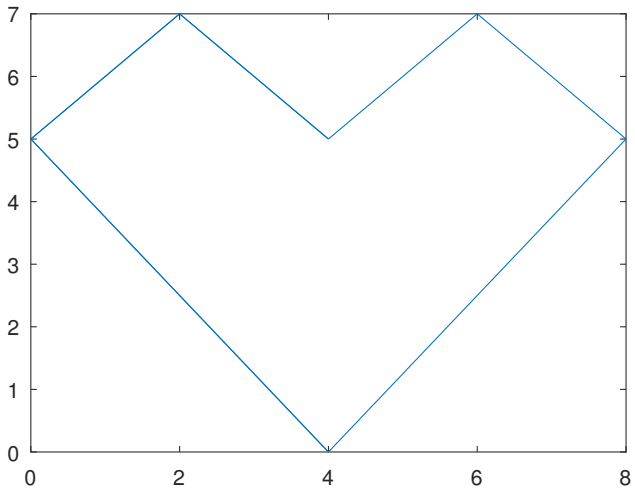



Figure 22.2: Simple use of the *gplot* command.

The dependencies between the nodes of a Cholesky factorization can be calculated in linear time without explicitly needing to calculate the Cholesky factorization by the `etree` command. This command returns the elimination tree of the matrix and can be displayed graphically by the command `treeplot (etree (A))` if `A` is symmetric or `treeplot (etree (A+A'))` otherwise.

`spy (`*x*`)`
`spy (...,` *markersize*`)`
`spy (...,` *line_spec*`)`
Plot the sparsity pattern of the sparse matrix *x*.

If the optional numeric argument *markersize* is given, it determines the size of the markers used in the plot.

If the optional string *line_spec* is given it is passed to `plot` and determines the appearance of the plot.

**See also:** [plot], page 334, [gplot], page 738.

*p* `= etree (`*S*`)`
*p* `= etree (`*S*`,` *typ*`)`
`[`*p*`,` *q*`] = etree (`*S*`,` *typ*`)`
Return the elimination tree for the matrix *S*.

By default *S* is assumed to be symmetric and the symmetric elimination tree is returned. The argument *typ* controls whether a symmetric or column elimination tree is returned. Valid values of *typ* are `"sym"` or `"col"`, for symmetric or column elimination tree respectively.

Called with a second argument, `etree` also returns the postorder permutations on the tree.

`etreeplot (A)`
`etreeplot (A, node_style, edge_style)`

Plot the elimination tree of the matrix $A$ or `A+A'` if $A$ in not symmetric.

The optional parameters *node_style* and *edge_style* define the output style.

**See also:** [treeplot], page 738, [gplot], page 738.

`gplot (A, xy)`
`gplot (A, xy, line_style)`
`[x, y] = gplot (A, xy)`

Plot a graph defined by $A$ and $xy$ in the graph theory sense.

$A$ is the adjacency matrix of the array to be plotted and $xy$ is an $n$-by-2 matrix containing the coordinates of the nodes of the graph.

The optional parameter *line_style* defines the output style for the plot. Called with no output arguments the graph is plotted directly. Otherwise, return the coordinates of the plot in $x$ and $y$.

**See also:** [treeplot], page 738, [etreeplot], page 738, [spy], page 737.

`treeplot (tree)`
`treeplot (tree, node_style, edge_style)`

Produce a graph of tree or forest.

The first argument is a row vector of parent indices.

The optional parameters *node_style* and *edge_style* define the output plot style.

The complexity of the algorithm is O(n) in terms of time and memory requirements.

**See also:** [etreeplot], page 738, [gplot], page 738.

`[x, y] = treelayout (tree)`
`[x, y] = treelayout (tree, permutation)`
`[x, y, h, s] = treelayout (...)`

treelayout lays out a tree or a forest.

The first argument *tree* is a vector of predecessors.

The optional parameter *permutation* is a postorder permutation.

The complexity of the algorithm is O(n) in terms of time and memory requirements.

**See also:** [etreeplot], page 738, [gplot], page 738, [treeplot], page 738.

## 22.1.4 Basic Operators and Functions on Sparse Matrices

### 22.1.4.1 Sparse Functions

Many Octave functions have been overloaded to work with either sparse or full matrices. There is no difference in calling convention when using an overloaded function with a sparse matrix, however, there is also no access to potentially sparse-specific features. At any time the sparse matrix specific version of a function can be used by explicitly calling its function name.

The table below lists all of the sparse functions of Octave. Note that the names of the specific sparse forms of the functions are typically the same as the general versions with a

*sp* prefix. In the table below, and in the rest of this article, the specific sparse versions of functions are used.

Generate sparse matrices:
*spalloc*, *spdiags*, *speye*, *sprand*, *sprandn*, *sprandsym*

Sparse matrix conversion:
*full*, *sparse*, *spconvert*

Manipulate sparse matrices
*issparse*, *nnz*, *nonzeros*, *nzmax*, *spfun*, *spones*, *spy*

Graph Theory:
*etree*, *etreeplot*, *gplot*, *treeplot*

Sparse matrix reordering:
*amd*, *ccolamd*, *colamd*, *colperm*, *csymamd*, *dmperm*, *symamd*, *randperm*, *symrcm*

Linear algebra:
*condest*, *eigs*, *matrix_type*, *normest*, *normest1*, *sprank*, *spaugment*, *svds*

Iterative techniques:
*ichol*, *ilu*, *pcg*, *pcr*

Miscellaneous:
*spparms*, *symbfact*, *spstats*

In addition all of the standard Octave mapper functions (i.e., basic math functions that take a single argument) such as *abs*, etc. can accept sparse matrices. The reader is referred to the documentation supplied with these functions within Octave itself for further details.

### 22.1.4.2 Return Types of Operators and Functions

The two basic reasons to use sparse matrices are to reduce the memory usage and to not have to do calculations on zero elements. The two are closely related in that the computation time on a sparse matrix operator or function is roughly linear with the number of nonzero elements.

Therefore, there is a certain density of nonzero elements of a matrix where it no longer makes sense to store it as a sparse matrix, but rather as a full matrix. For this reason operators and functions that have a high probability of returning a full matrix will always return one. For example adding a scalar constant to a sparse matrix will almost always make it a full matrix, and so the example,

```
speye (3) + 0
⇒   1  0  0
    0  1  0
    0  0  1
```

returns a full matrix as can be seen.

As all of the mixed operators and functions between full and sparse matrices exist, in general this does not cause any problems. However, one area where it does cause a problem is where a sparse matrix is promoted to a full matrix, where subsequent operations would resparsify the matrix. Such cases are rare, but can be artificially created, for example

`(fliplr (speye (3)) + speye (3)) - speye (3)` gives a full matrix when it should give a sparse one. In general, where such cases occur, they impose only a small memory penalty.

There is however one known case where this behavior of Octave's sparse matrices will cause a problem. That is in the handling of the *diag* function. Whether *diag* returns a sparse or full matrix depending on the type of its input arguments. So

```
a = diag (sparse ([1,2,3]), -1);
```

should return a sparse matrix. To ensure this actually happens, the *sparse* function, and other functions based on it like *speye*, always returns a sparse matrix, even if the memory used will be larger than its full representation.

### 22.1.4.3 Mathematical Considerations

The attempt has been made to make sparse matrices behave in exactly the same manner as there full counterparts. However, there are certain differences and especially differences with other products sparse implementations.

First, the `"./"` and `".^"` operators must be used with care. Consider what the examples

```
s = speye (4);
a1 = s .^ 2;
a2 = s .^ s;
a3 = s .^ -2;
a4 = s ./ 2;
a5 = 2 ./ s;
a6 = s ./ s;
```

will give. The first example of *s* raised to the power of 2 causes no problems. However *s* raised element-wise to itself involves a large number of terms `0 .^ 0` which is 1. There `s .^ s` is a full matrix.

Likewise `s .^ -2` involves terms like `0 .^ -2` which is infinity, and so `s .^ -2` is equally a full matrix.

For the `"./"` operator `s ./ 2` has no problems, but `2 ./ s` involves a large number of infinity terms as well and is equally a full matrix. The case of `s ./ s` involves terms like `0 ./ 0` which is a `NaN` and so this is equally a full matrix with the zero elements of *s* filled with `NaN` values.

The above behavior is consistent with full matrices, but is not consistent with sparse implementations in other products.

A particular problem of sparse matrices comes about due to the fact that as the zeros are not stored, the sign-bit of these zeros is equally not stored. In certain cases the sign-bit of zero is important. For example:

```
a = 0 ./ [-1, 1; 1, -1];
b = 1 ./ a
⇒ -Inf            Inf
    Inf           -Inf
c = 1 ./ sparse (a)
⇒  Inf            Inf
    Inf            Inf
```

To correct this behavior would mean that zero elements with a negative sign-bit would need to be stored in the matrix to ensure that their sign-bit was respected. This is not done

at this time, for reasons of efficiency, and so the user is warned that calculations where the sign-bit of zero is important must not be done using sparse matrices.

In general any function or operator used on a sparse matrix will result in a sparse matrix with the same or a larger number of nonzero elements than the original matrix. This is particularly true for the important case of sparse matrix factorizations. The usual way to address this is to reorder the matrix, such that its factorization is sparser than the factorization of the original matrix. That is the factorization of `L * U = P * S * Q` has sparser terms `L` and `U` than the equivalent factorization `L * U = S`.

Several functions are available to reorder depending on the type of the matrix to be factorized. If the matrix is symmetric positive-definite, then *symamd* or *csymamd* should be used. Otherwise *amd*, *colamd* or *ccolamd* should be used. For completeness the reordering functions *colperm* and *randperm* are also available.

See Figure 22.3, for an example of the structure of a simple positive definite matrix.

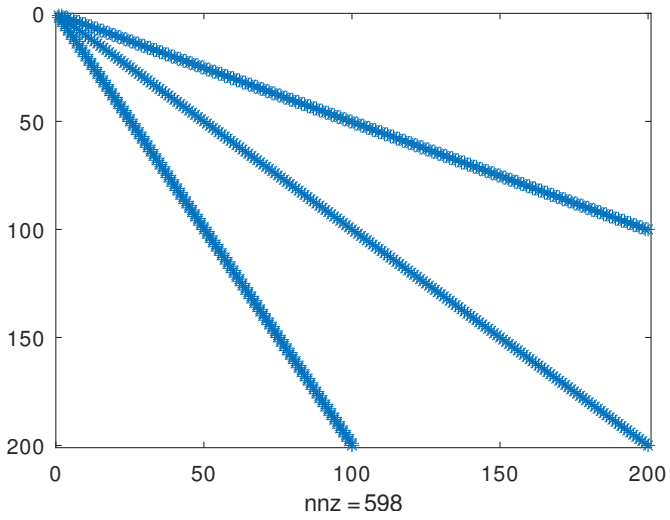



Figure 22.3: Structure of simple sparse matrix.

The standard Cholesky factorization of this matrix can be obtained by the same command that would be used for a full matrix. This can be visualized with the command `r = chol (A); spy (r);`. See Figure 22.4. The original matrix had 598 nonzero terms, while this Cholesky factorization has 10200, with only half of the symmetric matrix being stored. This is a significant level of fill in, and although not an issue for such a small test case, can represents a large overhead in working with other sparse matrices.

The appropriate sparsity preserving permutation of the original matrix is given by *symamd* and the factorization using this reordering can be visualized using the command `q = symamd (A); r = chol (A(q,q)); spy (r)`. This gives 399 nonzero terms which is a significant improvement.

The Cholesky factorization itself can be used to determine the appropriate sparsity preserving reordering of the matrix during the factorization, In that case this might be obtained with three return arguments as `[r, p, q] = chol (A); spy (r)`.

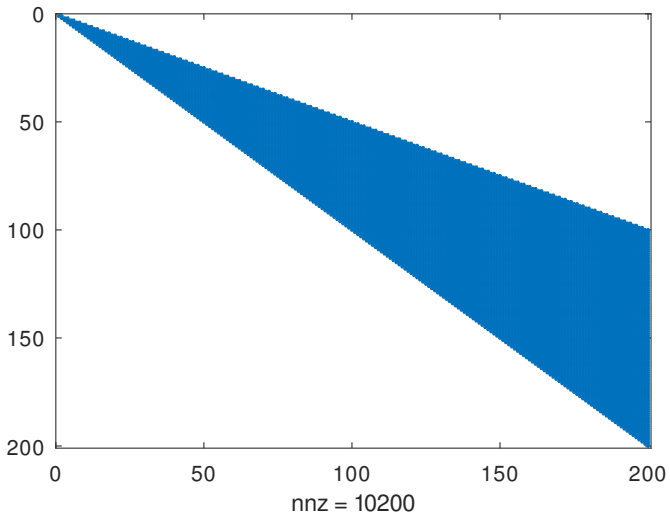



Figure 22.4: Structure of the unpermuted Cholesky factorization of the above matrix.

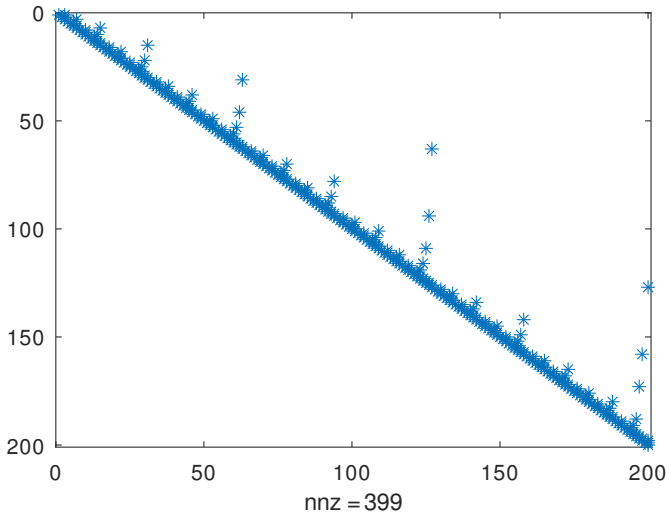



Figure 22.5: Structure of the permuted Cholesky factorization of the above matrix.

In the case of an asymmetric matrix, the appropriate sparsity preserving permutation is *colamd* and the factorization using this reordering can be visualized using the command `q = colamd (A); [l, u, p] = lu (A(:,q)); spy (l+u)`.

Finally, Octave implicitly reorders the matrix when using the div (/) and ldiv (\) operators, and so no the user does not need to explicitly reorder the matrix to maximize performance.

`p = amd (S)`
`p = amd (S, opts)`

Return the approximate minimum degree permutation of a matrix.

This is a permutation such that the Cholesky factorization of `S (p, p)` tends to be sparser than the Cholesky factorization of $S$ itself. `amd` is typically faster than `symamd` but serves a similar purpose.

The optional parameter *opts* is a structure that controls the behavior of `amd`. The fields of the structure are

*opts*.dense Determines what `amd` considers to be a dense row or column of the input matrix. Rows or columns with more than `max (16, (dense * sqrt (n)))` entries, where $n$ is the order of the matrix $S$, are ignored by `amd` during the calculation of the permutation. The value of dense must be a positive scalar and the default value is 10.0

*opts*.aggressive
If this value is a nonzero scalar, then `amd` performs aggressive absorption. The default is not to perform aggressive absorption.

The author of the code itself is Timothy A. Davis (see http://faculty.cse.tamu.edu/davis/suitesparse.html).

**See also:** [symamd], page 747, [colamd], page 744.

`p = ccolamd (S)`
`p = ccolamd (S, knobs)`
`p = ccolamd (S, knobs, cmember)`
`[p, stats] = ccolamd (...)`

Constrained column approximate minimum degree permutation.

`p = ccolamd (S)` returns the column approximate minimum degree permutation vector for the sparse matrix $S$. For a non-symmetric matrix $S$, `S(:, p)` tends to have sparser LU factors than $S$. `chol (S(:, p)' * S(:, p))` also tends to be sparser than `chol (S' * S)`. `p = ccolamd (S, 1)` optimizes the ordering for `lu (S(:, p))`. The ordering is followed by a column elimination tree post-ordering.

*knobs* is an optional 1-element to 5-element input vector, with a default value of `[0 10 10 1 0]` if not present or empty. Entries not present are set to their defaults.

`knobs(1)` if nonzero, the ordering is optimized for `lu (S(:, p))`. It will be a poor ordering for `chol (S(:, p)' * S(:, p))`. This is the most important knob for ccolamd.

`knobs(2)` if $S$ is m-by-n, rows with more than `max (16, knobs(2) * sqrt (n))` entries are ignored.

`knobs(3)` columns with more than `max (16, knobs(3) * sqrt (min (m, n)))` entries are ignored and ordered last in the output permutation (subject to the *cmember* constraints).

`knobs(4)` if nonzero, aggressive absorption is performed.

`knobs(5)` if nonzero, statistics and knobs are printed.

*cmember* is an optional vector of length $n$. It defines the constraints on the column ordering. If `cmember(j) = c`, then column $j$ is in constraint set $c$ ($c$ must be in the range 1 to n). In the output permutation $p$, all columns in set 1 appear first, followed by all columns in set 2, and so on. `cmember = ones (1,n)` if not present or empty. `ccolamd (S, [], 1 : n)` returns `1 : n`

`p = ccolamd (S)` is about the same as `p = colamd (S)`. *knobs* and its default values differ. `colamd` always does aggressive absorption, and it finds an ordering suitable for both `lu (S(:, p))` and `chol (S(:, p)' * S(:, p))`; it cannot optimize its ordering for `lu (S(:, p))` to the extent that `ccolamd (S, 1)` can.

*stats* is an optional 20-element output vector that provides data about the ordering and the validity of the input matrix $S$. Ordering statistics are in `stats(1 : 3)`. `stats(1)` and `stats(2)` are the number of dense or empty rows and columns ignored by CCOLAMD and `stats(3)` is the number of garbage collections performed on the internal data structure used by CCOLAMD (roughly of size `2.2 * nnz (S) + 4 * m + 7 * n` integers).

`stats(4 : 7)` provide information if CCOLAMD was able to continue. The matrix is OK if `stats(4)` is zero, or 1 if invalid. `stats(5)` is the rightmost column index that is unsorted or contains duplicate entries, or zero if no such column exists. `stats(6)` is the last seen duplicate or out-of-order row index in the column index given by `stats(5)`, or zero if no such row index exists. `stats(7)` is the number of duplicate or out-of-order row indices. `stats(8 : 20)` is always zero in the current version of CCOLAMD (reserved for future use).

The authors of the code itself are S. Larimore, T. Davis and S. Rajamanickam in collaboration with J. Bilbert and E. Ng. Supported by the National Science Foundation (DMS-9504974, DMS-9803599, CCR-0203270), and a grant from Sandia National Lab. See http://faculty.cse.tamu.edu/davis/suitesparse.html for ccolamd, csymamd, amd, colamd, symamd, and other related orderings.

**See also:** [colamd], page 744, [csymamd], page 745.

`p = colamd (S)`
`p = colamd (S, knobs)`
`[p, stats] = colamd (S)`
`[p, stats] = colamd (S, knobs)`

Compute the column approximate minimum degree permutation.

`p = colamd (S)` returns the column approximate minimum degree permutation vector for the sparse matrix $S$. For a non-symmetric matrix $S$, `S(:,p)` tends to have sparser LU factors than $S$. The Cholesky factorization of `S(:,p)' * S(:,p)` also tends to be sparser than that of `S' * S`.

*knobs* is an optional one- to three-element input vector. If $S$ is m-by-n, then rows with more than `max(16,knobs(1)*sqrt(n))` entries are ignored. Columns with more than `max (16,knobs(2)*sqrt(min(m,n)))` entries are removed prior to ordering, and ordered last in the output permutation $p$. Only completely dense rows or columns are removed if `knobs(1)` and `knobs(2)` are < 0, respectively. If `knobs(3)` is nonzero, *stats* and *knobs* are printed. The default is `knobs = [10 10 0]`. Note that *knobs* differs from earlier versions of colamd.

*stats* is an optional 20-element output vector that provides data about the ordering and the validity of the input matrix *S*. Ordering statistics are in `stats(1:3)`. `stats(1)` and `stats(2)` are the number of dense or empty rows and columns ignored by COLAMD and `stats(3)` is the number of garbage collections performed on the internal data structure used by COLAMD (roughly of size `2.2 * nnz(S) + 4 * m + 7 * n` integers).

Octave built-in functions are intended to generate valid sparse matrices, with no duplicate entries, with ascending row indices of the nonzeros in each column, with a non-negative number of entries in each column (!) and so on. If a matrix is invalid, then COLAMD may or may not be able to continue. If there are duplicate entries (a row index appears two or more times in the same column) or if the row indices in a column are out of order, then COLAMD can correct these errors by ignoring the duplicate entries and sorting each column of its internal copy of the matrix *S* (the input matrix *S* is not repaired, however). If a matrix is invalid in other ways then COLAMD cannot continue, an error message is printed, and no output arguments (*p* or *stats*) are returned. COLAMD is thus a simple way to check a sparse matrix to see if it's valid.

`stats(4:7)` provide information if COLAMD was able to continue. The matrix is OK if `stats(4)` is zero, or 1 if invalid. `stats(5)` is the rightmost column index that is unsorted or contains duplicate entries, or zero if no such column exists. `stats(6)` is the last seen duplicate or out-of-order row index in the column index given by `stats(5)`, or zero if no such row index exists. `stats(7)` is the number of duplicate or out-of-order row indices. `stats(8:20)` is always zero in the current version of COLAMD (reserved for future use).

The ordering is followed by a column elimination tree post-ordering.

The authors of the code itself are Stefan I. Larimore and Timothy A. Davis. The algorithm was developed in collaboration with John Gilbert, Xerox PARC, and Esmond Ng, Oak Ridge National Laboratory. (see http://faculty.cse.tamu.edu/davis/suitesparse.html)

**See also:** [colperm], page 745, [symamd], page 747, [ccolamd], page 743.

`p = colperm (s)`

Return the column permutations such that the columns of `s(:, p)` are ordered in terms of increasing number of nonzero elements.

If *s* is symmetric, then *p* is chosen such that `s(p, p)` orders the rows and columns with increasing number of nonzero elements.

`p = csymamd (S)`
`p = csymamd (S, knobs)`
`p = csymamd (S, knobs, cmember)`
`[p, stats] = csymamd (...)`

For a symmetric positive definite matrix *S*, return the permutation vector *p* such that `S(p,p)` tends to have a sparser Cholesky factor than *S*.

Sometimes `csymamd` works well for symmetric indefinite matrices too. The matrix *S* is assumed to be symmetric; only the strictly lower triangular part is referenced. *S* must be square. The ordering is followed by an elimination tree post-ordering.

*knobs* is an optional 1-element to 3-element input vector, with a default value of `[10 1 0]`. Entries not present are set to their defaults.

`knobs(1)` If $S$ is n-by-n, then rows and columns with more than `max(16,knobs(1)*sqrt(n))` entries are ignored, and ordered last in the output permutation (subject to the cmember constraints).

`knobs(2)` If nonzero, aggressive absorption is performed.

`knobs(3)` If nonzero, statistics and knobs are printed.

*cmember* is an optional vector of length n. It defines the constraints on the ordering. If `cmember(j) = S`, then row/column j is in constraint set $c$ ($c$ must be in the range 1 to n). In the output permutation $p$, rows/columns in set 1 appear first, followed by all rows/columns in set 2, and so on. `cmember = ones (1,n)` if not present or empty. `csymamd (S,[],1:n)` returns `1:n`.

`p = csymamd (S)` is about the same as `p = symamd (S)`. *knobs* and its default values differ.

`stats(4:7)` provide information if CCOLAMD was able to continue. The matrix is OK if `stats(4)` is zero, or 1 if invalid. `stats(5)` is the rightmost column index that is unsorted or contains duplicate entries, or zero if no such column exists. `stats(6)` is the last seen duplicate or out-of-order row index in the column index given by `stats(5)`, or zero if no such row index exists. `stats(7)` is the number of duplicate or out-of-order row indices. `stats(8:20)` is always zero in the current version of CCOLAMD (reserved for future use).

The authors of the code itself are S. Larimore, T. Davis and S. Rajamanickam in collaboration with J. Bilbert and E. Ng. Supported by the National Science Foundation (DMS-9504974, DMS-9803599, CCR-0203270), and a grant from Sandia National Lab. See http://faculty.cse.tamu.edu/davis/suitesparse.html for ccolamd, colamd, csymamd, amd, colamd, symamd, and other related orderings.

**See also:** [symamd], page 747, [ccolamd], page 743.

`p = dmperm (A)`
`[p, q, r, s, cc, rr] = dmperm (A)`

Perform a Dulmage-Mendelsohn permutation of the sparse matrix $A$.

With a single output argument `dmperm`, return a maximum matching $p$ such that `p(j) = i` if column $j$ is matched to row $i$, or 0 if column $j$ is unmatched. If $A$ is square and full structural rank, $p$ is a row permutation and `A(p,:)` has a zero-free diagonal. The structural rank of $A$ is `sprank(A) = sum(p>0)`.

Called with two or more output arguments, return the Dulmage-Mendelsohn decomposition of $A$. $p$ and $q$ are permutation vectors. *cc* and *rr* are vectors of length 5. `c = A(p,q)` is split into a 4-by-4 set of coarse blocks:

```
A11 A12 A13 A14
 0   0  A23 A24
 0   0   0  A34
 0   0   0  A44
```

where `A12`, `A23`, and `A34` are square with zero-free diagonals. The columns of `A11` are the unmatched columns, and the rows of `A44` are the unmatched rows. Any

of these blocks can be empty. In the "coarse" decomposition, the (i,j)-th block is `C(rr(i):rr(i+1)-1,cc(j):cc(j+1)-1)`. In terms of a linear system, `[A11 A12]` is the underdetermined part of the system (it is always rectangular and with more columns and rows, or 0-by-0), `A23` is the well-determined part of the system (it is always square), and `[A34 ; A44]` is the over-determined part of the system (it is always rectangular with more rows than columns, or 0-by-0).

The structural rank of $A$ is `sprank (A) = rr(4)-1`, which is an upper bound on the numerical rank of $A$. `sprank(A) = rank(full(sprand(A)))` with probability 1 in exact arithmetic.

The `A23` submatrix is further subdivided into block upper triangular form via the "fine" decomposition (the strongly-connected components of `A23`). If $A$ is square and structurally non-singular, `A23` is the entire matrix.

`C(r(i):r(i+1)-1,s(j):s(j+1)-1)` is the (i,j)-th block of the fine decomposition. The (1,1) block is the rectangular block `[A11 A12]`, unless this block is 0-by-0. The (b,b) block is the rectangular block `[A34 ; A44]`, unless this block is 0-by-0, where `b = length(r)-1`. All other blocks of the form `C(r(i):r(i+1)-1,s(i):s(i+1)-1)` are diagonal blocks of `A23`, and are square with a zero-free diagonal.

The method used is described in: A. Pothen & C.-J. Fan. "Computing the Block Triangular Form of a Sparse Matrix". *ACM Trans. Math. Software*, 16(4), pp. 303–324, 1990.

**See also:** [colamd], page 744, [ccolamd], page 743.

`p = symamd (S)`
`p = symamd (S, knobs)`
`[p, stats] = symamd (S)`
`[p, stats] = symamd (S, knobs)`

For a symmetric positive definite matrix $S$, returns the permutation vector p such that `S(p, p)` tends to have a sparser Cholesky factor than $S$.

Sometimes `symamd` works well for symmetric indefinite matrices too. The matrix $S$ is assumed to be symmetric; only the strictly lower triangular part is referenced. $S$ must be square.

*knobs* is an optional one- to two-element input vector. If $S$ is n-by-n, then rows and columns with more than `max (16,knobs(1)*sqrt(n))` entries are removed prior to ordering, and ordered last in the output permutation $p$. No rows/columns are removed if `knobs(1) < 0`. If `knobs(2)` is nonzero, *stats* and *knobs* are printed. The default is `knobs = [10 0]`. Note that *knobs* differs from earlier versions of `symamd`.

*stats* is an optional 20-element output vector that provides data about the ordering and the validity of the input matrix $S$. Ordering statistics are in `stats(1:3)`. `stats(1) = stats(2)` is the number of dense or empty rows and columns ignored by SYMAMD and `stats(3)` is the number of garbage collections performed on the internal data structure used by SYMAMD (roughly of size `8.4 * nnz (tril (S, -1)) + 9 * n` integers).

Octave built-in functions are intended to generate valid sparse matrices, with no duplicate entries, with ascending row indices of the nonzeros in each column, with a non-negative number of entries in each column (!) and so on. If a matrix is invalid,

then SYMAMD may or may not be able to continue. If there are duplicate entries (a row index appears two or more times in the same column) or if the row indices in a column are out of order, then SYMAMD can correct these errors by ignoring the duplicate entries and sorting each column of its internal copy of the matrix S (the input matrix S is not repaired, however). If a matrix is invalid in other ways then SYMAMD cannot continue, an error message is printed, and no output arguments (*p* or *stats*) are returned. SYMAMD is thus a simple way to check a sparse matrix to see if it's valid.

`stats(4:7)` provide information if SYMAMD was able to continue. The matrix is OK if `stats (4)` is zero, or 1 if invalid. `stats(5)` is the rightmost column index that is unsorted or contains duplicate entries, or zero if no such column exists. `stats(6)` is the last seen duplicate or out-of-order row index in the column index given by `stats(5)`, or zero if no such row index exists. `stats(7)` is the number of duplicate or out-of-order row indices. `stats(8:20)` is always zero in the current version of SYMAMD (reserved for future use).

The ordering is followed by a column elimination tree post-ordering.

The authors of the code itself are Stefan I. Larimore and Timothy A. Davis. The algorithm was developed in collaboration with John Gilbert, Xerox PARC, and Esmond Ng, Oak Ridge National Laboratory. (see http://faculty.cse.tamu.edu/davis/suitesparse.html)

**See also:** [colperm], page 745, [colamd], page 744.

`p = symrcm (S)`

Return the symmetric reverse Cuthill-McKee permutation of *S*.

*p* is a permutation vector such that `S(p, p)` tends to have its diagonal elements closer to the diagonal than *S*. This is a good preordering for LU or Cholesky factorization of matrices that come from "long, skinny" problems. It works for both symmetric and asymmetric *S*.

The algorithm represents a heuristic approach to the NP-complete bandwidth minimization problem. The implementation is based in the descriptions found in

E. Cuthill, J. McKee. "Reducing the Bandwidth of Sparse Symmetric Matrices". *Proceedings of the 24th ACM National Conference*, pp. 157–172, 1969, Brandon Press, New Jersey.

A. George, J.W.H. Liu. "Computer Solution of Large Sparse Positive Definite Systems", *Prentice Hall Series in Computational Mathematics*, ISBN 0-13-165274-5, 1981.

**See also:** [colperm], page 745, [colamd], page 744, [symamd], page 747.

## 22.2 Linear Algebra on Sparse Matrices

Octave includes a polymorphic solver for sparse matrices, where the exact solver used to factorize the matrix, depends on the properties of the sparse matrix itself. Generally, the cost of determining the matrix type is small relative to the cost of factorizing the matrix itself, but in any case the matrix type is cached once it is calculated, so that it is not re-determined each time it is used in a linear equation.

The selection tree for how the linear equation is solve is

1. If the matrix is diagonal, solve directly and goto 8
2. If the matrix is a permuted diagonal, solve directly taking into account the permutations. Goto 8
3. If the matrix is square, banded and if the band density is less than that given by `spparms ("bandden")` continue, else goto 4.
   a. If the matrix is tridiagonal and the right-hand side is not sparse continue, else goto 3b.
      1. If the matrix is Hermitian, with a positive real diagonal, attempt Cholesky factorization using LAPACK xPTSV.
      2. If the above failed or the matrix is not Hermitian with a positive real diagonal use Gaussian elimination with pivoting using LAPACK xGTSV, and goto 8.
   b. If the matrix is Hermitian with a positive real diagonal, attempt Cholesky factorization using LAPACK xPBTRF.
   c. if the above failed or the matrix is not Hermitian with a positive real diagonal use Gaussian elimination with pivoting using LAPACK xGBTRF, and goto 8.
4. If the matrix is upper or lower triangular perform a sparse forward or backward substitution, and goto 8
5. If the matrix is an upper triangular matrix with column permutations or lower triangular matrix with row permutations, perform a sparse forward or backward substitution, and goto 8
6. If the matrix is square, Hermitian with a real positive diagonal, attempt sparse Cholesky factorization using CHOLMOD.
7. If the sparse Cholesky factorization failed or the matrix is not Hermitian with a real positive diagonal, and the matrix is square, factorize, solve, and perform one refinement iteration using UMFPACK.
8. If the matrix is not square, or any of the previous solvers flags a singular or near singular matrix, find a minimum norm solution using CXSPARSE[3].

The band density is defined as the number of nonzero values in the band divided by the total number of values in the full band. The banded matrix solvers can be entirely disabled by using *spparms* to set `bandden` to 1 (i.e., `spparms ("bandden", 1)`).

The QR solver factorizes the problem with a Dulmage-Mendelsohn decomposition, to separate the problem into blocks that can be treated as over-determined, multiple well determined blocks, and a final over-determined block. For matrices with blocks of strongly connected nodes this is a big win as LU decomposition can be used for many blocks. It also significantly improves the chance of finding a solution to over-determined problems rather than just returning a vector of *NaN*'s.

All of the solvers above, can calculate an estimate of the condition number. This can be used to detect numerical stability problems in the solution and force a minimum norm solution to be used. However, for narrow banded, triangular or diagonal matrices, the cost

[3] The CHOLMOD, UMFPACK and CXSPARSE packages were written by Tim Davis and are available at http://faculty.cse.tamu.edu/davis/suitesparse.html

of calculating the condition number is significant, and can in fact exceed the cost of factoring the matrix. Therefore the condition number is not calculated in these cases, and Octave relies on simpler techniques to detect singular matrices or the underlying LAPACK code in the case of banded matrices.

The user can force the type of the matrix with the `matrix_type` function. This overcomes the cost of discovering the type of the matrix. However, it should be noted that identifying the type of the matrix incorrectly will lead to unpredictable results, and so `matrix_type` should be used with care.

```
nest = normest (A)
nest = normest (A, tol)
[nest, iter] = normest (...)
```

Estimate the 2-norm of the matrix $A$ using a power series analysis.

This is typically used for large matrices, where the cost of calculating `norm (A)` is prohibitive and an approximation to the 2-norm is acceptable.

*tol* is the tolerance to which the 2-norm is calculated. By default *tol* is 1e-6.

The optional output *iter* returns the number of iterations that were required for `normest` to converge.

**See also:** [normest1], page 750, [norm], page 654, [cond], page 648, [condest], page 751.

```
nest = normest1 (A)
nest = normest1 (A, t)
nest = normest1 (A, t, x0)
nest = normest1 (Afcn, t, x0, p1, p2, ...)
[nest, v] = normest1 (A, ...)
[nest, v, w] = normest1 (A, ...)
[nest, v, w, iter] = normest1 (A, ...)
```

Estimate the 1-norm of the matrix $A$ using a block algorithm.

`normest1` is best for large sparse matrices where only an estimate of the norm is required. For small to medium sized matrices, consider using `norm (A, 1)`. In addition, `normest1` can be used for the estimate of the 1-norm of a linear operator $A$ when matrix-vector products `A * x` and `A' * x` can be cheaply computed. In this case, instead of the matrix $A$, a function `Afcn (flag, x)` is used; it must return:

- the dimension $n$ of $A$, if *flag* is `"dim"`
- true if $A$ is a real operator, if *flag* is `"real"`
- the result `A * x`, if *flag* is `"notransp"`
- the result `A' * x`, if *flag* is `"transp"`

A typical case is $A$ defined by `b ^ m`, in which the result `A * x` can be computed without even forming explicitly `b ^ m` by:

```
y = x;
for i = 1:m
  y = b * y;
endfor
```

The parameters *p1*, *p2*, ... are arguments of `Afcn (flag, x, p1, p2, ...)`.

The default value for $t$ is 2. The algorithm requires matrix-matrix products with sizes $n$ x $n$ and $n$ x $t$.

The initial matrix *x0* should have columns of unit 1-norm. The default initial matrix *x0* has the first column `ones (n, 1) / n` and, if $t$ > 1, the remaining columns with random elements `-1 / n`, `1 / n`, divided by $n$.

On output, *nest* is the desired estimate, $v$ and $w$ are vectors such that `w = A * v`, with `norm (w, 1)` = $c$ * `norm (v, 1)`. *iter* contains in `iter(1)` the number of iterations (the maximum is hardcoded to 5) and in `iter(2)` the total number of products `A * x` or `A' * x` performed by the algorithm.

Algorithm Note: `normest1` uses random numbers during evaluation. Therefore, if consistent results are required, the `"state"` of the random generator should be fixed before invoking `normest1`.

Reference: N. J. Higham and F. Tisseur, "A block algorithm for matrix 1-norm estimation, with and application to 1-norm pseudospectra", *SIAM J. Matrix Anal. Appl.*, Vol. 21, No. 4, pp. 1185–1201, 2000.

**See also:** [normest], page 750, [norm], page 654, [cond], page 648, [condest], page 751.

```
cest = condest (A)
cest = condest (A, t)
cest = condest (A, Ainvfcn)
cest = condest (A, Ainvfcn, t)
cest = condest (A, Ainvfcn, t, p1, p2, ...)
cest = condest (Afcn, Ainvfcn)
cest = condest (Afcn, Ainvfcn, t)
cest = condest (Afcn, Ainvfcn, t, p1, p2, ...)
[cest, v] = condest (...)
```

Estimate the 1-norm condition number of a square matrix $A$ using $t$ test vectors and a randomized 1-norm estimator.

The optional input $t$ specifies the number of test vectors (default 5).

The input may be a matrix $A$ (the algorithm is particularly appropriate for large, sparse matrices). Alternatively, the behavior of the matrix can be defined implicitly by functions. When using an implicit definition, `condest` requires the following functions:

– `Afcn (flag, x)` which must return
  - the dimension $n$ of $A$, if *flag* is `"dim"`
  - true if $A$ is a real operator, if *flag* is `"real"`
  - the result `A * x`, if *flag* is "notransp"
  - the result `A' * x`, if *flag* is "transp"

– `Ainvfcn (flag, x)` which must return
  - the dimension $n$ of `inv (A)`, if *flag* is `"dim"`
  - true if `inv (A)` is a real operator, if *flag* is `"real"`
  - the result `inv (A) * x`, if *flag* is "notransp"
  - the result `inv (A)' * x`, if *flag* is "transp"

Any parameters *p1*, *p2*, ... are additional arguments of `Afcn (flag, x, p1, p2, ...)` and `Ainvfcn (flag, x, p1, p2, ...)`.

The principal output is the 1-norm condition number estimate *cest*.

The optional second output *v* is an approximate null vector; it satisfies the equation `norm (A*v, 1) == norm (A, 1) * norm (v, 1) / cest`.

Algorithm Note: `condest` uses a randomized algorithm to approximate the 1-norms. Therefore, if consistent results are required, the `"state"` of the random generator should be fixed before invoking `condest`.

References:

- N.J. Higham and F. Tisseur, "A Block Algorithm for Matrix 1-Norm Estimation, with an Application to 1-Norm Pseudospectra", *SIAM Journal on Matrix Analysis and Applications*, Vol. 21, Iss. 4, pp. 1185–1201, 2000, https://dx.doi.org/10.1137/S0895479899356080.
- N.J. Higham and F. Tisseur, *A Block Algorithm for Matrix 1-Norm Estimation, with an Application to 1-Norm Pseudospectra*, https://citeseer.ist.psu.edu/223007.html.

**See also:** [cond], page 648, [rcond], page 656, [norm], page 654, [normest1], page 750, [normest], page 750.

```
  spparms ()
vals = spparms ()
[keys, vals] = spparms ()
val = spparms (key)
  spparms (vals)
  spparms ("default")
  spparms ("tight")
  spparms (key, val)
```

Query or set the parameters used by the sparse solvers and factorization functions.

The first four calls above get information about the current settings, while the others change the current settings. The parameters are stored as pairs of keys and values, where the values are all floats and the keys are one of the following strings:

'`spumoni`' Printing level of debugging information of the solvers (default 0)

'`ths_rel`' Included for compatibility. Not used. (default 1)

'`ths_abs`' Included for compatibility. Not used. (default 1)

'`exact_d`' Included for compatibility. Not used. (default 0)

'`supernd`' Included for compatibility. Not used. (default 3)

'`rreduce`' Included for compatibility. Not used. (default 3)

'`wh_frac`' Included for compatibility. Not used. (default 0.5)

'`autommd`' Flag whether the LU/QR and the '\' and '/' operators will automatically use the sparsity preserving mmd functions (default 1)

'`autoamd`' Flag whether the LU and the '\' and '/' operators will automatically use the sparsity preserving amd functions (default 1)

'piv_tol' The pivot tolerance of the UMFPACK solvers (default 0.1)

'sym_tol' The pivot tolerance of the UMFPACK symmetric solvers (default 0.001)

'bandden' The density of nonzero elements in a banded matrix before it is treated by the LAPACK banded solvers (default 0.5)

'umfpack' Flag whether the UMFPACK or mmd solvers are used for the LU, '\' and '/' operations (default 1)

The value of individual keys can be set with `spparms (key, val)`. The default values can be restored with the special keyword `"default"`. The special keyword `"tight"` can be used to set the mmd solvers to attempt a sparser solution at the potential cost of longer running time.

**See also:** [chol], page 657, [colamd], page 744, [lu], page 659, [qr], page 661, [symamd], page 747.

`p = sprank (S)`

Calculate the structural rank of the sparse matrix *S*.

Note that only the structure of the matrix is used in this calculation based on a Dulmage-Mendelsohn permutation to block triangular form. As such the numerical rank of the matrix *S* is bounded by `sprank (S) >= rank (S)`. Ignoring floating point errors `sprank (S) == rank (S)`.

**See also:** [dmperm], page 746.

`[count, h, parent, post, R] = symbfact (S)`
`[...] = symbfact (S, typ)`
`[...] = symbfact (S, typ, mode)`

Perform a symbolic factorization analysis of the sparse matrix *S*.

The input variables are

*S* *S* is a real or complex sparse matrix.

*typ* Is the type of the factorization and can be one of

`"sym"` (default)
Factorize *S*. Assumes *S* is symmetric and uses the upper triangular portion of the matrix.

`"col"` Factorize `S' * S`.

`"row"` Factorize `S * S'`.

`"lo"` Factorize `S'`. Assumes *S* is symmetric and uses the lower triangular portion of the matrix.

*mode* When *mode* is unspecified return the Cholesky factorization for *R*. If *mode* is `"lower"` or `"L"` then return the conjugate transpose `R'` which is a lower triangular factor. The conjugate transpose version is faster and uses less memory, but still returns the same values for all other outputs: *count*, *h*, *parent*, and *post*.

The output variables are:

*count* The row counts of the Cholesky factorization as determined by *typ*. The computational difficulty of performing the true factorization using `chol` is `sum (count .^ 2)`.

*h* The height of the elimination tree.

*parent* The elimination tree itself.

*post* A sparse boolean matrix whose structure is that of the Cholesky factorization as determined by *typ*.

**See also:** [chol], page 657, [etree], page 737, [treelayout], page 738.

For non square matrices, the user can also utilize the `spaugment` function to find a least squares solution to a linear equation.

`s = spaugment (A, c)`

Create the augmented matrix of *A*.

This is given by

```
[c * eye(m, m), A;
             A', zeros(n, n)]
```

This is related to the least squares solution of `A \ b`, by

```
s * [ r / c; x] = [ b, zeros(n, columns(b)) ]
```

where *r* is the residual error

```
r = b - A * x
```

As the matrix *s* is symmetric indefinite it can be factorized with `lu`, and the minimum norm solution can therefore be found without the need for a `qr` factorization. As the residual error will be `zeros (m, m)` for underdetermined problems, and example can be

```
m = 11; n = 10; mn = max (m, n);
A = spdiags ([ones(mn,1), 10*ones(mn,1), -ones(mn,1)],
             [-1, 0, 1], m, n);
x0 = A \ ones (m,1);
s = spaugment (A);
[L, U, P, Q] = lu (s);
x1 = Q * (U \ (L \ (P  * [ones(m,1); zeros(n,1)])));
x1 = x1(end - n + 1 : end);
```

To find the solution of an overdetermined problem needs an estimate of the residual error *r* and so it is more complex to formulate a minimum norm solution using the `spaugment` function.

In general the left division operator is more stable and faster than using the `spaugment` function.

**See also:** [mldivide], page 172.

Finally, the function `eigs` can be used to calculate a limited number of eigenvalues and eigenvectors based on a selection criteria and likewise for `svds` which calculates a limited number of singular values and vectors.

`d = eigs (A)`
`d = eigs (A, k)`
`d = eigs (A, k, sigma)`
`d = eigs (A, k, sigma, opts)`
`d = eigs (A, B)`
`d = eigs (A, B, k)`
`d = eigs (A, B, k, sigma)`
`d = eigs (A, B, k, sigma, opts)`
`d = eigs (Af, n)`
`d = eigs (Af, n, k)`
`d = eigs (Af, n, k, sigma)`
`d = eigs (Af, n, k, sigma, opts)`
`d = eigs (Af, n, B)`
`d = eigs (Af, n, B, k)`
`d = eigs (Af, n, B, k, sigma)`
`d = eigs (Af, n, B, k, sigma, opts)`
`[V, D] = eigs (...)`
`[V, D, flag] = eigs (...)`

Calculate a limited number of eigenvalues and eigenvectors based on a selection criteria.

By default, `eigs` solve the equation $A\nu = \lambda\nu$, where $\lambda$ is a scalar representing one of the eigenvalues, and $\nu$ is the corresponding eigenvector. If given the positive definite matrix $B$ then `eigs` solves the general eigenvalue equation $A\nu = \lambda B\nu$.

The input $A$ is a square matrix of dimension $n$-by-$n$. Typically, $A$ is also large and sparse.

The input $B$ for the generalized eigenvalue problem is a square matrix with the same size as $A$ ($n$-by-$n$). Typically, $B$ is also large and sparse.

The number of eigenvalues and eigenvectors to calculate is given by $k$ and defaults to 6.

The argument *sigma* determines which eigenvalues are returned. *sigma* can be either a scalar or a string. When *sigma* is a scalar, the $k$ eigenvalues closest to *sigma* are returned. If *sigma* is a string, it must be one of the following values.

`"lm"` Largest Magnitude (default).

`"sm"` Smallest Magnitude.

`"la"` Largest Algebraic (valid only for real symmetric problems).

`"sa"` Smallest Algebraic (valid only for real symmetric problems).

`"be"` Both Ends, with one more from the high-end if $k$ is odd (valid only for real symmetric problems).

`"lr"` Largest Real part (valid only for complex or unsymmetric problems).

`"sr"` Smallest Real part (valid only for complex or unsymmetric problems).

`"li"` Largest Imaginary part (valid only for complex or unsymmetric problems).

`"si"` Smallest Imaginary part (valid only for complex or unsymmetric problems).

If *opts* is given, it is a structure defining possible options that `eigs` should use. The fields of the *opts* structure are:

`issym` If *Af* is given then this flag (true/false) determines whether the function *Af* defines a symmetric problem. It is ignored if a matrix *A* is given. The default is false.

`isreal` If *Af* is given then this flag (true/false) determines whether the function *Af* defines a real problem. It is ignored if a matrix *A* is given. The default is true.

`tol` Defines the required convergence tolerance, calculated as `tol * norm (A)`. The default is `eps`.

`maxit` The maximum number of iterations. The default is 300.

`p` The number of Lanczos basis vectors to use. More vectors will result in faster convergence, but a greater use of memory. The optimal value of `p` is problem dependent and should be in the range *k* + 1 to *n*. The default value is `2 * k`.

`v0` The starting vector for the algorithm. An initial vector close to the final vector will speed up convergence. The default is for ARPACK to randomly generate a starting vector. If specified, `v0` must be an *n*-by-1 vector where *n* = `rows (A)`.

`disp` The level of diagnostic printout (0|1|2). If `disp` is 0 then diagnostics are disabled. The default value is 0.

`cholB` If the generalized eigenvalue problem is being calculated, this flag (true/false) specifies whether the *B* input represents `chol (B)` or simply the matrix *B*. The default is false.

`permB` The permutation vector of the Cholesky factorization for *B* if `cholB` is true. It is obtained by `[R, ~, permB] = chol (B, "vector")`. The default is `1:n`.

It is also possible to represent *A* by a function denoted *Af*. *Af* must be followed by a scalar argument *n* defining the length of the vector argument accepted by *Af*. *Af* can be a function handle, an inline function, or a string. When *Af* is a string it holds the name of the function to use.

*Af* is a function of the form `y = Af (x)` where the required return value of *Af* is determined by the value of *sigma*. The four possible forms are

`A * x` if *sigma* is not given or is a string other than "sm".

`A \ x` if *sigma* is 0 or "sm".

`(A - sigma * I) \ x`
if *sigma* is a scalar not equal to 0; `I` is the identity matrix of the same size as *A*.

`(A - sigma * B) \ x`
: for the general eigenvalue problem.

The return arguments and their form depend on the number of return arguments requested. For a single return argument, a column vector $d$ of length $k$ is returned containing the $k$ eigenvalues that have been found. For two return arguments, $V$ is an $n$-by-$k$ matrix whose columns are the $k$ eigenvectors corresponding to the returned eigenvalues. The eigenvalues themselves are returned in $D$ in the form of a $k$-by-$k$ matrix, where the elements on the diagonal are the eigenvalues.

The third return argument *flag* returns the status of the convergence. If *flag* is 0 then all eigenvalues have converged. Any other value indicates a failure to converge.

Programming Notes: For small problems, $n$ < 500, consider using `eig (full (A))`.

If ARPACK fails to converge consider increasing the number of Lanczos vectors (*opt*.p), increasing the number of iterations (*opt*.maxiter), or decreasing the tolerance (*opt*.tol).

Reference: This function is based on the ARPACK package, written by R. Lehoucq, K. Maschhoff, D. Sorensen, and C. Yang. For more information see http://www.caam.rice.edu/software/ARPACK/.

**See also:** [eig], page 649, [svds], page 757.

`s = svds (A)`
`s = svds (A, k)`
`s = svds (A, k, sigma)`
`s = svds (A, k, sigma, opts)`
`[u, s, v] = svds (...)`
`[u, s, v, flag] = svds (...)`
Find a few singular values of the matrix $A$.

The singular values are calculated using

```
[m, n] = size (A);
s = eigs ([sparse(m, m), A;
                         A', sparse(n, n)])
```

The eigenvalues returned by `eigs` correspond to the singular values of $A$. The number of singular values to calculate is given by $k$ and defaults to 6.

The argument *sigma* specifies which singular values to find. When *sigma* is the string `'L'`, the default, the largest singular values of $A$ are found. Otherwise, *sigma* must be a real scalar and the singular values closest to *sigma* are found. As a corollary, `sigma = 0` finds the smallest singular values. Note that for relatively small values of *sigma*, there is a chance that the requested number of singular values will not be found. In that case *sigma* should be increased.

*opts* is a structure defining options that `svds` will pass to `eigs`. The possible fields of this structure are documented in `eigs`. By default, `svds` sets the following three fields:

`tol`
: The required convergence tolerance for the singular values. The default value is 1e-10. `eigs` is passed `tol / sqrt (2)`.

`maxit`
: The maximum number of iterations. The default is 300.

`disp` The level of diagnostic printout (0|1|2). If `disp` is 0 then diagnostics are disabled. The default value is 0.

If more than one output is requested then `svds` will return an approximation of the singular value decomposition of $A$

```
A_approx = u*s*v'
```

where $A$_approx is a matrix of size $A$ but only rank $k$.

*flag* returns 0 if the algorithm has successfully converged, and 1 otherwise. The test for convergence is

```
norm (A*v - u*s, 1) <= tol * norm (A, 1)
```

`svds` is best for finding only a few singular values from a large sparse matrix. Otherwise, `svd (full (A))` will likely be more efficient.

**See also:** [svd], page 669, [eigs], page 754.

## 22.3 Iterative Techniques Applied to Sparse Matrices

The left division `\` and right division `/` operators, discussed in the previous section, use direct solvers to resolve a linear equation of the form `x = A \ b` or `x = b / A`. Octave also includes a number of functions to solve sparse linear equations using iterative techniques.

```
x = pcg (A, b, tol, maxit, m1, m2, x0, ...)
x = pcg (A, b, tol, maxit, M, [], x0, ...)
[x, flag, relres, iter, resvec, eigest] = pcg (A, b, ...)
```

Solve the linear system of equations `A * x = b` by means of the Preconditioned Conjugate Gradient iterative method.

The input arguments are:

- $A$ is the matrix of the linear system and it must be square. $A$ can be passed as a matrix, function handle, or inline function `Afcn` such that `Afcn(x) = A * x`. Additional parameters to `Afcn` may be passed after *x0*.

  $A$ has to be Hermitian and Positive Definite (HPD). If `pcg` detects $A$ not to be positive definite, a warning is printed and the *flag* output is set.

- $b$ is the right-hand side vector.

- *tol* is the required relative tolerance for the residual error, `b - A * x`. The iteration stops if `norm (b - A * x)` $\leq$ *tol* `* norm (b)`. If *tol* is omitted or empty, then a tolerance of 1e-6 is used.

- *maxit* is the maximum allowed number of iterations; if *maxit* is omitted or empty then a value of 20 is used.

- $m$ is a HPD preconditioning matrix. For any decomposition `m = p1 * p2` such that `inv (p1) * A * inv (p2)` is HPD, the conjugate gradient method is formally applied to the linear system `inv (p1) * A * inv (p2) * y = inv (p1) * b`, with `x = inv (p2) * y` (split preconditioning). In practice, at each iteration of the conjugate gradient method a linear system with matrix $m$ is solved with `mldivide`. If a particular factorization `m = m1 * m2` is available (for instance, an incomplete Cholesky factorization of $a$), the two matrices *m1* and *m2* can be passed and the relative linear systems are solved with the `mldivide` operator. Note that a proper

choice of the preconditioner may dramatically improve the overall performance of the method. Instead of matrices *m1* and *m2*, the user may pass two functions which return the results of applying the inverse of *m1* and *m2* to a vector. If *m1* is omitted or empty `[]`, then no preconditioning is applied. If no factorization of *m* is available, *m2* can be omitted or left [], and the input variable *m1* can be used to pass the preconditioner *m*.

- *x0* is the initial guess. If *x0* is omitted or empty then the function sets *x0* to a zero vector by default.

The arguments which follow *x0* are treated as parameters, and passed in an appropriate manner to any of the functions (*A* or *m1* or *m2*) that have been given to `pcg`. See the examples below for further details.

The output arguments are:

- *x* is the computed approximation to the solution of `A * x = b`. If the algorithm did not converge, then *x* is the iteration which has the minimum residual.
- *flag* reports on the convergence:
  - 0: The algorithm converged to within the prescribed tolerance.
  - 1: The algorithm did not converge and it reached the maximum number of iterations.
  - 2: The preconditioner matrix is singular.
  - 3: The algorithm stagnated, i.e., the absolute value of the difference between the current iteration *x* and the previous is less than `eps * norm (x,2)`.
  - 4: The algorithm detects that the input (preconditioned) matrix is not HPD.
- *relres* is the ratio of the final residual to its initial value, measured in the Euclidean norm.
- *iter* indicates the iteration of *x* which it was computed. Since the output *x* corresponds to the minimal residual solution, the total number of iterations that the method performed is given by `length(resvec) - 1`.
- *resvec* describes the convergence history of the method. `resvec (i, 1)` is the Euclidean norm of the residual, and `resvec (i, 2)` is the preconditioned residual norm, after the (*i*-1)-th iteration, `i = 1, 2, ..., iter+1`. The preconditioned residual norm is defined as `r' * (m \ r)` where `r = b - A * x`, see also the description of *m*. If *eigest* is not required, only `resvec (:, 1)` is returned.
- *eigest* returns the estimate for the smallest `eigest(1)` and largest `eigest(2)` eigenvalues of the preconditioned matrix `P = m \ A`. In particular, if no preconditioning is used, the estimates for the extreme eigenvalues of *A* are returned. `eigest(1)` is an overestimate and `eigest(2)` is an underestimate, so that `eigest(2) / eigest(1)` is a lower bound for `cond (P, 2)`, which nevertheless in the limit should theoretically be equal to the actual value of the condition number.

Let us consider a trivial problem with a tridiagonal matrix

```
n = 10;
A = toeplitz (sparse ([1, 1], [1, 2], [2, 1], 1, n));
b = A * ones (n, 1);
M1 = ichol (A); # in this tridiagonal case it corresponds to chol (A)'
M2 = M1';
M = M1 * M2;
Afcn = @(x) A * x;
Mfcn = @(x) M \ x;
M1fcn = @(x) M1 \ x;
M2fcn = @(x) M2 \ x;
```

EXAMPLE 1: Simplest use of `pcg`

```
x = pcg (A, b)
```

EXAMPLE 2: `pcg` with a function which computes *A * x*

```
x = pcg (Afcn, b)
```

EXAMPLE 3: `pcg` with a preconditioner matrix *M*

```
x = pcg (A, b, 1e-06, 100, M)
```

EXAMPLE 4: `pcg` with a function as preconditioner

```
x = pcg (Afcn, b, 1e-6, 100, Mfcn)
```

EXAMPLE 5: `pcg` with preconditioner matrices *M1* and *M2*

```
x = pcg (A, b, 1e-6, 100, M1, M2)
```

EXAMPLE 6: `pcg` with functions as preconditioners

```
x = pcg (Afcn, b, 1e-6, 100, M1fcn, M2fcn)
```

EXAMPLE 7: `pcg` with as input a function requiring an argument

```
  function y = Ap (A, x, p) # compute A^p * x
     y = x;
     for i = 1:p
       y = A * y;
     endfor
  endfunction
Apfcn = @(x, p) Ap (A, x, p);
x = pcg (Apfcn, b, [], [], [], [], [], 2);
```

EXAMPLE 8: explicit example to show that `pcg` uses a split preconditioner

```
M1 = ichol (A + 0.1 * eye (n)); # factorization of A perturbed
M2 = M1';
M = M1 * M2;

## reference solution computed by pcg after two iterations
[x_ref, fl] = pcg (A, b, [], 2, M)

## split preconditioning
[y, fl] = pcg ((M1 \ A) / M2, M1 \ b, [], 2)
x = M2 \ y # compare x and x_ref
```

References:

1. C.T. Kelley, *Iterative Methods for Linear and Nonlinear Equations*, SIAM, 1995. (the base PCG algorithm)

2. Y. Saad, *Iterative Methods for Sparse Linear Systems*, PWS, 1996. (condition number estimate from PCG) Revised version of this book is available online at https://www-users.cs.umn.edu/~saad/books.html

**See also:** [sparse], page 732, [pcr], page 761, [gmres], page 682, [bicg], page 676, [bicgstab], page 678, [cgs], page 680.

`x = pcr (A, b, tol, maxit, m, x0, ...)`
`[x, flag, relres, iter, resvec] = pcr (...)`

Solve the linear system of equations `A * x = b` by means of the Preconditioned Conjugate Residuals iterative method.

The input arguments are

- *A* can be either a square (preferably sparse) matrix or a function handle, inline function or string containing the name of a function which computes `A * x`. In principle *A* should be symmetric and non-singular; if `pcr` finds *A* to be numerically singular, you will get a warning message and the *flag* output parameter will be set.
- *b* is the right hand side vector.
- *tol* is the required relative tolerance for the residual error, `b - A * x`. The iteration stops if `norm (b - A * x) <= tol * norm (b - A * x0)`. If *tol* is empty or is omitted, the function sets `tol = 1e-6` by default.
- *maxit* is the maximum allowable number of iterations; if `[]` is supplied for *maxit*, or `pcr` has less arguments, a default value equal to 20 is used.
- *m* is the (left) preconditioning matrix, so that the iteration is (theoretically) equivalent to solving by `pcr P * x = m \ b`, with `P = m \ A`. Note that a proper choice of the preconditioner may dramatically improve the overall performance of the method. Instead of matrix *m*, the user may pass a function which returns the results of applying the inverse of *m* to a vector (usually this is the preferred way of using the preconditioner). If `[]` is supplied for *m*, or *m* is omitted, no preconditioning is applied.
- *x0* is the initial guess. If *x0* is empty or omitted, the function sets *x0* to a zero vector by default.

The arguments which follow *x0* are treated as parameters, and passed in a proper way to any of the functions (*A* or *m*) which are passed to `pcr`. See the examples below for further details.

The output arguments are

- *x* is the computed approximation to the solution of `A * x = b`.
- *flag* reports on the convergence. `flag = 0` means the solution converged and the tolerance criterion given by *tol* is satisfied. `flag = 1` means that the *maxit* limit for the iteration count was reached. `flag = 3` reports a `pcr` breakdown, see [1] for details.
- *relres* is the ratio of the final residual to its initial value, measured in the Euclidean norm.
- *iter* is the actual number of iterations performed.

- *resvec* describes the convergence history of the method, so that `resvec (i)` contains the Euclidean norms of the residual after the (*i*-1)-th iteration, `i = 1,2, ..., iter+1`.

Let us consider a trivial problem with a diagonal matrix (we exploit the sparsity of A)

```
n = 10;
A = sparse (diag (1:n));
b = rand (N, 1);
```

Example 1: Simplest use of `pcr`

```
x = pcr (A, b)
```

Example 2: `pcr` with a function which computes `A * x`.

```
function y = apply_a (x)
  y = [1:10]' .* x;
endfunction

x = pcr ("apply_a", b)
```

Example 3: Preconditioned iteration, with full diagnostics. The preconditioner (quite strange, because even the original matrix *A* is trivial) is defined as a function

```
function y = apply_m (x)
  k = floor (length (x) - 2);
  y = x;
  y(1:k) = x(1:k) ./ [1:k]';
endfunction

[x, flag, relres, iter, resvec] = ...
                    pcr (A, b, [], [], "apply_m")
semilogy ([1:iter+1], resvec);
```

Example 4: Finally, a preconditioner which depends on a parameter *k*.

```
function y = apply_m (x, varargin)
  k = varargin{1};
  y = x;
  y(1:k) = x(1:k) ./ [1:k]';
endfunction

[x, flag, relres, iter, resvec] = ...
                    pcr (A, b, [], [], "apply_m"', [], 3)
```

Reference:

W. Hackbusch, *Iterative Solution of Large Sparse Systems of Equations*, section 9.5.4, Springer, 1994.

**See also:** [sparse], page 732, [pcg], page 758.

The speed with which an iterative solver converges to a solution can be accelerated with the use of a pre-conditioning matrix *M*. In this case the linear equation `M^-1 * x = M^-1 * A \ b` is solved instead. Typical pre-conditioning matrices are partial factorizations of the original matrix.

`L = ichol (A)`
`L = ichol (A, opts)`

Compute the incomplete Cholesky factorization of the sparse square matrix $A$.

By default, `ichol` uses only the lower triangle of $A$ and produces a lower triangular factor $L$ such that `L*L'` approximates $A$.

The factor given by this routine may be useful as a preconditioner for a system of linear equations being solved by iterative methods such as PCG (Preconditioned Conjugate Gradient).

The factorization may be modified by passing options in a structure *opts*. The option name is a field of the structure and the setting is the value of field. Names and specifiers are case sensitive.

type — Type of factorization.

`"nofill"` (default)
: Incomplete Cholesky factorization with no fill-in (IC(0)).

`"ict"`
: Incomplete Cholesky factorization with threshold dropping (ICT).

diagcomp — A non-negative scalar *alpha* for incomplete Cholesky factorization of `A + alpha * diag (diag (A))` instead of $A$. This can be useful when $A$ is not positive definite. The default value is 0.

droptol — A non-negative scalar specifying the drop tolerance for factorization if performing ICT. The default value is 0 which produces the complete Cholesky factorization.

Non-diagonal entries of $L$ are set to 0 unless

`abs (L(i,j)) >= droptol * norm (A(j:end, j), 1)`.

michol — Modified incomplete Cholesky factorization:

`"off"` (default)
: Row and column sums are not necessarily preserved.

`"on"`
: The diagonal of $L$ is modified so that row (and column) sums are preserved even when elements have been dropped during the factorization. The relationship preserved is: `A * e = L * L' * e`, where e is a vector of ones.

shape

`"lower"` (default)
: Use only the lower triangle of $A$ and return a lower triangular factor $L$ such that `L*L'` approximates $A$.

`"upper"`
: Use only the upper triangle of $A$ and return an upper triangular factor $U$ such that `U'*U` approximates $A$.

EXAMPLES

The following problem demonstrates how to factorize a sample symmetric positive definite matrix with the full Cholesky decomposition and with the incomplete one.

```
A = [ 0.37, -0.05,  -0.05,  -0.07;
     -0.05,  0.116,  0.0,   -0.05;
     -0.05,  0.0,    0.116, -0.05;
     -0.07, -0.05,  -0.05,   0.202];
A = sparse (A);
nnz (tril (A))
ans =  9
L = chol (A, "lower");
nnz (L)
ans =  10
norm (A - L * L', "fro") / norm (A, "fro")
ans =  1.1993e-16
opts.type = "nofill";
L = ichol (A, opts);
nnz (L)
ans =  9
norm (A - L * L', "fro") / norm (A, "fro")
ans =  0.019736
```

Another example for decomposition is a finite difference matrix used to solve a boundary value problem on the unit square.

```
nx = 400; ny = 200;
hx = 1 / (nx + 1); hy = 1 / (ny + 1);
Dxx = spdiags ([ones(nx, 1), -2*ones(nx, 1), ones(nx, 1)],
               [-1 0 1 ], nx, nx) / (hx ^ 2);
Dyy = spdiags ([ones(ny, 1), -2*ones(ny, 1), ones(ny, 1)],
               [-1 0 1 ], ny, ny) / (hy ^ 2);
A = -kron (Dxx, speye (ny)) - kron (speye (nx), Dyy);
nnz (tril (A))
ans =  239400
opts.type = "nofill";
L = ichol (A, opts);
nnz (tril (A))
ans =  239400
norm (A - L * L', "fro") / norm (A, "fro")
ans =  0.062327
```

References for implemented algorithms:

[1] Y. Saad, "Preconditioning Techniques", *Iterative Methods for Sparse Linear Systems*, PWS Publishing Company, 1996.

[2] M. Jones, P. Plassmann, *An Improved Incomplete Cholesky Factorization*, 1992.

**See also:** [chol], page 657, [ilu], page 764, [pcg], page 758.

*LUA* = ilu (*A*)
*LUA* = ilu (*A*, *opts*)

`[L, U] = ilu (...)`
`[L, U, P] = ilu (...)`

Compute the incomplete LU factorization of the sparse square matrix $A$.

`ilu` returns a unit lower triangular matrix $L$, an upper triangular matrix $U$, and optionally a permutation matrix $P$, such that `L*U` approximates `P*A`.

The factors given by this routine may be useful as preconditioners for a system of linear equations being solved by iterative methods such as BICG (BiConjugate Gradients) or GMRES (Generalized Minimum Residual Method).

The factorization may be modified by passing options in a structure *opts*. The option name is a field of the structure and the setting is the value of field. Names and specifiers are case sensitive.

`type` Type of factorization.

`"nofill"` (default)
ILU factorization with no fill-in (ILU(0)).
Additional supported options: `milu`.

`"crout"` Crout version of ILU factorization (ILUC).
Additional supported options: `milu`, `droptol`.

`"ilutp"` ILU factorization with threshold and pivoting.
Additional supported options: `milu`, `droptol`, `udiag`, `thresh`.

`droptol` A non-negative scalar specifying the drop tolerance for factorization. The default value is 0 which produces the complete LU factorization.

Non-diagonal entries of $U$ are set to 0 unless
`abs (U(i,j)) >= droptol * norm (A(:,j))`.

Non-diagonal entries of $L$ are set to 0 unless
`abs (L(i,j)) >= droptol * norm (A(:,j))/U(j,j)`.

`milu` Modified incomplete LU factorization:

`"row"` Row-sum modified incomplete LU factorization. The factorization preserves row sums: `A * e = L * U * e`, where e is a vector of ones.

`"col"` Column-sum modified incomplete LU factorization. The factorization preserves column sums: `e' * A = e' * L * U`.

`"off"` (default)
Row and column sums are not necessarily preserved.

`udiag` If true, any zeros on the diagonal of the upper triangular factor are replaced by the local drop tolerance `droptol * norm (A(:,j))/U(j,j)`. The default is false.

`thresh` Pivot threshold for factorization. It can range between 0 (diagonal pivoting) and 1 (default), where the maximum magnitude entry in the column is chosen to be the pivot.

If `ilu` is called with just one output, the returned matrix is `L + U - speye (size (A))`, where $L$ is unit lower triangular and $U$ is upper triangular.

With two outputs, `ilu` returns a unit lower triangular matrix $L$ and an upper triangular matrix $U$. For *opts*.type == `"ilutp"`, one of the factors is permuted based on the value of *opts*.milu. When *opts*.milu == `"row"`, $U$ is a column permuted upper triangular factor. Otherwise, $L$ is a row-permuted unit lower triangular factor.

If there are three named outputs and *opts*.milu != `"row"`, $P$ is returned such that $L$ and $U$ are incomplete factors of `P*A`. When *opts*.milu == `"row"`, $P$ is returned such that $L$ and $U$ are incomplete factors of `A*P`.

EXAMPLES

```
A = gallery ("neumann", 1600) + speye (1600);
opts.type = "nofill";
nnz (A)
ans = 7840

nnz (lu (A))
ans = 126478

nnz (ilu (A, opts))
ans = 7840
```

This shows that $A$ has 7,840 nonzeros, the complete LU factorization has 126,478 nonzeros, and the incomplete LU factorization, with 0 level of fill-in, has 7,840 nonzeros, the same amount as $A$. Taken from: https://www.mathworks.com/help/matlab/ref/ilu.html

```
A = gallery ("wathen", 10, 10);
b = sum (A, 2);
tol = 1e-8;
maxit = 50;
opts.type = "crout";
opts.droptol = 1e-4;
[L, U] = ilu (A, opts);
x = bicg (A, b, tol, maxit, L, U);
norm (A * x - b, inf)
```

This example uses ILU as preconditioner for a random FEM-Matrix, which has a large condition number. Without $L$ and $U$ BICG would not converge.

**See also:** [lu], page 659, [ichol], page 762, [bicg], page 676, [gmres], page 682.

## 22.4 Real Life Example using Sparse Matrices

A common application for sparse matrices is in the solution of Finite Element Models. Finite element models allow numerical solution of partial differential equations that do not have closed form solutions, typically because of the complex shape of the domain.

In order to motivate this application, we consider the boundary value Laplace equation. This system can model scalar potential fields, such as heat or electrical potential. Given a medium $\Omega$ with boundary $\partial\Omega$. At all points on the $\partial\Omega$ the boundary conditions are

known, and we wish to calculate the potential in Ω. Boundary conditions may specify the potential (Dirichlet boundary condition), its normal derivative across the boundary (Neumann boundary condition), or a weighted sum of the potential and its derivative (Cauchy boundary condition).

In a thermal model, we want to calculate the temperature in Ω and know the boundary temperature (Dirichlet condition) or heat flux (from which we can calculate the Neumann condition by dividing by the thermal conductivity at the boundary). Similarly, in an electrical model, we want to calculate the voltage in Ω and know the boundary voltage (Dirichlet) or current (Neumann condition after diving by the electrical conductivity). In an electrical model, it is common for much of the boundary to be electrically isolated; this is a Neumann boundary condition with the current equal to zero.

The simplest finite element models will divide Ω into simplexes (triangles in 2-D, pyramids in 3-D). We take as a 3-D example a cylindrical liquid filled tank with a small non-conductive ball from the EIDORS project[4]. This is model is designed to reflect an application of electrical impedance tomography, where current patterns are applied to such a tank in order to image the internal conductivity distribution. In order to describe the FEM geometry, we have a matrix of vertices `nodes` and simplices `elems`.

The following example creates a simple rectangular 2-D electrically conductive medium with 10 V and 20 V imposed on opposite sides (Dirichlet boundary conditions). All other edges are electrically isolated.

```
node_y = [1;1.2;1.5;1.8;2]*ones(1,11);
node_x = ones(5,1)*[1,1.05,1.1,1.2, ...
          1.3,1.5,1.7,1.8,1.9,1.95,2];
nodes = [node_x(:), node_y(:)];

[h,w] = size (node_x);
elems = [];
for idx = 1:w-1
  widx = (idx-1)*h;
  elems = [elems; ...
    widx+[(1:h-1);(2:h);h+(1:h-1)]'; ...
    widx+[(2:h);h+(2:h);h+(1:h-1)]' ];
endfor

E = size (elems,1);  # No. of simplices
N = size (nodes,1);  # No. of vertices
D = size (elems,2);  # dimensions+1
```

This creates a N-by-2 matrix `nodes` and a E-by-3 matrix `elems` with values, which define finite element triangles:

[4] EIDORS - Electrical Impedance Tomography and Diffuse optical Tomography Reconstruction Software http://eidors3d.sourceforge.net

```
nodes(1:7,:)'
  1.00 1.00 1.00 1.00 1.00 1.05 1.05 ...
  1.00 1.20 1.50 1.80 2.00 1.00 1.20 ...

elems(1:7,:)'
  1    2    3    4    2    3    4 ...
  2    3    4    5    7    8    9 ...
  6    7    8    9    6    7    8 ...
```

Using a first order FEM, we approximate the electrical conductivity distribution in $\Omega$ as constant on each simplex (represented by the vector `conductivity`). Based on the finite element geometry, we first calculate a system (or stiffness) matrix for each simplex (represented as 3-by-3 elements on the diagonal of the element-wise system matrix `SE`). Based on `SE` and a N-by-DE connectivity matrix `C`, representing the connections between simplices and vertices, the global connectivity matrix `S` is calculated.

```
## Element conductivity
conductivity = [1*ones(1,16), ...
       2*ones(1,48), 1*ones(1,16)];

## Connectivity matrix
C = sparse ((1:D*E), reshape (elems', ...
       D*E, 1), 1, D*E, N);

## Calculate system matrix
Siidx = floor ([0:D*E-1]'/D) * D * ...
       ones(1,D) + ones(D*E,1)*(1:D) ;
Sjidx = [1:D*E]'*ones (1,D);
Sdata = zeros (D*E,D);
dfact = factorial (D-1);
for j = 1:E
   a = inv ([ones(D,1), ...
       nodes(elems(j,:), :)]);
   const = conductivity(j) * 2 / ...
       dfact / abs (det (a));
   Sdata(D*(j-1)+(1:D),:) = const * ...
       a(2:D,:)' * a(2:D,:);
endfor
## Element-wise system matrix
SE = sparse(Siidx,Sjidx,Sdata);
## Global system matrix
S = C'* SE *C;
```

The system matrix acts like the conductivity $S$ in Ohm's law $SV = I$. Based on the Dirichlet and Neumann boundary conditions, we are able to solve for the voltages at each vertex `V`.

```
## Dirichlet boundary conditions
D_nodes = [1:5, 51:55];
D_value = [10*ones(1,5), 20*ones(1,5)];
```

```
V = zeros (N,1);
V(D_nodes) = D_value;
idx = 1:N; # vertices without Dirichlet
           # boundary condns
idx(D_nodes) = [];

## Neumann boundary conditions.  Note that
## N_value must be normalized by the
## boundary length and element conductivity
N_nodes = [];
N_value = [];

Q = zeros (N,1);
Q(N_nodes) = N_value;

V(idx) = S(idx,idx) \ ( Q(idx) - ...
          S(idx,D_nodes) * V(D_nodes));
```

Finally, in order to display the solution, we show each solved voltage value in the z-axis for each simplex vertex. See Figure 22.6.

```
elemx = elems(:,[1,2,3,1])';
xelems = reshape (nodes(elemx, 1), 4, E);
yelems = reshape (nodes(elemx, 2), 4, E);
velems = reshape (V(elemx), 4, E);
plot3 (xelems,yelems,velems,"k");
print "grid.eps";
```

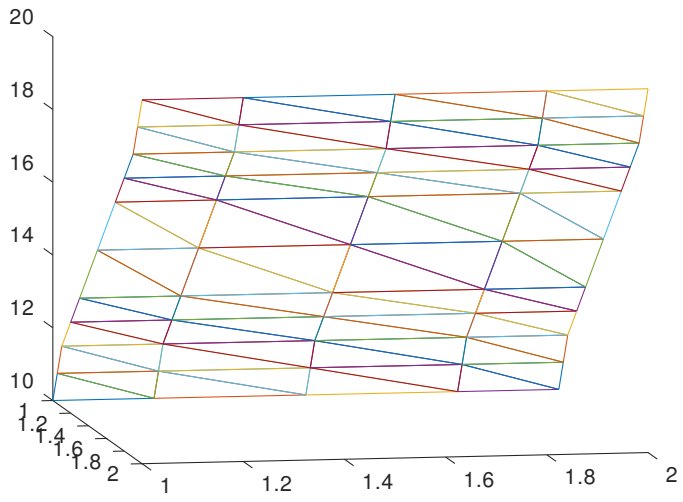



Figure 22.6: Example finite element model the showing triangular elements. The height of each vertex corresponds to the solution value.

# 23 Numerical Integration

Octave comes with several built-in functions for computing the integral of a function numerically (termed quadrature). These functions all solve 1-dimensional integration problems.

## 23.1 Functions of One Variable

Octave supports five different adaptive quadrature algorithms for computing the integral

$$\int_a^b f(x)dx$$

of a function $f$ over the interval from $a$ to $b$. These are

`quad` Numerical integration based on Gaussian quadrature.

`quadv` Numerical integration using an adaptive vectorized Simpson's rule.

`quadl` Numerical integration using an adaptive Lobatto rule.

`quadgk` Numerical integration using an adaptive Gauss-Konrod rule.

`quadcc` Numerical integration using adaptive Clenshaw-Curtis rules.

In addition, the following functions are also provided:

`integral` A compatibility wrapper function that will choose between `quadv` and `quadgk` depending on the integrand and options chosen.

`trapz, cumtrapz`
Numerical integration of data using the trapezoidal method.

The best quadrature algorithm to use depends on the integrand. If you have empirical data, rather than a function, the choice is `trapz` or `cumtrapz`. If you are uncertain about the characteristics of the integrand, `quadcc` will be the most robust as it can handle discontinuities, singularities, oscillatory functions, and infinite intervals. When the integrand is smooth `quadgk` may be the fastest of the algorithms.

| **Function** | **Characteristics** |
|---|---|
| quad | Low accuracy with nonsmooth integrands |
| quadv | Medium accuracy with smooth integrands |
| quadl | Medium accuracy with smooth integrands. Slower than quadgk. |
| quadgk | Medium accuracy (1e-6 – 1e-9) with smooth integrands.<br>Handles oscillatory functions and infinite bounds |
| quadcc | Low to High accuracy with nonsmooth/smooth integrands<br>Handles oscillatory functions, singularities, and infinite bounds |

Here is an example of using `quad` to integrate the function

$$f(x) = x \sin(1/x)\sqrt{|1 - x|}$$

from $x = 0$ to $x = 3$.

This is a fairly difficult integration (plot the function over the range of integration to see why).

The first step is to define the function:

```
function y = f (x)
  y = x .* sin (1./x) .* sqrt (abs (1 - x));
endfunction
```

Note the use of the 'dot' forms of the operators. This is not necessary for the `quad` integrator, but is required by the other integrators. In any case, it makes it much easier to generate a set of points for plotting because it is possible to call the function with a vector argument to produce a vector result.

The second step is to call quad with the limits of integration:

```
[q, ier, nfun, err] = quad ("f", 0, 3)
     ⇒ 1.9819
     ⇒ 1
     ⇒ 5061
     ⇒ 1.1522e-07
```

Although `quad` returns a nonzero value for *ier*, the result is reasonably accurate (to see why, examine what happens to the result if you move the lower bound to 0.1, then 0.01, then 0.001, etc.).

The function `"f"` can be the string name of a function or a function handle. These options make it quite easy to do integration without having to fully define a function in an m-file. For example:

```
# Verify gamma function = (n-1)! for n = 4
f = @(x) x.^3 .* exp (-x);
quadcc (f, 0, Inf)
     ⇒ 6.0000
```

`q = quad (f, a, b)`
`q = quad (f, a, b, tol)`
`q = quad (f, a, b, tol, sing)`
`[q, ier, nfev, err] = quad (...)`

Numerically evaluate the integral of $f$ from $a$ to $b$ using Fortran routines from QUADPACK.

$f$ is a function handle, inline function, or a string containing the name of the function to evaluate. The function must have the form `y = f (x)` where $y$ and $x$ are scalars.

$a$ and $b$ are the lower and upper limits of integration. Either or both may be infinite.

The optional argument *tol* is a vector that specifies the desired accuracy of the result. The first element of the vector is the desired absolute tolerance, and the second element is the desired relative tolerance. To choose a relative test only, set the absolute tolerance to zero. To choose an absolute test only, set the relative tolerance to zero. Both tolerances default to `sqrt (eps)` or approximately 1.5e-8.

The optional argument *sing* is a vector of values at which the integrand is known to be singular.

The result of the integration is returned in $q$.

*ier* contains an integer error code (0 indicates a successful integration).

*nfev* indicates the number of function evaluations that were made.

*err* contains an estimate of the error in the solution.

The function `quad_options` can set other optional parameters for `quad`.

Note: because `quad` is written in Fortran it cannot be called recursively. This prevents its use in integrating over more than one variable by routines `dblquad` and `triplequad`.

**See also:** [quad_options], page 773, [quadv], page 773, [quadl], page 774, [quadgk], page 775, [quadcc], page 776, [trapz], page 779, [dblquad], page 781, [triplequad], page 782.

`quad_options ()`
`val = quad_options (opt)`
`quad_options (opt, val)`

Query or set options for the function `quad`.

When called with no arguments, the names of all available options and their current values are displayed.

Given one argument, return the value of the option *opt*.

When called with two arguments, `quad_options` sets the option *opt* to value *val*.

Options include

`"absolute tolerance"`
: Absolute tolerance; may be zero for pure relative error test.

`"relative tolerance"`
: Non-negative relative tolerance. If the absolute tolerance is zero, the relative tolerance must be greater than or equal to `max (50*eps, 0.5e-28)`.

`"single precision absolute tolerance"`
: Absolute tolerance for single precision; may be zero for pure relative error test.

`"single precision relative tolerance"`
: Non-negative relative tolerance for single precision. If the absolute tolerance is zero, the relative tolerance must be greater than or equal to `max (50*eps, 0.5e-28)`.

`q = quadv (f, a, b)`
`q = quadv (f, a, b, tol)`
`q = quadv (f, a, b, tol, trace)`
`q = quadv (f, a, b, tol, trace, p1, p2, ...)`
`[q, nfev] = quadv (...)`

Numerically evaluate the integral of $f$ from $a$ to $b$ using an adaptive Simpson's rule.

$f$ is a function handle, inline function, or string containing the name of the function to evaluate. `quadv` is a vectorized version of `quad` and the function defined by $f$ must accept a scalar or vector as input and return a scalar, vector, or array as output.

$a$ and $b$ are the lower and upper limits of integration. Both limits must be finite.

The optional argument *tol* defines the absolute tolerance used to stop the adaptation procedure. The default value is 1e-6.

The algorithm used by `quadv` involves recursively subdividing the integration interval and applying Simpson's rule on each subinterval. If *trace* is true then after computing each of these partial integrals display: (1) the total number of function evaluations, (2) the left end of the subinterval, (3) the length of the subinterval, (4) the approximation of the integral over the subinterval.

Additional arguments *p1*, etc., are passed directly to the function *f*. To use default values for *tol* and *trace*, one may pass empty matrices ([]).

The result of the integration is returned in *q*.

The optional output *nfev* indicates the total number of function evaluations performed.

Note: `quadv` is written in Octave's scripting language and can be used recursively in `dblquad` and `triplequad`, unlike the `quad` function.

**See also:** [quad], page 772, [quadl], page 774, [quadgk], page 775, [quadcc], page 776, [trapz], page 779, [dblquad], page 781, [triplequad], page 782, [integral], page 778, [integral2], page 784, [integral3], page 786.

```
q = quadl (f, a, b)
q = quadl (f, a, b, tol)
q = quadl (f, a, b, tol, trace)
q = quadl (f, a, b, tol, trace, p1, p2, ...)
[q, nfev] = quadl (...)
```

Numerically evaluate the integral of *f* from *a* to *b* using an adaptive Lobatto rule.

*f* is a function handle, inline function, or string containing the name of the function to evaluate. The function *f* must be vectorized and return a vector of output values when given a vector of input values.

*a* and *b* are the lower and upper limits of integration. Both limits must be finite.

The optional argument *tol* defines the absolute tolerance with which to perform the integration. The default value is 1e-6.

The algorithm used by `quadl` involves recursively subdividing the integration interval. If *trace* is defined then for each subinterval display: (1) the total number of function evaluations, (2) the left end of the subinterval, (3) the length of the subinterval, (4) the approximation of the integral over the subinterval.

Additional arguments *p1*, etc., are passed directly to the function *f*. To use default values for *tol* and *trace*, one may pass empty matrices ([]).

The result of the integration is returned in *q*.

The optional output *nfev* indicates the total number of function evaluations performed.

Reference: W. Gander and W. Gautschi, *Adaptive Quadrature - Revisited*, BIT Vol. 40, No. 1, March 2000, pp. 84–101. https://www.inf.ethz.ch/personal/gander/

**See also:** [quad], page 772, [quadv], page 773, [quadgk], page 775, [quadcc], page 776, [trapz], page 779, [dblquad], page 781, [triplequad], page 782, [integral], page 778, [integral2], page 784, [integral3], page 786.

`q = quadgk (f, a, b)`
`q = quadgk (f, a, b, abstol)`
`q = quadgk (f, a, b, abstol, trace)`
`q = quadgk (f, a, b, "prop", val, ...)`
`[q, err] = quadgk (...)`

Numerically evaluate the integral of $f$ from $a$ to $b$ using adaptive Gauss-Kronrod quadrature.

$f$ is a function handle, inline function, or string containing the name of the function to evaluate. The function $f$ must be vectorized and return a vector of output values when given a vector of input values (See property `"ArrayValued"` for an exception to this rule).

$a$ and $b$ are the lower and upper limits of integration. Either or both limits may be infinite or contain weak end singularities. Variable transformation will be used to treat any infinite intervals and weaken the singularities. For example:

```
quadgk (@(x) 1 ./ (sqrt (x) .* (x + 1)), 0, Inf)
```

Note that the formulation of the integrand uses the element-by-element operator `./` and all user functions to `quadgk` should do the same.

The optional argument *abstol* defines the absolute tolerance used to stop the integration procedure. The default value is 1e-10 (1e-5 for single).

The algorithm used by `quadgk` involves subdividing the integration interval and evaluating each subinterval. If *trace* is true then after computing each of these partial integrals display: (1) the number of subintervals at this step, (2) the current estimate of the error *err*, (3) the current estimate for the integral $q$.

The behavior of the algorithm can be configured by passing arguments to `quadgk` as pairs `"prop", val`. Valid properties are

`AbsTol` Define the absolute error tolerance for the quadrature. The default absolute tolerance is 1e-10 (1e-5 for single).

`RelTol` Define the relative error tolerance for the quadrature. The default relative tolerance is 1e-6 (1e-4 for single).

`ArrayValued`
When set to true, the function $f$ produces an array output for a scalar input. The default is false which requires that $f$ produce an output that is the same size as the input. For example,

```
quadgk (@(x) x .^ (1:5), 0, 2, "ArrayValued", 1)
```

will integrate `[x.^1, x.^2, x.^3, x.^4, x.^5]` in one function call rather than having to repeatedly define a single anonymous function and use a normal invocation of `quadgk`.

`WayPoints`
Specify points which will become endpoints for subintervals in the algorithm which can result in significantly improved estimation of the error in the integral, faster computation, or both. It can be useful to specify more subintervals around a region where the integrand is rapidly changing or to flag locations where there is a discontinuity in the first derivative of

the function. For example, the signum function has a discontinuity at `x == 0` and by specifying a waypoint

```
quadgk (@(x) sign (x), -0.5, 1, "Waypoints", [0])
```

the error bound is reduced from 4e-7 to 1e-13.

If the function has **singularities** within the region of integration those should not be addressed with waypoints. Instead, the overall integral should be decomposed into a sum of several smaller integrals such that the singularity occurs as one of the bounds of integration in the call to `quadgk`.

If any of the waypoints are complex then contour integration is performed as documented below.

`MaxIntervalCount`

`quadgk` initially subdivides the interval on which to perform the quadrature into 10 intervals or, if WayPoints are given, at each waypoint. Subintervals that have an unacceptable error are subdivided and re-evaluated. If the number of subintervals exceeds 650 subintervals at any point then a poor convergence is signaled and the current estimate of the integral is returned. The property `"MaxIntervalCount"` can be used to alter the number of subintervals that can exist before exiting.

`Trace` If logically true `quadgk` prints information on the convergence of the quadrature at each iteration.

If any of *a*, *b*, or *waypoints* is complex then the quadrature is treated as a contour integral along a piecewise linear path defined by `[a, waypoints(1), waypoints(2), ..., b]`. In this case the integral is assumed to have no edge singularities. For example,

```
quadgk (@(z) log (z), 1+1i, 1+1i, "WayPoints",
        [-1+1i, -1-1i, +1-1i])
```

integrates `log (z)` along the square defined by `[1+1i, -1+1i, -1-1i, +1-1i]`.

The result of the integration is returned in *q*.

*err* is an approximate bound on the error in the integral `abs (q - I)`, where *I* is the exact value of the integral. If the adaptive integration did not converge, the value of *err* will be larger than the requested tolerance. If only a single output is requested then a warning will be emitted when the requested tolerance is not met. If the second output *err* is requested then no warning is issued and it is the responsibility of the programmer to inspect and determine whether the results are satisfactory.

Reference: L.F. Shampine, "Vectorized adaptive quadrature in MATLAB", *Journal of Computational and Applied Mathematics*, Vol. 211, Issue 2, pp. 131–140, Feb 2008.

**See also:** [quad], page 772, [quadv], page 773, [quadl], page 774, [quadcc], page 776, [trapz], page 779, [dblquad], page 781, [triplequad], page 782, [integral], page 778, [integral2], page 784, [integral3], page 786.

```
q = quadcc (f, a, b)
q = quadcc (f, a, b, tol)
```

`q = quadcc (f, a, b, tol, sing)`
`[q, err, nr_points] = quadcc (...)`

Numerically evaluate the integral of $f$ from $a$ to $b$ using doubly-adaptive Clenshaw-Curtis quadrature.

$f$ is a function handle, inline function, or string containing the name of the function to evaluate. The function $f$ must be vectorized and must return a vector of output values if given a vector of input values. For example,

```
f = @(x) x .* sin (1./x) .* sqrt (abs (1 - x));
```

which uses the element-by-element "dot" form for all operators.

$a$ and $b$ are the lower and upper limits of integration. Either or both limits may be infinite. `quadcc` handles an infinite limit by substituting the variable of integration with `x = tan (pi/2*u)`.

The optional argument *tol* is a 1- or 2-element vector that specifies the desired accuracy of the result. The first element of the vector is the desired absolute tolerance, and the second element is the desired relative tolerance. To choose a relative test only, set the absolute tolerance to zero. To choose an absolute test only, set the relative tolerance to zero. The default absolute tolerance is 1e-10 (1e-5 for single), and the default relative tolerance is 1e-6 (1e-4 for single).

The optional argument *sing* contains a list of points where the integrand has known singularities, or discontinuities in any of its derivatives, inside the integration interval. For the example above, which has a discontinuity at x=1, the call to `quadcc` would be as follows

```
int = quadcc (f, a, b, [], [ 1 ]);
```

The result of the integration is returned in $q$.

*err* is an estimate of the absolute integration error.

*nr_points* is the number of points at which the integrand was evaluated.

If the adaptive integration did not converge, the value of *err* will be larger than the requested tolerance. If only a single output is requested then a warning will be emitted when the requested tolerance is not met. If the second output *err* is requested then no warning is issued and it is the responsibility of the programmer to inspect and determine whether the results are satisfactory.

`quadcc` is capable of dealing with non-numeric values of the integrand such as `NaN` or `Inf`. If the integral diverges, and `quadcc` detects this, then a warning is issued and `Inf` or `-Inf` is returned.

Note: `quadcc` is a general purpose quadrature algorithm and, as such, may be less efficient for a smooth or otherwise well-behaved integrand than other methods such as `quadgk`.

The algorithm uses Clenshaw-Curtis quadrature rules of increasing degree in each interval and bisects the interval if either the function does not appear to be smooth or a rule of maximum degree has been reached. The error estimate is computed from the L2-norm of the difference between two successive interpolations of the integrand over the nodes of the respective quadrature rules.

Reference: P. Gonnet, "Increasing the Reliability of Adaptive Quadrature Using Explicit Interpolants", *ACM Transactions on Mathematical Software*, Vol. 37, Issue 3, Article No. 3, 2010.

**See also:** [quad], page 772, [quadv], page 773, [quadl], page 774, [quadgk], page 775, [trapz], page 779, [dblquad], page 781, [triplequad], page 782.

`q = integral (f, a, b)`
`q = integral (f, a, b, prop, val, ...)`
`[q, err] = integral (...)`

Numerically evaluate the integral of $f$ from $a$ to $b$ using adaptive quadrature.

`integral` is a wrapper for `quadcc` (general real-valued, scalar integrands and limits), and `quadgk` (integrals with specified integration paths and array-valued integrands) that is intended to provide MATLAB compatibility. More control of the numerical integration may be achievable by calling the various quadrature functions directly.

$f$ is a function handle, inline function, or string containing the name of the function to evaluate. The function $f$ must be vectorized and return a vector of output values when given a vector of input values.

$a$ and $b$ are the lower and upper limits of integration. Either or both limits may be infinite or contain weak end singularities. If either or both limits are complex, `integral` will perform a straight line path integral. Alternatively, a complex domain path can be specified using the `"Waypoints"` option (see below).

Additional optional parameters can be specified using `"property", value` pairs. Valid properties are:

`Waypoints`
: Specifies points to be used in defining subintervals of the quadrature algorithm, or if $a$, $b$, or *waypoints* are complex then the quadrature is calculated as a contour integral along a piecewise continuous path. For more detail, see [quadgk], page 775.

`ArrayValued`
: `integral` expects $f$ to return a scalar value unless *arrayvalued* is specified as true. This option will cause `integral` to perform the integration over the entire array and return $q$ with the same dimensions as returned by $f$. For more detail see [quadgk], page 775.

`AbsTol`
: Define the absolute error tolerance for the quadrature. The default absolute tolerance is 1e-10 (1e-5 for single).

`RelTol`
: Define the relative error tolerance for the quadrature. The default relative tolerance is 1e-6 (1e-4 for single).

The optional output *err* contains the absolute error estimate used by the called integrator.

Adaptive quadrature is used to minimize the estimate of error until the following is satisfied:

$$error \leq \max\left(AbsTol, RelTol \cdot |q|\right)$$

Known MATLAB incompatibilities:

1. If tolerances are left unspecified, and any integration limits or waypoints are of type `single`, then Octave's integral functions automatically reduce the default absolute and relative error tolerances as specified above. If tighter tolerances are desired they must be specified. MATLAB leaves the tighter tolerances appropriate for `double` inputs in place regardless of the class of the integration limits.

**See also:** [integral2], page 784, [integral3], page 786, [quad], page 772, [quadgk], page 775, [quadv], page 773, [quadl], page 774, [quadcc], page 776, [trapz], page 779, [dblquad], page 781, [triplequad], page 782.

Sometimes one does not have the function, but only the raw (x, y) points from which to perform an integration. This can occur when collecting data in an experiment. The `trapz` function can integrate these values as shown in the following example where "data" has been collected on the cosine function over the range [0, pi/2).

```
x = 0:0.1:pi/2;  # Uniformly spaced points
y = cos (x);
trapz (x, y)
     ⇒ 0.99666
```

The answer is reasonably close to the exact value of 1. Ordinary quadrature is sensitive to the characteristics of the integrand. Empirical integration depends not just on the integrand, but also on the particular points chosen to represent the function. Repeating the example above with the sine function over the range [0, pi/2) produces far inferior results.

```
x = 0:0.1:pi/2;  # Uniformly spaced points
y = sin (x);
trapz (x, y)
     ⇒ 0.92849
```

However, a slightly different choice of data points can change the result significantly. The same integration, with the same number of points, but spaced differently produces a more accurate answer.

```
x = linspace (0, pi/2, 16);  # Uniformly spaced, but including endpoint
y = sin (x);
trapz (x, y)
     ⇒ 0.99909
```

In general there may be no way of knowing the best distribution of points ahead of time. Or the points may come from an experiment where there is no freedom to select the best distribution. In any case, one must remain aware of this issue when using `trapz`.

`q = trapz (y)`
`q = trapz (x, y)`
`q = trapz (...,` *`dim`*`)`

Numerically evaluate the integral of points *y* using the trapezoidal method.

`trapz (y)` computes the integral of *y* along the first non-singleton dimension. When the argument *x* is omitted an equally spaced *x* vector with unit spacing (1) is assumed. `trapz (x, y)` evaluates the integral with respect to the spacing in *x* and the values in *y*. This is useful if the points in *y* have been sampled unevenly.

If the optional *dim* argument is given, operate along this dimension.

Application Note: If $x$ is not specified then unit spacing will be used. To scale the integral to the correct value you must multiply by the actual spacing value (deltaX). As an example, the integral of $x^3$ over the range [0, 1] is $x^4/4$ or 0.25. The following code uses `trapz` to calculate the integral in three different ways.

```
x = 0:0.1:1;
y = x.^3;
## No scaling
q = trapz (y)
  ⇒ q = 2.5250
## Approximation to integral by scaling
q * 0.1
  ⇒ 0.25250
## Same result by specifying x
trapz (x, y)
  ⇒ 0.25250
```

**See also:** [cumtrapz], page 780.

`q = cumtrapz (y)`
`q = cumtrapz (x, y)`
`q = cumtrapz (..., dim)`

Cumulative numerical integration of points $y$ using the trapezoidal method.

`cumtrapz (y)` computes the cumulative integral of $y$ along the first non-singleton dimension. Where `trapz` reports only the overall integral sum, `cumtrapz` reports the current partial sum value at each point of $y$.

When the argument $x$ is omitted an equally spaced $x$ vector with unit spacing (1) is assumed. `cumtrapz (x, y)` evaluates the integral with respect to the spacing in $x$ and the values in $y$. This is useful if the points in $y$ have been sampled unevenly.

If the optional *dim* argument is given, operate along this dimension.

Application Note: If $x$ is not specified then unit spacing will be used. To scale the integral to the correct value you must multiply by the actual spacing value (deltaX).

**See also:** [trapz], page 779, [cumsum], page 613.

## 23.2 Orthogonal Collocation

`[r, amat, bmat, q] = colloc (n, "left", "right")`

Compute derivative and integral weight matrices for orthogonal collocation.

Reference: J. Villadsen, M. L. Michelsen, *Solution of Differential Equation Models by Polynomial Approximation*.

Here is an example of using `colloc` to generate weight matrices for solving the second order differential equation $u' - \alpha u'' = 0$ with the boundary conditions $u(0) = 0$ and $u(1) = 1$.

First, we can generate the weight matrices for $n$ points (including the endpoints of the interval), and incorporate the boundary conditions in the right hand side (for a specific value of $\alpha$).

```
n = 7;
alpha = 0.1;
[r, a, b] = colloc (n-2, "left", "right");
at = a(2:n-1,2:n-1);
bt = b(2:n-1,2:n-1);
rhs = alpha * b(2:n-1,n) - a(2:n-1,n);
```

Then the solution at the roots $r$ is

```
u = [ 0; (at - alpha * bt) \ rhs; 1]
     ⇒ [ 0.00; 0.004; 0.01 0.00; 0.12; 0.62; 1.00 ]
```

## 23.3 Functions of Multiple Variables

Octave includes several functions for computing the integral of functions of multiple variables. This procedure can generally be performed by creating a function that integrates $f$ with respect to $x$, and then integrates that function with respect to $y$. This procedure can be performed manually using the following example which integrates the function:

$$f(x, y) = \sin(\pi x y)\sqrt{xy}$$

for $x$ and $y$ between 0 and 1.

Using `quadgk` in the example below, a double integration can be performed. (Note that any of the 1-D quadrature functions can be used in this fashion except for `quad` since it is written in Fortran and cannot be called recursively.)

```
function q = g(y)
  q = ones (size (y));
  for i = 1:length (y)
    f = @(x) sin (pi*x.*y(i)) .* sqrt (x.*y(i));
    q(i) = quadgk (f, 0, 1);
  endfor
endfunction

I = quadgk ("g", 0, 1)
      ⇒ 0.30022
```

The algorithm above is implemented in the function `dblquad` for integrals over two variables. The 3-D equivalent of this process is implemented in `triplequad` for integrals over three variables. As an example, the result above can be replicated with a call to `dblquad` as shown below.

```
I = dblquad (@(x, y) sin (pi*x.*y) .* sqrt (x.*y), 0, 1, 0, 1)
      ⇒ 0.30022
```

*q* = dblquad (*f*, *xa*, *xb*, *ya*, *yb*)
*q* = dblquad (*f*, *xa*, *xb*, *ya*, *yb*, *tol*)
*q* = dblquad (*f*, *xa*, *xb*, *ya*, *yb*, *tol*, *quadf*)
*q* = dblquad (*f*, *xa*, *xb*, *ya*, *yb*, *tol*, *quadf*, ...)

Numerically evaluate the double integral of *f*.

*f* is a function handle, inline function, or string containing the name of the function to evaluate. The function *f* must have the form $z = f(x, y)$ where *x* is a vector and *y* is a scalar. It should return a vector of the same length and orientation as *x*.

*xa*, *ya* and *xb*, *yb* are the lower and upper limits of integration for x and y respectively. The underlying integrator determines whether infinite bounds are accepted.

The optional argument *tol* defines the absolute tolerance used to integrate each sub-integral. The default value is 1e-6.

The optional argument *quadf* specifies which underlying integrator function to use. Any choice but `quad` is available and the default is `quadcc`.

Additional arguments, are passed directly to *f*. To use the default value for *tol* or *quadf* one may pass `':'` or an empty matrix ([]).

**See also:** [integral2], page 784, [integral3], page 786, [triplequad], page 782, [quad], page 772, [quadv], page 773, [quadl], page 774, [quadgk], page 775, [quadcc], page 776, [trapz], page 779.

`q = triplequad (f, xa, xb, ya, yb, za, zb)`
`q = triplequad (f, xa, xb, ya, yb, za, zb, tol)`
`q = triplequad (f, xa, xb, ya, yb, za, zb, tol, quadf)`
`q = triplequad (f, xa, xb, ya, yb, za, zb, tol, quadf, ...)`

Numerically evaluate the triple integral of *f*.

*f* is a function handle, inline function, or string containing the name of the function to evaluate. The function *f* must have the form $w = f(x, y, z)$ where either *x* or *y* is a vector and the remaining inputs are scalars. It should return a vector of the same length and orientation as *x* or *y*.

*xa*, *ya*, *za* and *xb*, *yb*, *zb* are the lower and upper limits of integration for x, y, and z respectively. The underlying integrator determines whether infinite bounds are accepted.

The optional argument *tol* defines the absolute tolerance used to integrate each sub-integral. The default value is 1e-6.

The optional argument *quadf* specifies which underlying integrator function to use. Any choice but `quad` is available and the default is `quadcc`.

Additional arguments, are passed directly to *f*. To use the default value for *tol* or *quadf* one may pass `':'` or an empty matrix ([]).

**See also:** [integral3], page 786, [integral2], page 784, [dblquad], page 781, [quad], page 772, [quadv], page 773, [quadl], page 774, [quadgk], page 775, [quadcc], page 776, [trapz], page 779.

The recursive algorithm for quadrature presented above is referred to as `"iterated"`. A separate 2-D integration method is implemented in the function `quad2d`. This function performs a `"tiled"` integration by subdividing the integration domain into rectangular regions and performing separate integrations over those domains. The domains are further subdivided in areas requiring refinement to reach the desired numerical accuracy. For certain functions this method can be faster than the 2-D iteration used in the other functions above.

`q = quad2d (f, xa, xb, ya, yb)`
`q = quad2d (f, xa, xb, ya, yb, prop, val, ...)`
`[q, err, iter] = quad2d (...)`

Numerically evaluate the two-dimensional integral of $f$ using adaptive quadrature over the two-dimensional domain defined by $xa$, $xb$, $ya$, $yb$ using tiled integration. Additionally, $ya$ and $yb$ may be scalar functions of $x$, allowing for the integration over non-rectangular domains.

$f$ is a function handle, inline function, or string containing the name of the function to evaluate. The function $f$ must be of the form $z = f(x, y)$, and all operations must be vectorized such that $x$ and $y$ accept array inputs and return array outputs of the same size. (It can be assumed that $x$ and $y$ will either be same-size arrays or one will be a scalar.) The underlying integrators will input arrays of integration points into $f$ and/or use internal vector expansions to speed computation that can produce unpredictable results if $f$ is not restricted to elementwise operations. For integrands where this is unavoidable, the `("Vectorized") option described below may produce`more reliable results.

Additional optional parameters can be specified using `"property", value` pairs. Valid properties are:

`AbsTol` Define the absolute error tolerance for the quadrature. The default value is 1e-10 (1e-5 for single).

`RelTol` Define the relative error tolerance for the quadrature. The default value is 1e-6 (1e-4 for single).

`MaxFunEvals`
The maximum number of function calls to the vectorized function *f*. The default value is 5000.

`Singular` Enable/disable transforms to weaken singularities on the edge of the integration domain. The default value is *true*.

`Vectorized`
Enable or disable vectorized integration. A value of `false` forces Octave to use only scalar inputs when calling the integrand, which enables integrands $f(x, y)$ that have not been vectorized or only accept scalar values of $x$ or $y$. The default value is `true`. Note that this is achieved by wrapping $f(x, y)$ with the function `arrayfun`, which may significantly decrease computation speed.

`FailurePlot`
If `quad2d` fails to converge to the desired error tolerance before MaxFunEvals is reached, a plot of the areas that still need refinement is created. The default value is *false*.

Adaptive quadrature is used to minimize the estimate of error until the following is satisfied:

$$error \leq \max\left(AbsTol, RelTol \cdot |q|\right)$$

The optional output *err* is an approximate bound on the error in the integral `abs (q - I)`, where $I$ is the exact value of the integral. The optional output *iter* is the number of vectorized function calls to the function $f$ that were used.

Example 1 : integrate a rectangular region in x-y plane

```
f = @(x,y) 2*ones (size (x));
q = quad2d (f, 0, 1, 0, 1)
  ⇒ q =  2
```

The result is a volume, which for this constant-value integrand, is just *Length* * *Width* * *Height*.

Example 2 : integrate a triangular region in x-y plane

```
f = @(x,y) 2*ones (size (x));
ymax = @(x) 1 - x;
q = quad2d (f, 0, 1, 0, ymax)
  ⇒ q =  1
```

The result is a volume, which for this constant-value integrand $f = 2$, is the Triangle Area x Height or `1/2 * Base * Width * Height`.

Example 3 : integrate a non-vectorized function over a square region

```
f = @(x,y) sinc (x) * sinc (y));
q = quad2d (f, -1, 1, -1, 1)
  ⇒ q =  12.328  (incorrect)
q = quad2d (f, -1, 1, -1, 1, "Vectorized", false)
  ⇒ q =  1.390 (correct)
f = @(x,y) sinc (x) .* sinc (y);
q = quad2d (f, -1, 1, -1, 1)
  ⇒ q =  1.390  (correct)
```

The first result is incorrect as the non-elementwise operator between the sinc functions in $f$ create unintended matrix multiplications between the internal integration arrays used by `quad2d`. In the second result, setting `"Vectorized"` to false forces `quad2d` to perform scalar internal operations to compute the integral, resulting in the correct numerical result at the cost of about a 20x increase in computation time. In the third result, vectorizing the integrand $f$ using the elementwise multiplication operator gets the correct result without increasing computation time.

Programming Notes: If there are singularities within the integration region it is best to split the integral and place the singularities on the boundary.

Known MATLAB incompatibility: If tolerances are left unspecified, and any integration limits are of type `single`, then Octave's integral functions automatically reduce the default absolute and relative error tolerances as specified above. If tighter tolerances are desired they must be specified. MATLAB leaves the tighter tolerances appropriate for `double` inputs in place regardless of the class of the integration limits.

Reference: L.F. Shampine, "MATLAB program for quadrature in 2D", *Applied Mathematics and Computation*, Vol. 1, pp. 266–274, 2008.

**See also:** [integral2], page 784, [dblquad], page 781, [integral], page 778, [quad], page 772, [quadgk], page 775, [quadv], page 773, [quadl], page 774, [quadcc], page 776, [trapz], page 779, [integral3], page 786, [triplequad], page 782.

Finally, the functions `integral2` and `integral3` are provided as general 2-D and 3-D integration functions. They will auto-select between iterated and tiled integration methods and, unlike `dblquad` and `triplequad`, will work with non-rectangular integration domains.

`q = integral2 (f, xa, xb, ya, yb)`
`q = integral2 (f, xa, xb, ya, yb, prop, val, ...)`
`[q, err] = integral2 (...)`

Numerically evaluate the two-dimensional integral of $f$ using adaptive quadrature over the two-dimensional domain defined by $xa$, $xb$, $ya$, $yb$ (scalars may be finite or infinite). Additionally, $ya$ and $yb$ may be scalar functions of $x$, allowing for integration over non-rectangular domains.

$f$ is a function handle, inline function, or string containing the name of the function to evaluate. The function $f$ must be of the form $z = f(x, y)$, and all operations must be vectorized such that $x$ and $y$ accept array inputs and return array outputs of the same size. (It can be assumed that $x$ and $y$ will either be same-size arrays or one will be a scalar.) The underlying integrators will input arrays of integration points into $f$ and/or use internal vector expansions to speed computation that can produce unpredictable results if $f$ is not restricted to elementwise operations. For integrands where this is unavoidable, the `("Vectorized") option described below may produce`more reliable results.

Additional optional parameters can be specified using `"property", value` pairs. Valid properties are:

`AbsTol` Define the absolute error tolerance for the quadrature. The default value is 1e-10 (1e-5 for single).

`RelTol` Define the relative error tolerance for the quadrature. The default value is 1e-6 (1e-4 for single).

`Method` Specify the two-dimensional integration method to be used, with valid options being `"auto"` (default), `"tiled"`, or `"iterated"`. When using `"auto"`, Octave will choose the `"tiled"` method unless any of the integration limits are infinite.

`Vectorized`
Enable or disable vectorized integration. A value of `false` forces Octave to use only scalar inputs when calling the integrand, which enables integrands $f(x, y)$ that have not been vectorized or only accept scalar values of $x$ or $y$. The default value is `true`. Note that this is achieved by wrapping $f(x, y)$ with the function `arrayfun`, which may significantly decrease computation speed.

Adaptive quadrature is used to minimize the estimate of error until the following is satisfied:

$$error \leq \max\left(AbsTol, RelTol \cdot |q|\right)$$

*err* is an approximate bound on the error in the integral `abs (q - I)`, where $I$ is the exact value of the integral.

Example 1 : integrate a rectangular region in x-y plane

```
f = @(x,y) 2*ones (size (x));
q = integral2 (f, 0, 1, 0, 1)
  ⇒ q =  2
```

The result is a volume, which for this constant-value integrand, is just *Length* * *Width* * *Height*.

Example 2 : integrate a triangular region in x-y plane

```
f = @(x,y) 2*ones (size (x));
ymax = @(x) 1 - x;
q = integral2 (f, 0, 1, 0, ymax)
  ⇒ q =  1
```

The result is a volume, which for this constant-value integrand $f = 2$, is the Triangle Area x Height or `1/2 * Base * Width * Height`.

Example 3 : integrate a non-vectorized function over a square region

```
f = @(x,y) sinc (x) * sinc (y));
q = integral2 (f, -1, 1, -1, 1)
  ⇒ q =  12.328  (incorrect)
q = integral2 (f, -1, 1, -1, 1, "Vectorized", false)
  ⇒ q =  1.390 (correct)
f = @(x,y) sinc (x) .* sinc (y);
q = integral2 (f, -1, 1, -1, 1)
  ⇒ q =  1.390  (correct)
```

The first result is incorrect as the non-elementwise operator between the sinc functions in *f* create unintended matrix multiplications between the internal integration arrays used by `integral2`. In the second result, setting `"Vectorized"` to false forces `integral2` to perform scalar internal operations to compute the integral, resulting in the correct numerical result at the cost of about a 20x increase in computation time. In the third result, vectorizing the integrand *f* using the elementwise multiplication operator gets the correct result without increasing computation time.

Programming Notes: If there are singularities within the integration region it is best to split the integral and place the singularities on the boundary.

Known MATLAB incompatibility: If tolerances are left unspecified, and any integration limits are of type `single`, then Octave's integral functions automatically reduce the default absolute and relative error tolerances as specified above. If tighter tolerances are desired they must be specified. MATLAB leaves the tighter tolerances appropriate for `double` inputs in place regardless of the class of the integration limits.

Reference: L.F. Shampine, "MATLAB program for quadrature in 2D", *Applied Mathematics and Computation*, Vol. 1, pp. 266–274, 2008.

**See also:** [quad2d], page 782, [dblquad], page 781, [integral], page 778, [quad], page 772, [quadgk], page 775, [quadv], page 773, [quadl], page 774, [quadcc], page 776, [trapz], page 779, [integral3], page 786, [triplequad], page 782.

`q = integral3 (f, xa, xb, ya, yb, za, zb)`
`q = integral3 (f, xa, xb, ya, yb, za, zb, prop, val, ...)`

Numerically evaluate the three-dimensional integral of *f* using adaptive quadrature over the three-dimensional domain defined by *xa*, *xb*, *ya*, *yb*, *za*, *zb* (scalars may be finite or infinite). Additionally, *ya* and *yb* may be scalar functions of *x* and *za*, and *zb* maybe be scalar functions of *x* and *y*, allowing for integration over non-rectangular domains.

$f$ is a function handle, inline function, or string containing the name of the function to evaluate. The function $f$ must be of the form $z = f(x, y, z)$, and all operations must be vectorized such that $x$, $y$, and $z$ accept array inputs and return array outputs of the same size. (It can be assumed that $x$, $y$, and $z$ will either be same-size arrays or scalars.) The underlying integrators will input arrays of integration points into $f$ and/or use internal vector expansions to speed computation that can produce unpredictable results if $f$ is not restricted to elementwise operations. For integrands where this is unavoidable, the `("Vectorized") option`described below may produce more reliable results.

Additional optional parameters can be specified using `"property", value` pairs. Valid properties are:

`AbsTol` Define the absolute error tolerance for the quadrature. The default value is 1e-10 (1e-5 for single).

`RelTol` Define the relative error tolerance for the quadrature. The default value is 1e-6 (1e-4 for single).

`Method` Specify the two-dimensional integration method to be used, with valid options being `"auto"` (default), `"tiled"`, or `"iterated"`. When using `"auto"`, Octave will choose the `"tiled"` method unless any of the integration limits are infinite.

`Vectorized`
Enable or disable vectorized integration. A value of `false` forces Octave to use only scalar inputs when calling the integrand, which enables integrands $f(x, y, z)$ that have not been vectorized or only accept scalar values of $x$, $y$, or $z$. The default value is `true`. Note that this is achieved by wrapping $f(x, y, z)$ with the function `arrayfun`, which may significantly decrease computation speed.

Adaptive quadrature is used to minimize the estimate of error until the following is satisfied:

$$error \leq \max\left(AbsTol, RelTol \cdot |q|\right)$$

*err* is an approximate bound on the error in the integral `abs (q - I)`, where $I$ is the exact value of the integral.

Example 1 : integrate over a rectangular volume

```
f = @(x,y,z) ones (size (x));
q = integral3 (f, 0, 1, 0, 1, 0, 1)
  ⇒ q =  1.00000
```

For this constant-value integrand, the result is a volume which is just *`Length * Width * Height`*.

Example 2 : integrate over a spherical volume

```
f = @(x,y) ones (size (x));
ymax = @(x) sqrt (1 - x.^2);
zmax = @(x,y) sqrt (1 - x.^2 - y.^2);
q = integral3 (f, 0, 1, 0, ymax, 0, zmax)
  ⇒ q =  0.52360
```

For this constant-value integrand, the result is a volume which is 1/8th of a unit sphere or `1/8 * 4/3 * pi`.

Example 3 : integrate a non-vectorized function over a cubic volume

```
f = @(x,y) sinc (x) * sinc (y), * sinc (z);
q = integral3 (f, -1, 1, -1, 1, -1, 1)
  ⇒ q =  14.535  (incorrect)
q = integral3 (f, -1, 1, -1, 1, -1, 1, "Vectorized", false)
  ⇒ q =  1.6388 (correct)
f = @(x,y,z) sinc (x) .* sinc (y), .* sinc (z);
q = integral3 (f, -1, 1, -1, 1, -1, 1)
  ⇒ q =  1.6388  (correct)
```

The first result is incorrect as the non-elementwise operator between the sinc functions in $f$ create unintended matrix multiplications between the internal integration arrays used by `integral3`. In the second result, setting `"Vectorized"` to false forces `integral3` to perform scalar internal operations to compute the integral, resulting in the correct numerical result at the cost of about a 30x increase in computation time. In the third result, vectorizing the integrand $f$ using the elementwise multiplication operator gets the correct result without increasing computation time.

Programming Notes: If there are singularities within the integration region it is best to split the integral and place the singularities on the boundary.

Known MATLAB incompatibility: If tolerances are left unspecified, and any integration limits are of type `single`, then Octave's integral functions automatically reduce the default absolute and relative error tolerances as specified above. If tighter tolerances are desired they must be specified. MATLAB leaves the tighter tolerances appropriate for `double` inputs in place regardless of the class of the integration limits.

Reference: L.F. Shampine, "MATLAB program for quadrature in 2D", *Applied Mathematics and Computation*, Vol. 1, pp. 266–274, 2008.

**See also:** [triplequad], page 782, [integral], page 778, [quad], page 772, [quadgk], page 775, [quadv], page 773, [quadl], page 774, [quadcc], page 776, [trapz], page 779, [integral2], page 784, [quad2d], page 782, [dblquad], page 781.

The above integrations can be fairly slow, and that problem increases exponentially with the dimensionality of the integral. Another possible solution for 2-D integration is to use Orthogonal Collocation as described in the previous section (see Section 23.2 [Orthogonal Collocation], page 780). The integral of a function $f(x, y)$ for $x$ and $y$ between 0 and 1 can be approximated using $n$ points by

$$\int_0^1 \int_0^1 f(x,y)dxdy \approx \sum_{i=1}^{n} \sum_{j=1}^{n} q_i q_j f(r_i, r_j),$$

where $q$ and $r$ is as returned by `colloc (n)`. The generalization to more than two variables is straight forward. The following code computes the studied integral using $n = 8$ points.

```
f = @(x,y) sin (pi*x*y') .* sqrt (x*y');
n = 8;
[t, ~, ~, q] = colloc (n);
I = q'*f(t,t)*q;
      ⇒ 0.30022
```

It should be noted that the number of points determines the quality of the approximation. If the integration needs to be performed between $a$ and $b$, instead of 0 and 1, then a change of variables is needed.

# 24 Differential Equations

Octave has built-in functions for solving ordinary differential equations (ODEs), and differential-algebraic equations (DAEs).

## 24.1 Ordinary Differential Equations

The function `lsode` can be used to solve ODEs of the form

$$\frac{dx}{dt} = f(x, t)$$

using Hindmarsh's ODE solver LSODE.

`[x, istate, msg] = lsode (fcn, x_0, t)`
`[x, istate, msg] = lsode (fcn, x_0, t, t_crit)`

Ordinary Differential Equation (ODE) solver.

The set of differential equations to solve is

$$\frac{dx}{dt} = f(x, t)$$

with

$$x(t_0) = x_0$$

The solution is returned in the matrix *x*, with each row corresponding to an element of the vector *t*. The first element of *t* should be $t_0$ and should correspond to the initial state of the system *x_0*, so that the first row of the output is *x_0*.

The first argument, *fcn*, is a string, inline, or function handle that names the function $f$ to call to compute the vector of right hand sides for the set of equations. The function must have the form

```
xdot = f (x, t)
```

in which *xdot* and *x* are vectors and *t* is a scalar.

If *fcn* is a two-element string array or a two-element cell array of strings, inline functions, or function handles, the first element names the function $f$ described above, and the second element names a function to compute the Jacobian of $f$. The Jacobian function must have the form

```
jac = j (x, t)
```

in which *jac* is the matrix of partial derivatives

$$J = \frac{\partial f_i}{\partial x_j} = \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \frac{\partial f_1}{\partial x_2} & \cdots & \frac{\partial f_1}{\partial x_N} \\ \frac{\partial f_2}{\partial x_1} & \frac{\partial f_2}{\partial x_2} & \cdots & \frac{\partial f_2}{\partial x_N} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial f_M}{\partial x_1} & \frac{\partial f_M}{\partial x_2} & \cdots & \frac{\partial f_M}{\partial x_N} \end{bmatrix}$$

The second argument specifies the initial state of the system $x_0$. The third argument is a vector, *t*, specifying the time values for which a solution is sought.

The fourth argument is optional, and may be used to specify a set of times that the ODE solver should not integrate past. It is useful for avoiding difficulties with singularities and points where there is a discontinuity in the derivative.

After a successful computation, the value of *istate* will be 2 (consistent with the Fortran version of LSODE).

If the computation is not successful, *istate* will be something other than 2 and *msg* will contain additional information.

You can use the function `lsode_options` to set optional parameters for `lsode`.

See Alan C. Hindmarsh, "ODEPACK, A Systematized Collection of ODE Solvers", *Scientific Computing*, R. S. Stepleman, editor, 1983. or https://computing.llnl.gov/projects/odepack for more information about the inner workings of `lsode`.

Example: Solve the Van der Pol equation

```
fvdp = @(y,t) [y(2); (1 - y(1)^2) * y(2) - y(1)];
t = linspace (0, 20, 100);
y = lsode (fvdp, [2; 0], t);
```

**See also:** [daspk], page 794, [dassl], page 798, [dasrt], page 800.

`lsode_options ()`
`val = lsode_options (opt)`
`lsode_options (opt, val)`

Query or set options for the function `lsode`.

When called with no arguments, the names of all available options and their current values are displayed.

Given one argument, return the value of the option *opt*.

When called with two arguments, `lsode_options` sets the option *opt* to value *val*.

Options include

`"absolute tolerance"`
: Absolute tolerance. May be either vector or scalar. If a vector, it must match the dimension of the state vector.

`"relative tolerance"`
: Relative tolerance parameter. Unlike the absolute tolerance, this parameter may only be a scalar.

  The local error test applied at each integration step is

  ```
  abs (local error in x(i)) <= ...
      rtol * abs (y(i)) + atol(i)
  ```

`"integration method"`
: A string specifying the method of integration to use to solve the ODE system. Valid values are

  `"adams"`
  `"non-stiff"`
  : No Jacobian used (even if it is available).

`"bdf"`

`"stiff"` Use stiff backward differentiation formula (BDF) method. If a function to compute the Jacobian is not supplied, `lsode` will compute a finite difference approximation of the Jacobian matrix.

`"initial step size"`
The step size to be attempted on the first step (default is determined automatically).

`"maximum order"`
Restrict the maximum order of the solution method. If using the Adams method, this option must be between 1 and 12. Otherwise, it must be between 1 and 5, inclusive.

`"maximum step size"`
Setting the maximum stepsize will avoid passing over very large regions (default is not specified).

`"minimum step size"`
The minimum absolute step size allowed (default is 0).

`"step limit"`
Maximum number of steps allowed (default is 100000).

`"jacobian type"`
A string specifying the type of Jacobian used with the stiff backward differentiation formula (BDF) integration method. Valid values are

`"full"` The default. All partial derivatives are approximated or used from the user-supplied Jacobian function.

`"banded"` Only the diagonal and the number of lower and upper subdiagonals specified by the options `"lower jacobian subdiagonals"` and `"upper jacobian subdiagonals"`, respectively, are approximated or used from the user-supplied Jacobian function. A user-supplied Jacobian function may set all other partial derivatives to arbitrary values.

`"diagonal"`
If a Jacobian function is supplied by the user, this setting has no effect. A Jacobian approximated by `lsode` is restricted to the diagonal, where each partial derivative is computed by applying a finite change to all elements of the state together; if the real Jacobian is indeed always diagonal, this has the same effect as applying the finite change only to the respective element of the state, but is more efficient.

`"lower jacobian subdiagonals"`
Number of lower subdiagonals used if option `"jacobian type"` is set to `"banded"`. The default is zero.

`"upper jacobian subdiagonals"`
: Number of upper subdiagonals used if option `"jacobian type"` is set to `"banded"`. The default is zero.

Here is an example of solving a set of three differential equations using `lsode`. Given the function

```
## oregonator differential equation
function xdot = f (x, t)

  xdot = zeros (3,1);

  xdot(1) = 77.27 * (x(2) - x(1)*x(2) + x(1) ...
            - 8.375e-06*x(1)^2);
  xdot(2) = (x(3) - x(1)*x(2) - x(2)) / 77.27;
  xdot(3) = 0.161*(x(1) - x(3));

endfunction
```

and the initial condition `x0 = [ 4; 1.1; 4 ]`, the set of equations can be integrated using the command

```
t = linspace (0, 500, 1000);

y = lsode ("f", x0, t);
```

If you try this, you will see that the value of the result changes dramatically between $t = 0$ and 5, and again around $t = 305$. A more efficient set of output points might be

```
t = [0, logspace(-1, log10(303), 150), ...
       logspace(log10(304), log10(500), 150)];
```

An m-file for the differential equation used above is included with the Octave distribution in the examples directory under the name `oregonator.m`.

## 24.2 Differential-Algebraic Equations

The function `daspk` can be used to solve DAEs of the form

$$0 = f(\dot{x}, x, t), \qquad x(t=0) = x_0, \dot{x}(t=0) = \dot{x}_0$$

where $\dot{x} = \frac{dx}{dt}$ is the derivative of $x$. The equation is solved using Petzold's DAE solver DASPK.

`[x, xdot, istate, msg] = daspk (fcn, x_0, xdot_0, t, t_crit)`
: Solve a set of differential-algebraic equations.

  `daspk` solves the set of equations

$$0 = f(x, \dot{x}, t)$$

  with

$$x(t_0) = x_0, \dot{x}(t_0) = \dot{x}_0$$

The solution is returned in the matrices *x* and *xdot*, with each row in the result matrices corresponding to one of the elements in the vector *t*. The first element of *t* should be $t_0$ and correspond to the initial state of the system *x_0* and its derivative *xdot_0*, so that the first row of the output *x* is *x_0* and the first row of the output *xdot* is *xdot_0*.

The first argument, *fcn*, is a string, inline, or function handle that names the function *f* to call to compute the vector of residuals for the set of equations. It must have the form

```
res = f (x, xdot, t)
```

in which *x*, *xdot*, and *res* are vectors, and *t* is a scalar.

If *fcn* is a two-element string array or a two-element cell array of strings, inline functions, or function handles, the first element names the function *f* described above, and the second element names a function to compute the modified Jacobian

$$J = \frac{\partial f}{\partial x} + c\frac{\partial f}{\partial \dot{x}}$$

The modified Jacobian function must have the form

```
jac = j (x, xdot, t, c)
```

The second and third arguments to `daspk` specify the initial condition of the states and their derivatives, and the fourth argument specifies a vector of output times at which the solution is desired, including the time corresponding to the initial condition.

The set of initial states and derivatives are not strictly required to be consistent. If they are not consistent, you must use the `daspk_options` function to provide additional information so that `daspk` can compute a consistent starting point.

The fifth argument is optional, and may be used to specify a set of times that the DAE solver should not integrate past. It is useful for avoiding difficulties with singularities and points where there is a discontinuity in the derivative.

After a successful computation, the value of *istate* will be greater than zero (consistent with the Fortran version of DASPK).

If the computation is not successful, the value of *istate* will be less than zero and *msg* will contain additional information.

You can use the function `daspk_options` to set optional parameters for `daspk`.

**See also:** [dassl], page 798.

```
daspk_options ()
val = daspk_options (opt)
daspk_options (opt, val)
```

Query or set options for the function `daspk`.

When called with no arguments, the names of all available options and their current values are displayed.

Given one argument, return the value of the option *opt*.

When called with two arguments, `daspk_options` sets the option *opt* to value *val*.

Options include

`"absolute tolerance"`
: Absolute tolerance. May be either vector or scalar. If a vector, it must match the dimension of the state vector, and the relative tolerance must also be a vector of the same length.

`"relative tolerance"`
: Relative tolerance. May be either vector or scalar. If a vector, it must match the dimension of the state vector, and the absolute tolerance must also be a vector of the same length.

  The local error test applied at each integration step is

```
abs (local error in x(i))
     <= rtol(i) * abs (Y(i)) + atol(i)
```

`"compute consistent initial condition"`
: Denoting the differential variables in the state vector by '`Y_d`' and the algebraic variables by '`Y_a`', `ddaspk` can solve one of two initialization problems:

  1. Given Y_d, calculate Y_a and Y'_d
  2. Given Y', calculate Y.

  In either case, initial values for the given components are input, and initial guesses for the unknown components must also be provided as input. Set this option to 1 to solve the first problem, or 2 to solve the second (the default is 0, so you must provide a set of initial conditions that are consistent).

  If this option is set to a nonzero value, you must also set the `"algebraic variables"` option to declare which variables in the problem are algebraic.

`"use initial condition heuristics"`
: Set to a nonzero value to use the initial condition heuristics options described below.

`"initial condition heuristics"`
: A vector of the following parameters that can be used to control the initial condition calculation.

  `MXNIT` Maximum number of Newton iterations (default is 5).

  `MXNJ` Maximum number of Jacobian evaluations (default is 6).

  `MXNH` Maximum number of values of the artificial stepsize parameter to be tried if the `"compute consistent initial condition"` option has been set to 1 (default is 5).

  Note that the maximum total number of Newton iterations allowed is `MXNIT*MXNJ*MXNH` if the `"compute consistent initial condition"` option has been set to 1 and `MXNIT*MXNJ` if it is set to 2.

`LSOFF` Set to a nonzero value to disable the linesearch algorithm (default is 0).

`STPTOL` Minimum scaled step in linesearch algorithm (default is eps^(2/3)).

`EPINIT` Swing factor in the Newton iteration convergence test. The test is applied to the residual vector, premultiplied by the approximate Jacobian. For convergence, the weighted RMS norm of this vector (scaled by the error weights) must be less than `EPINIT*EPCON`, where `EPCON` = 0.33 is the analogous test constant used in the time steps. The default is `EPINIT` = 0.01.

`"print initial condition info"`
Set this option to a nonzero value to display detailed information about the initial condition calculation (default is 0).

`"exclude algebraic variables from error test"`
Set to a nonzero value to exclude algebraic variables from the error test. You must also set the `"algebraic variables"` option to declare which variables in the problem are algebraic (default is 0).

`"algebraic variables"`
A vector of the same length as the state vector. A nonzero element indicates that the corresponding element of the state vector is an algebraic variable (i.e., its derivative does not appear explicitly in the equation set).

This option is required by the `"compute consistent initial condition"` and `"exclude algebraic variables from error test"` options.

`"enforce inequality constraints"`
Set to one of the following values to enforce the inequality constraints specified by the `"inequality constraint types"` option (default is 0).

1. To have constraint checking only in the initial condition calculation.
2. To enforce constraint checking during the integration.
3. To enforce both options 1 and 2.

`"inequality constraint types"`
A vector of the same length as the state specifying the type of inequality constraint. Each element of the vector corresponds to an element of the state and should be assigned one of the following codes

-2 Less than zero.

-1 Less than or equal to zero.

0 Not constrained.

1 Greater than or equal to zero.

2 Greater than zero.

This option only has an effect if the `"enforce inequality constraints"` option is nonzero.

`"initial step size"`
: Differential-algebraic problems may occasionally suffer from severe scaling difficulties on the first step. If you know a great deal about the scaling of your problem, you can help to alleviate this problem by specifying an initial stepsize (default is computed automatically).

`"maximum order"`
: Restrict the maximum order of the solution method. This option must be between 1 and 5, inclusive (default is 5).

`"maximum step size"`
: Setting the maximum stepsize will avoid passing over very large regions (default is not specified).

Octave also includes DASSL, an earlier version of DASPK, and DASRT, which can be used to solve DAEs with constraints (stopping conditions).

`[`*x*`, `*xdot*`, `*istate*`, `*msg*`] = dassl (`*fcn*`, `*x_0*`, `*xdot_0*`, `*t*`, `*t_crit*`)

Solve a set of differential-algebraic equations.

`dassl` solves the set of equations

$$0 = f(x, \dot{x}, t)$$

with

$$x(t_0) = x_0, \dot{x}(t_0) = \dot{x}_0$$

The solution is returned in the matrices *x* and *xdot*, with each row in the result matrices corresponding to one of the elements in the vector *t*. The first element of *t* should be $t_0$ and correspond to the initial state of the system *x_0* and its derivative *xdot_0*, so that the first row of the output *x* is *x_0* and the first row of the output *xdot* is *xdot_0*.

The first argument, *fcn*, is a string, inline, or function handle that names the function $f$ to call to compute the vector of residuals for the set of equations. It must have the form

```
res = f (x, xdot, t)
```

in which *x*, *xdot*, and *res* are vectors, and *t* is a scalar.

If *fcn* is a two-element string array or a two-element cell array of strings, inline functions, or function handles, the first element names the function $f$ described above, and the second element names a function to compute the modified Jacobian

$$J = \frac{\partial f}{\partial x} + c\frac{\partial f}{\partial \dot{x}}$$

The modified Jacobian function must have the form

```
jac = j (x, xdot, t, c)
```

The second and third arguments to `dassl` specify the initial condition of the states and their derivatives, and the fourth argument specifies a vector of output times at which the solution is desired, including the time corresponding to the initial condition.

The set of initial states and derivatives are not strictly required to be consistent. In practice, however, DASSL is not very good at determining a consistent set for you, so it is best if you ensure that the initial values result in the function evaluating to zero.

The fifth argument is optional, and may be used to specify a set of times that the DAE solver should not integrate past. It is useful for avoiding difficulties with singularities and points where there is a discontinuity in the derivative.

After a successful computation, the value of *istate* will be greater than zero (consistent with the Fortran version of DASSL).

If the computation is not successful, the value of *istate* will be less than zero and *msg* will contain additional information.

You can use the function `dassl_options` to set optional parameters for `dassl`.

**See also:** [daspk], page 794, [dasrt], page 800, [lsode], page 791.

`dassl_options ()`
`val = dassl_options (`*`opt`*`)`
`dassl_options (`*`opt`*`, `*`val`*`)`

Query or set options for the function `dassl`.

When called with no arguments, the names of all available options and their current values are displayed.

Given one argument, return the value of the option *opt*.

When called with two arguments, `dassl_options` sets the option *opt* to value *val*.

Options include

`"absolute tolerance"`
: Absolute tolerance. May be either vector or scalar. If a vector, it must match the dimension of the state vector, and the relative tolerance must also be a vector of the same length.

`"relative tolerance"`
: Relative tolerance. May be either vector or scalar. If a vector, it must match the dimension of the state vector, and the absolute tolerance must also be a vector of the same length.

  The local error test applied at each integration step is

```
abs (local error in x(i))
     <= rtol(i) * abs (Y(i)) + atol(i)
```

`"compute consistent initial condition"`
: If nonzero, `dassl` will attempt to compute a consistent set of initial conditions. This is generally not reliable, so it is best to provide a consistent set and leave this option set to zero.

`"enforce nonnegativity constraints"`
: If you know that the solutions to your equations will always be non-negative, it may help to set this parameter to a nonzero value. However,

it is probably best to try leaving this option set to zero first, and only setting it to a nonzero value if that doesn't work very well.

`"initial step size"`
: Differential-algebraic problems may occasionally suffer from severe scaling difficulties on the first step. If you know a great deal about the scaling of your problem, you can help to alleviate this problem by specifying an initial stepsize.

`"maximum order"`
: Restrict the maximum order of the solution method. This option must be between 1 and 5, inclusive.

`"maximum step size"`
: Setting the maximum stepsize will avoid passing over very large regions (default is not specified).

`"step limit"`
: Maximum number of integration steps to attempt on a single call to the underlying Fortran code.

`[x, xdot, t_out, istat, msg] = dasrt (fcn, g, x_0, xdot_0, t)`
`... = dasrt (fcn, g, x_0, xdot_0, t, t_crit)`
`... = dasrt (fcn, x_0, xdot_0, t)`
`... = dasrt (fcn, x_0, xdot_0, t, t_crit)`

Solve a set of differential-algebraic equations.

`dasrt` solves the set of equations

$$0 = f(x, \dot{x}, t)$$

with

$$x(t_0) = x_0, \dot{x}(t_0) = \dot{x}_0$$

with functional stopping criteria (root solving).

The solution is returned in the matrices *x* and *xdot*, with each row in the result matrices corresponding to one of the elements in the vector *t_out*. The first element of *t* should be $t_0$ and correspond to the initial state of the system *x_0* and its derivative *xdot_0*, so that the first row of the output *x* is *x_0* and the first row of the output *xdot* is *xdot_0*.

The vector *t* provides an upper limit on the length of the integration. If the stopping condition is met, the vector *t_out* will be shorter than *t*, and the final element of *t_out* will be the point at which the stopping condition was met, and may not correspond to any element of the vector *t*.

The first argument, *fcn*, is a string, inline, or function handle that names the function *f* to call to compute the vector of residuals for the set of equations. It must have the form

```
res = f (x, xdot, t)
```

in which *x*, *xdot*, and *res* are vectors, and *t* is a scalar.

If *fcn* is a two-element string array or a two-element cell array of strings, inline functions, or function handles, the first element names the function $f$ described above, and the second element names a function to compute the modified Jacobian

$$J = \frac{\partial f}{\partial x} + c\frac{\partial f}{\partial \dot{x}}$$

The modified Jacobian function must have the form

```
jac = j (x, xdot, t, c)
```

The optional second argument names a function that defines the constraint functions whose roots are desired during the integration. This function must have the form

```
g_out = g (x, t)
```

and return a vector of the constraint function values. If the value of any of the constraint functions changes sign, DASRT will attempt to stop the integration at the point of the sign change.

If the name of the constraint function is omitted, `dasrt` solves the same problem as `daspk` or `dassl`.

Note that because of numerical errors in the constraint functions due to round-off and integration error, DASRT may return false roots, or return the same root at two or more nearly equal values of $T$. If such false roots are suspected, the user should consider smaller error tolerances or higher precision in the evaluation of the constraint functions.

If a root of some constraint function defines the end of the problem, the input to DASRT should nevertheless allow integration to a point slightly past that root, so that DASRT can locate the root by interpolation.

The third and fourth arguments to `dasrt` specify the initial condition of the states and their derivatives, and the fourth argument specifies a vector of output times at which the solution is desired, including the time corresponding to the initial condition.

The set of initial states and derivatives are not strictly required to be consistent. In practice, however, DASSL is not very good at determining a consistent set for you, so it is best if you ensure that the initial values result in the function evaluating to zero.

The sixth argument is optional, and may be used to specify a set of times that the DAE solver should not integrate past. It is useful for avoiding difficulties with singularities and points where there is a discontinuity in the derivative.

After a successful computation, the value of *istate* will be greater than zero (consistent with the Fortran version of DASSL).

If the computation is not successful, the value of *istate* will be less than zero and *msg* will contain additional information.

You can use the function `dasrt_options` to set optional parameters for `dasrt`.

**See also:** [dasrt_options], page 802, [daspk], page 794, [dasrt], page 800, [lsode], page 791.

`dasrt_options ()`
`val = dasrt_options (`*opt*`)`
`dasrt_options (`*opt*`, `*val*`)`

Query or set options for the function `dasrt`.

When called with no arguments, the names of all available options and their current values are displayed.

Given one argument, return the value of the option *opt*.

When called with two arguments, `dasrt_options` sets the option *opt* to value *val*.

Options include

`"absolute tolerance"`
: Absolute tolerance. May be either vector or scalar. If a vector, it must match the dimension of the state vector, and the relative tolerance must also be a vector of the same length.

`"relative tolerance"`
: Relative tolerance. May be either vector or scalar. If a vector, it must match the dimension of the state vector, and the absolute tolerance must also be a vector of the same length.

  The local error test applied at each integration step is

```
abs (local error in x(i)) <= ...
    rtol(i) * abs (Y(i)) + atol(i)
```

`"initial step size"`
: Differential-algebraic problems may occasionally suffer from severe scaling difficulties on the first step. If you know a great deal about the scaling of your problem, you can help to alleviate this problem by specifying an initial stepsize.

`"maximum order"`
: Restrict the maximum order of the solution method. This option must be between 1 and 5, inclusive.

`"maximum step size"`
: Setting the maximum stepsize will avoid passing over very large regions.

`"step limit"`
: Maximum number of integration steps to attempt on a single call to the underlying Fortran code.

See K. E. Brenan, et al., *Numerical Solution of Initial-Value Problems in Differential-Algebraic Equations*, North-Holland, 1989, DOI: https://doi.org/10.1137/1.9781611971224, for more information about the implementation of DASSL.

## 24.3 Matlab-compatible solvers

Octave also provides a set of solvers for initial value problems for ordinary differential equations (ODEs) that have a MATLAB-compatible interface. The options for this class of methods are set using the functions.

- [odeset], page 809,

- [odeget], page 811,

Currently implemented solvers are:

- Runge-Kutta methods
  - [ode45], page 803, integrates a system of non-stiff ODEs or index-1 differential-algebraic equations (DAEs) using the high-order, variable-step Dormand-Prince method. It requires six function evaluations per integration step, but may take larger steps on smooth problems than `ode23`: potentially offering improved efficiency at smaller tolerances.
  - [ode23], page 804, integrates a system of non-stiff ODEs or (or index-1 DAEs). It uses the third-order Bogacki-Shampine method and adapts the local step size in order to satisfy a user-specified tolerance. The solver requires three function evaluations per integration step.
  - [ode23s], page 805, integrates a system of stiff ODEs (or index-1 DAEs) using a modified second-order Rosenbrock method.
- Linear multistep methods
  - [ode15s], page 806, integrates a system of stiff ODEs (or index-1 DAEs) using a variable step, variable order method based on Backward Difference Formulas (BDF).
  - [ode15i], page 807, integrates a system of fully-implicit ODEs (or index-1 DAEs) using the same variable step, variable order method as `ode15s`. [decic], page 808, can be used to compute consistent initial conditions for `ode15i`.

Detailed information on the solvers are given in L. F. Shampine and M. W. Reichelt, "The MATLAB ODE Suite", *SIAM Journal on Scientific Computing*, Vol. 18, pp. 1–22, 1997, DOI: https://doi.org/10.1137/S1064827594276424.

```
[t, y] = ode45 (fcn, trange, init)
[t, y] = ode45 (fcn, trange, init, ode_opt)
[t, y, te, ye, ie] = ode45 (...)
solution = ode45 (...)
ode45 (...)
```

Solve a set of non-stiff Ordinary Differential Equations (non-stiff ODEs) with the well known explicit Dormand-Prince method of order 4.

*fcn* is a function handle, inline function, or string containing the name of the function that defines the ODE: `y' = f(t,y)`. The function must accept two inputs where the first is time $t$ and the second is a column vector of unknowns $y$.

*trange* specifies the time interval over which the ODE will be evaluated. Typically, it is a two-element vector specifying the initial and final times (`[tinit, tfinal]`). If there are more than two elements then the solution will also be evaluated at these intermediate time instances.

By default, `ode45` uses an adaptive timestep with the `integrate_adaptive` algorithm. The tolerance for the timestep computation may be changed by using the options `"RelTol"` and `"AbsTol"`.

*init* contains the initial value for the unknowns. If it is a row vector then the solution $y$ will be a matrix in which each column is the solution for the corresponding initial value in *init*.

The optional fourth argument *ode_opt* specifies non-default options to the ODE solver. It is a structure generated by `odeset`.

The function typically returns two outputs. Variable *t* is a column vector and contains the times where the solution was found. The output *y* is a matrix in which each column refers to a different unknown of the problem and each row corresponds to a time in *t*.

The output can also be returned as a structure *solution* which has a field *x* containing a row vector of times where the solution was evaluated and a field *y* containing the solution matrix such that each column corresponds to a time in *x*. Use `fieldnames (solution)` to see the other fields and additional information returned.

If no output arguments are requested, and no `"OutputFcn"` is specified in *ode_opt*, then the `"OutputFcn"` is set to `odeplot` and the results of the solver are plotted immediately.

If using the `"Events"` option then three additional outputs may be returned. *te* holds the time when an Event function returned a zero. *ye* holds the value of the solution at time *te*. *ie* contains an index indicating which Event function was triggered in the case of multiple Event functions.

Example: Solve the Van der Pol equation

```
fvdp = @(t,y) [y(2); (1 - y(1)^2) * y(2) - y(1)];
[t,y] = ode45 (fvdp, [0, 20], [2, 0]);
```

**See also:** [odeset], page 809, [odeget], page 811, [ode23], page 804, [ode15s], page 806.

`[t, y] = ode23 (fcn, trange, init)`
`[t, y] = ode23 (fcn, trange, init, ode_opt)`
`[t, y, te, ye, ie] = ode23 (...)`
`solution = ode23 (...)`
`ode23 (...)`

Solve a set of non-stiff Ordinary Differential Equations (non-stiff ODEs) with the well known explicit Bogacki-Shampine method of order 3.

*fcn* is a function handle, inline function, or string containing the name of the function that defines the ODE: `y' = f(t,y)`. The function must accept two inputs where the first is time *t* and the second is a column vector of unknowns *y*.

*trange* specifies the time interval over which the ODE will be evaluated. Typically, it is a two-element vector specifying the initial and final times (`[tinit, tfinal]`). If there are more than two elements then the solution will also be evaluated at these intermediate time instances.

By default, `ode23` uses an adaptive timestep with the `integrate_adaptive` algorithm. The tolerance for the timestep computation may be changed by using the options `"RelTol"` and `"AbsTol"`.

*init* contains the initial value for the unknowns. If it is a row vector then the solution *y* will be a matrix in which each column is the solution for the corresponding initial value in *init*.

The optional fourth argument *ode_opt* specifies non-default options to the ODE solver. It is a structure generated by `odeset`.

The function typically returns two outputs. Variable *t* is a column vector and contains the times where the solution was found. The output *y* is a matrix in which each column refers to a different unknown of the problem and each row corresponds to a time in *t*.

The output can also be returned as a structure *solution* which has a field *x* containing a row vector of times where the solution was evaluated and a field *y* containing the solution matrix such that each column corresponds to a time in *x*. Use `fieldnames (solution)` to see the other fields and additional information returned.

If no output arguments are requested, and no `"OutputFcn"` is specified in *ode_opt*, then the `"OutputFcn"` is set to `odeplot` and the results of the solver are plotted immediately.

If using the `"Events"` option then three additional outputs may be returned. *te* holds the time when an Event function returned a zero. *ye* holds the value of the solution at time *te*. *ie* contains an index indicating which Event function was triggered in the case of multiple Event functions.

Example: Solve the Van der Pol equation

```
fvdp = @(t,y) [y(2); (1 - y(1)^2) * y(2) - y(1)];
[t,y] = ode23 (fvdp, [0, 20], [2, 0]);
```

Reference: For the definition of this method see https://en.wikipedia.org/wiki/List_of_Runge%E2%80%93Kutta_methods.

**See also:** [odeset], page 809, [odeget], page 811, [ode45], page 803, [ode15s], page 806.

`[t, y] = ode23s (fcn, trange, init)`
`[t, y] = ode23s (fcn, trange, init, ode_opt)`
`[t, y] = ode23s (..., par1, par2, ...)`
`[t, y, te, ye, ie] = ode23s (...)`
`solution = ode23s (...)`

Solve a set of stiff Ordinary Differential Equations (stiff ODEs) with a Rosenbrock method of order (2,3).

*fcn* is a function handle, inline function, or string containing the name of the function that defines the ODE: `M y' = f(t,y)`. The function must accept two inputs where the first is time *t* and the second is a column vector of unknowns *y*. *M* is a constant mass matrix, non-singular and possibly sparse. Set the field `"Mass"` in *odeopts* using *odeset* to specify a mass matrix.

*trange* specifies the time interval over which the ODE will be evaluated. Typically, it is a two-element vector specifying the initial and final times (`[tinit, tfinal]`). If there are more than two elements then the solution will also be evaluated at these intermediate time instances using an interpolation procedure of the same order as the one of the solver.

By default, `ode23s` uses an adaptive timestep with the `integrate_adaptive` algorithm. The tolerance for the timestep computation may be changed by using the options `"RelTol"` and `"AbsTol"`.

*init* contains the initial value for the unknowns. If it is a row vector then the solution *y* will be a matrix in which each column is the solution for the corresponding initial value in *init*.

The optional fourth argument *ode_opt* specifies non-default options to the ODE solver. It is a structure generated by `odeset`. `ode23s` will ignore the following options: `"BDF"`, `"InitialSlope"`, `"MassSingular"`, `"MStateDependence"`, `"MvPattern"`, `"MaxOrder"`, `"Non-negative"`.

The function typically returns two outputs. Variable *t* is a column vector and contains the times where the solution was found. The output *y* is a matrix in which each column refers to a different unknown of the problem and each row corresponds to a time in *t*. If *trange* specifies intermediate time steps, only those will be returned.

The output can also be returned as a structure *solution* which has a field *x* containing a row vector of times where the solution was evaluated and a field *y* containing the solution matrix such that each column corresponds to a time in *x*. Use `fieldnames (solution)` to see the other fields and additional information returned.

If using the `"Events"` option then three additional outputs may be returned. *te* holds the time when an Event function returned a zero. *ye* holds the value of the solution at time *te*. *ie* contains an index indicating which Event function was triggered in the case of multiple Event functions.

Example: Solve the stiff Van der Pol equation

```
f = @(t,y) [y(2); 1000*(1 - y(1)^2) * y(2) - y(1)];
opt = odeset ('Mass', [1 0; 0 1], 'MaxStep', 1e-1);
[vt, vy] = ode23s (f, [0 2000], [2 0], opt);
```

**See also:** [odeset], page 809, [daspk], page 794, [dassl], page 798.

`[t, y] = ode15s (fcn, trange, y0)`
`[t, y] = ode15s (fcn, trange, y0, ode_opt)`
`[t, y, te, ye, ie] = ode15s (...)`
`solution = ode15s (...)`
`ode15s (...)`

Solve a set of stiff Ordinary Differential Equations (ODEs) or stiff semi-explicit index 1 Differential Algebraic Equations (DAEs).

`ode15s` uses a variable step, variable order BDF (Backward Differentiation Formula) method that ranges from order 1 to 5.

*fcn* is a function handle, inline function, or string containing the name of the function that defines the ODE: `y' = f(t,y)`. The function must accept two inputs where the first is time *t* and the second is a column vector of unknowns *y*.

*trange* specifies the time interval over which the ODE will be evaluated. Typically, it is a two-element vector specifying the initial and final times (`[tinit, tfinal]`). If there are more than two elements then the solution will also be evaluated at these intermediate time instances.

*init* contains the initial value for the unknowns. If it is a row vector then the solution *y* will be a matrix in which each column is the solution for the corresponding initial value in *init*.

The optional fourth argument *ode_opt* specifies non-default options to the ODE solver. It is a structure generated by `odeset`.

The function typically returns two outputs. Variable *t* is a column vector and contains the times where the solution was found. The output *y* is a matrix in which each

column refers to a different unknown of the problem and each row corresponds to a time in *t*.

The output can also be returned as a structure *solution* which has a field *x* containing a row vector of times where the solution was evaluated and a field *y* containing the solution matrix such that each column corresponds to a time in *x*. Use `fieldnames (solution)` to see the other fields and additional information returned.

If no output arguments are requested, and no `"OutputFcn"` is specified in *ode_opt*, then the `"OutputFcn"` is set to `odeplot` and the results of the solver are plotted immediately.

If using the `"Events"` option then three additional outputs may be returned. *te* holds the time when an Event function returned a zero. *ye* holds the value of the solution at time *te*. *ie* contains an index indicating which Event function was triggered in the case of multiple Event functions.

Example: Solve Robertson's equations:

```
function r = robertson_dae (t, y)
  r = [ -0.04*y(1) + 1e4*y(2)*y(3)
         0.04*y(1) - 1e4*y(2)*y(3) - 3e7*y(2)^2
y(1) + y(2) + y(3) - 1 ];
endfunction
opt = odeset ("Mass", [1 0 0; 0 1 0; 0 0 0], "MStateDependence", "none");
[t,y] = ode15s (@robertson_dae, [0, 1e3], [1; 0; 0], opt);
```

**See also:** [decic], page 808, [odeset], page 809, [odeget], page 811, [ode23], page 804, [ode45], page 803.

`[t, y] = ode15i (fcn, trange, y0, yp0)`
`[t, y] = ode15i (fcn, trange, y0, yp0, ode_opt)`
`[t, y, te, ye, ie] = ode15i (...)`
`solution = ode15i (...)`
`ode15i (...)`

Solve a set of fully-implicit Ordinary Differential Equations (ODEs) or index 1 Differential Algebraic Equations (DAEs).

`ode15i` uses a variable step, variable order BDF (Backward Differentiation Formula) method that ranges from order 1 to 5.

*fcn* is a function handle, inline function, or string containing the name of the function that defines the ODE: `0 = f(t,y,yp)`. The function must accept three inputs where the first is time *t*, the second is the function value *y* (a column vector), and the third is the derivative value *yp* (a column vector).

*trange* specifies the time interval over which the ODE will be evaluated. Typically, it is a two-element vector specifying the initial and final times (`[tinit, tfinal]`). If there are more than two elements then the solution will also be evaluated at these intermediate time instances.

*y0* and *yp0* contain the initial values for the unknowns *y* and *yp*. If they are row vectors then the solution *y* will be a matrix in which each column is the solution for the corresponding initial value in *y0* and *yp0*.

*y0* and *yp0* must be consistent initial conditions, meaning that `f(t,y0,yp0) = 0` is satisfied. The function `decic` may be used to compute consistent initial conditions given initial guesses.

The optional fifth argument *ode_opt* specifies non-default options to the ODE solver. It is a structure generated by `odeset`.

The function typically returns two outputs. Variable *t* is a column vector and contains the times where the solution was found. The output *y* is a matrix in which each column refers to a different unknown of the problem and each row corresponds to a time in *t*.

The output can also be returned as a structure *solution* which has a field *x* containing a row vector of times where the solution was evaluated and a field *y* containing the solution matrix such that each column corresponds to a time in x. Use `fieldnames (solution)` to see the other fields and additional information returned.

If no output arguments are requested, and no `"OutputFcn"` is specified in *ode_opt*, then the `"OutputFcn"` is set to `odeplot` and the results of the solver are plotted immediately.

If using the `"Events"` option then three additional outputs may be returned. *te* holds the time when an Event function returned a zero. *ye* holds the value of the solution at time *te*. *ie* contains an index indicating which Event function was triggered in the case of multiple Event functions.

Example: Solve Robertson's equations:

```
function r = robertson_dae (t, y, yp)
  r = [ -(yp(1) + 0.04*y(1) - 1e4*y(2)*y(3))
        -(yp(2) - 0.04*y(1) + 1e4*y(2)*y(3) + 3e7*y(2)^2)
y(1) + y(2) + y(3) - 1 ];
endfunction
[t,y] = ode15i (@robertson_dae, [0, 1e3], [1; 0; 0], [-1e-4; 1e-4; 0]);
```

**See also:** [decic], page 808, [odeset], page 809, [odeget], page 811.

`[y0_new, yp0_new] = decic (fcn, t0, y0, fixed_y0, yp0, fixed_yp0)`
`[y0_new, yp0_new] = decic (fcn, t0, y0, fixed_y0, yp0, fixed_yp0, options)`
`[y0_new, yp0_new, resnorm] = decic (...)`

Compute consistent implicit ODE initial conditions *y0_new* and *yp0_new* given initial guesses *y0* and *yp0*.

A maximum of `length (y0)` components between *fixed_y0* and *fixed_yp0* may be chosen as fixed values.

*fcn* is a function handle. The function must accept three inputs where the first is time *t*, the second is a column vector of unknowns *y*, and the third is a column vector of unknowns *yp*.

*t0* is the initial time such that `fcn(t0, y0_new, yp0_new) = 0`, specified as a scalar.

*y0* is a vector used as the initial guess for *y*.

*fixed_y0* is a vector which specifies the components of *y0* to hold fixed. Choose a maximum of `length (y0)` components between *fixed_y0* and *fixed_yp0* as fixed values. Set *fixed_y0*(i) component to 1 if you want to fix the value of *y0*(i). Set *fixed_y0*(i) component to 0 if you want to allow the value of *y0*(i) to change.

*yp0* is a vector used as the initial guess for *yp*.

*fixed_yp0* is a vector which specifies the components of *yp0* to hold fixed. Choose a maximum of `length (yp0)` components between *fixed_y0* and *fixed_yp0* as fixed

values. Set *fixed_yp0*(i) component to 1 if you want to fix the value of *yp0*(i). Set *fixed_yp0*(i) component to 0 if you want to allow the value of *yp0*(i) to change.

The optional seventh argument *options* is a structure array. Use `odeset` to generate this structure. The relevant options are `RelTol` and `AbsTol` which specify the error thresholds used to compute the initial conditions.

The function typically returns two outputs. Variable *y0_new* is a column vector and contains the consistent initial value of *y*. The output *yp0_new* is a column vector and contains the consistent initial value of *yp*.

The optional third output *resnorm* is the norm of the vector of residuals. If *resnorm* is small, `decic` has successfully computed the initial conditions. If the value of *resnorm* is large, use `RelTol` and `AbsTol` to adjust it.

Example: Compute initial conditions for Robertson's equations:

```
function r = robertson_dae (t, y, yp)
  r = [ -(yp(1) + 0.04*y(1) - 1e4*y(2)*y(3))
        -(yp(2) - 0.04*y(1) + 1e4*y(2)*y(3) + 3e7*y(2)^2)
y(1) + y(2) + y(3) - 1 ];
endfunction
[y0_new,yp0_new] = decic (@robertson_dae, 0, [1; 0; 0], [1; 1; 0],
[-1e-4; 1; 0], [0; 0; 0]);
```

**See also:** [ode15i], page 807, [odeset], page 809.

*odestruct* = odeset ()
*odestruct* = odeset ("*field1*", *value1*, "*field2*", *value2*, ...)
*odestruct* = odeset (*oldstruct*, "*field1*", *value1*, "*field2*", *value2*, ...)
*odestruct* = odeset (*oldstruct*, *newstruct*)
odeset ()

Create or modify an ODE options structure.

When called with no input argument and one output argument, return a new ODE options structure that contains all possible fields initialized to their default values. If no output argument is requested, display a list of the common ODE solver options along with their default value.

If called with name-value input argument pairs "*field1*", "*value1*", "*field2*", "*value2*", ... return a new ODE options structure with all the most common option fields initialized, **and** set the values of the fields "*field1*", "*field2*", ... to the values *value1*, *value2*, ...

If called with an input structure *oldstruct* then overwrite the values of the options "*field1*", "*field2*", ... with new values *value1*, *value2*, ... and return the modified structure.

When called with two input ODE options structures *oldstruct* and *newstruct* overwrite all values from the structure *oldstruct* with new values from the structure *newstruct*. Empty values in *newstruct* will not overwrite values in *oldstruct*.

The most commonly used ODE options, which are always assigned a value by `odeset`, are the following:

`AbsTol`: positive scalar | vector, def. `1e-6`
Absolute error tolerance.

`BDF`: {`"off"`} | `"on"`
: Use BDF formulas in implicit multistep methods. *Note*: This option is not yet implemented.

`Events`: function_handle
: Event function. An event function must have the form `[value, isterminal, direction] = my_events_f (t, y)`

`InitialSlope`: vector
: Consistent initial slope vector for DAE solvers.

`InitialStep`: positive scalar
: Initial time step size.

`Jacobian`: matrix | function_handle
: Jacobian matrix, specified as a constant matrix or a function of time and state.

`JConstant`: {`"off"`} | `"on"`
: Specify whether the Jacobian is a constant matrix or depends on the state.

`JPattern`: sparse matrix
: If the Jacobian matrix is sparse and non-constant but maintains a constant sparsity pattern, specify the sparsity pattern.

`Mass`: matrix | function_handle
: Mass matrix, specified as a constant matrix or a function of time and state.

`MassSingular`: {`"maybe"`} | `"yes"` | `"on"`
: Specify whether the mass matrix is singular.

`MaxOrder`: {`5`} | `4` | `3` | `2` | `1`
: Maximum order of formula.

`MaxStep`: positive scalar
: Maximum time step value.

`MStateDependence`: {`"weak"`} | `"none"` | `"strong"`
: Specify whether the mass matrix depends on the state or only on time.

`MvPattern`: sparse matrix
: If the mass matrix is sparse and non-constant but maintains a constant sparsity pattern, specify the sparsity pattern. *Note*: This option is not yet implemented.

`NonNegative`: scalar | vector
: Specify elements of the state vector that are expected to remain non-negative during the simulation.

`NormControl`: {`"off"`} | `"on"`
: Control error relative to the 2-norm of the solution, rather than its absolute value.

`OutputFcn`: function_handle
: Function to monitor the state during the simulation. For the form of the function to use see [odeplot], page 811.

`OutputSel`: scalar | vector
: Indices of elements of the state vector to be passed to the output monitoring function.

`Refine`: positive scalar
: Specify whether output should be returned only at the end of each time step or also at intermediate time instances. The value should be a scalar indicating the number of equally spaced time points to use within each timestep at which to return output.

`RelTol`: positive scalar
: Relative error tolerance.

`Stats`: {"off"} | "on"
: Print solver statistics after simulation.

`Vectorized`: {"off"} | "on"
: Specify whether `odefcn` can be passed multiple values of the state at once.

Field names that are not in the above list are also accepted and added to the result structure.

**See also:** [odeget], page 811.

*val* = odeget (*ode_opt*, *field*)
*val* = odeget (*ode_opt*, *field*, *default*)
Query the value of the property *field* in the ODE options structure *ode_opt*.

If called with two input arguments and the first input argument *ode_opt* is an ODE option structure and the second input argument *field* is a string specifying an option name, then return the option value *val* corresponding to *field* from *ode_opt*.

If called with an optional third input argument, and *field* is not set in the structure *ode_opt*, then return the default value *default* instead.

**See also:** [odeset], page 809.

*stop_solve* = odeplot (*t*, *y*, *flag*)
Open a new figure window and plot the solution of an ode problem at each time step during the integration.

The types and values of the input parameters $t$ and $y$ depend on the input *flag* that is of type string. Valid values of *flag* are:

`"init"`
: The input $t$ must be a column vector of length 2 with the first and last time step (`[tfirst tlast]`). The input $y$ contains the initial conditions for the ode problem (*y0*).

`""`
: The input $t$ must be a scalar double or vector specifying the time(s) for which the solution in input $y$ was calculated.

`"done"` The inputs should be empty, but are ignored if they are present.

`odeplot` always returns false, i.e., don't stop the ode solver.

Example: solve an anonymous implementation of the `"Van der Pol"` equation and display the results while solving.

```
fvdp = @(t,y) [y(2); (1 - y(1)^2) * y(2) - y(1)];

opt = odeset ("OutputFcn", @odeplot, "RelTol", 1e-6);
sol = ode45 (fvdp, [0 20], [2 0], opt);
```

Background Information: This function is called by an ode solver function if it was specified in the `"OutputFcn"` property of an options structure created with `odeset`. The ode solver will initially call the function with the syntax `odeplot ([`*`tfirst`*`, `*`tlast`*`], `*`y0`*`, "init")`. The function initializes internal variables, creates a new figure window, and sets the x limits of the plot. Subsequently, at each time step during the integration the ode solver calls `odeplot (`*`t`*`, `*`y`*`, [])`. At the end of the solution the ode solver calls `odeplot ([], [], "done")` so that odeplot can perform any clean-up actions required.

**See also:** [odeset], page 809, [odeget], page 811, [ode23], page 804, [ode45], page 803.

# 25 Optimization

Octave comes with support for solving various kinds of optimization problems. Specifically Octave can solve problems in Linear Programming, Quadratic Programming, Nonlinear Programming, and Linear Least Squares Minimization.

## 25.1 Linear Programming

Octave can solve Linear Programming problems using the `glpk` function. That is, Octave can solve

$$\min_x c^T x$$

subject to the linear constraints $Ax = b$ where $x \geq 0$.

The `glpk` function also supports variations of this problem.

`[`*`xopt`*`, `*`fmin`*`, `*`errnum`*`, `*`extra`*`] = glpk (`*`c`*`, `*`A`*`, `*`b`*`, `*`lb`*`, `*`ub`*`, `*`ctype`*`, `*`vartype`*`, `*`sense`*`, `*`param`*`)`

Solve a linear program using the GNU GLPK library.

Given three arguments, `glpk` solves the following standard LP:

$$\min_x C^T x$$

subject to

$$Ax = b \qquad x \geq 0$$

but may also solve problems of the form

$$[\min_x \mid \max_x] C^T x$$

subject to

$$Ax [= \mid \leq \mid \geq] b \qquad LB \leq x \leq UB$$

Input arguments:

*c* — A column array containing the objective function coefficients.

*A* — A matrix containing the constraints coefficients.

*b* — A column array containing the right-hand side value for each constraint in the constraint matrix.

*lb* — An array containing the lower bound on each of the variables. If *lb* is not supplied, the default lower bound for the variables is zero.

*ub* — An array containing the upper bound on each of the variables. If *ub* is not supplied, the default upper bound is assumed to be infinite.

*ctype* — An array of characters containing the sense of each constraint in the constraint matrix. Each element of the array may be one of the following values

`"F"` — A free (unbounded) constraint (the constraint is ignored).

`"U"` An inequality constraint with an upper bound (`A(i,:)*x <= b(i)`).

`"S"` An equality constraint (`A(i,:)*x = b(i)`).

`"L"` An inequality with a lower bound (`A(i,:)*x >= b(i)`).

`"D"` An inequality constraint with both upper and lower bounds (`A(i,:)*x >= -b(i)`) *and* (`A(i,:)*x <= b(i)`).

*vartype* A column array containing the types of the variables.

`"C"` A continuous variable.

`"I"` An integer variable.

*sense* If *sense* is 1, the problem is a minimization. If *sense* is -1, the problem is a maximization. The default value is 1.

*param* A structure containing the following parameters used to define the behavior of solver. Missing elements in the structure take on default values, so you only need to set the elements that you wish to change from the default.

Integer parameters:

`msglev (default: 1)`
Level of messages output by solver routines:

0 (`GLP_MSG_OFF`)
No output.

1 (`GLP_MSG_ERR`)
Error and warning messages only.

2 (`GLP_MSG_ON`)
Normal output.

3 (`GLP_MSG_ALL`)
Full output (includes informational messages).

`scale (default: 16)`
Scaling option. The values can be combined with the bitwise OR operator and may be the following:

1 (`GLP_SF_GM`)
Geometric mean scaling.

16 (`GLP_SF_EQ`)
Equilibration scaling.

32 (`GLP_SF_2N`)
Round scale factors to power of two.

64 (`GLP_SF_SKIP`)
Skip if problem is well scaled.

Alternatively, a value of 128 (`GLP_SF_AUTO`) may be also specified, in which case the routine chooses the scaling options automatically.

`dual (default: 1)`
Simplex method option:

1 (`GLP_PRIMAL`)
Use two-phase primal simplex.

2 (`GLP_DUALP`)
Use two-phase dual simplex, and if it fails, switch to the primal simplex.

3 (`GLP_DUAL`)
Use two-phase dual simplex.

`price (default: 34)`
Pricing option (for both primal and dual simplex):

17 (`GLP_PT_STD`)
Textbook pricing.

34 (`GLP_PT_PSE`)
Steepest edge pricing.

`itlim (default: intmax)`
Simplex iterations limit. It is decreased by one each time when one simplex iteration has been performed, and reaching zero value signals the solver to stop the search.

`outfrq (default: 200)`
Output frequency, in iterations. This parameter specifies how frequently the solver sends information about the solution to the standard output.

`branch (default: 4)`
Branching technique option (for MIP only):

1 (`GLP_BR_FFV`)
First fractional variable.

2 (`GLP_BR_LFV`)
Last fractional variable.

3 (`GLP_BR_MFV`)
Most fractional variable.

4 (`GLP_BR_DTH`)
Heuristic by Driebeck and Tomlin.

5 (`GLP_BR_PCH`)
Hybrid pseudocost heuristic.

`btrack (default: 4)`
: Backtracking technique option (for MIP only):

  1 (`GLP_BT_DFS`)
  : Depth first search.

  2 (`GLP_BT_BFS`)
  : Breadth first search.

  3 (`GLP_BT_BLB`)
  : Best local bound.

  4 (`GLP_BT_BPH`)
  : Best projection heuristic.

`presol (default: 1)`
: If this flag is set, the simplex solver uses the built-in LP presolver. Otherwise the LP presolver is not used.

`lpsolver (default: 1)`
: Select which solver to use. If the problem is a MIP problem this flag will be ignored.

  1 Revised simplex method.

  2 Interior point method.

`rtest (default: 34)`
: Ratio test technique:

  17 (`GLP_RT_STD`)
  : Standard ("textbook").

  34 (`GLP_RT_HAR`)
  : Harris' two-pass ratio test.

`tmlim (default: intmax)`
: Searching time limit, in milliseconds.

`outdly (default: 0)`
: Output delay, in seconds. This parameter specifies how long the solver should delay sending information about the solution to the standard output.

`save (default: 0)`
: If this parameter is nonzero, save a copy of the problem in CPLEX LP format to the file `"outpb.lp"`. There is currently no way to change the name of the output file.

Real parameters:

`tolbnd (default: 1e-7)`
: Relative tolerance used to check if the current basic solution is primal feasible. It is not recommended that you change this parameter unless you have a detailed understanding of its purpose.

`toldj (default: 1e-7)`
: Absolute tolerance used to check if the current basic solution is dual feasible. It is not recommended that you change this parameter unless you have a detailed understanding of its purpose.

`tolpiv (default: 1e-9)`
: Relative tolerance used to choose eligible pivotal elements of the simplex table. It is not recommended that you change this parameter unless you have a detailed understanding of its purpose.

`objll (default: -DBL_MAX)`
: Lower limit of the objective function. If the objective function reaches this limit and continues decreasing, the solver stops the search. This parameter is used in the dual simplex method only.

`objul (default: +DBL_MAX)`
: Upper limit of the objective function. If the objective function reaches this limit and continues increasing, the solver stops the search. This parameter is used in the dual simplex only.

`tolint (default: 1e-5)`
: Relative tolerance used to check if the current basic solution is integer feasible. It is not recommended that you change this parameter unless you have a detailed understanding of its purpose.

`tolobj (default: 1e-7)`
: Relative tolerance used to check if the value of the objective function is not better than in the best known integer feasible solution. It is not recommended that you change this parameter unless you have a detailed understanding of its purpose.

Output values:

*xopt* The optimizer (the value of the decision variables at the optimum).

*fopt* The optimum value of the objective function.

*errnum* Error code.

0 No error.

1 (`GLP_EBADB`)
: Invalid basis.

2 (`GLP_ESING`)
: Singular matrix.

3 (`GLP_ECOND`)
: Ill-conditioned matrix.

4 (`GLP_EBOUND`)
: Invalid bounds.

5 (`GLP_EFAIL`)
: Solver failed.

6 (`GLP_EOBJLL`)
: Objective function lower limit reached.

7 (`GLP_EOBJUL`)
: Objective function upper limit reached.

8 (`GLP_EITLIM`)
: Iterations limit exhausted.

9 (`GLP_ETMLIM`)
: Time limit exhausted.

10 (`GLP_ENOPFS`)
: No primal feasible solution.

11 (`GLP_ENODFS`)
: No dual feasible solution.

12 (`GLP_EROOT`)
: Root LP optimum not provided.

13 (`GLP_ESTOP`)
: Search terminated by application.

14 (`GLP_EMIPGAP`)
: Relative MIP gap tolerance reached.

15 (`GLP_ENOFEAS`)
: No primal/dual feasible solution.

16 (`GLP_ENOCVG`)
: No convergence.

17 (`GLP_EINSTAB`)
: Numerical instability.

18 (`GLP_EDATA`)
: Invalid data.

19 (`GLP_ERANGE`)
: Result out of range.

*extra* A data structure containing the following fields:

`lambda` Dual variables.

`redcosts` Reduced Costs.

`time` Time (in seconds) used for solving LP/MIP problem.

`status` Status of the optimization.

1 (`GLP_UNDEF`)
: Solution status is undefined.

2 (GLP_FEAS)
: Solution is feasible.

3 (GLP_INFEAS)
: Solution is infeasible.

4 (GLP_NOFEAS)
: Problem has no feasible solution.

5 (GLP_OPT)
: Solution is optimal.

6 (GLP_UNBND)
: Problem has no unbounded solution.

Example:

```
c = [10, 6, 4]';
A = [ 1, 1, 1;
     10, 4, 5;
      2, 2, 6];
b = [100, 600, 300]';
lb = [0, 0, 0]';
ub = [];
ctype = "UUU";
vartype = "CCC";
s = -1;

param.msglev = 1;
param.itlim = 100;

[xmin, fmin, status, extra] = ...
   glpk (c, A, b, lb, ub, ctype, vartype, s, param);
```

## 25.2 Quadratic Programming

Octave can also solve Quadratic Programming problems, this is

$$\min_x \frac{1}{2} x^T H x + x^T q$$

subject to

$$Ax = b \qquad lb \leq x \leq ub \qquad A_{lb} \leq A_{in} x \leq A_{ub}$$

[*x*, *obj*, *info*, *lambda*] = qp (*x0*, *H*)
[*x*, *obj*, *info*, *lambda*] = qp (*x0*, *H*, *q*)
[*x*, *obj*, *info*, *lambda*] = qp (*x0*, *H*, *q*, *A*, *b*)
[*x*, *obj*, *info*, *lambda*] = qp (*x0*, *H*, *q*, *A*, *b*, *lb*, *ub*)
[*x*, *obj*, *info*, *lambda*] = qp (*x0*, *H*, *q*, *A*, *b*, *lb*, *ub*, *A_lb*, *A_in*, *A_ub*)
[*x*, *obj*, *info*, *lambda*] = qp (*...*, *options*)
: Solve a quadratic program (QP).

Solve the quadratic program defined by

$$\min_x \frac{1}{2} x^T H x + x^T q$$

subject to

$$Ax = b \qquad lb \leq x \leq ub \qquad A_{lb} \leq A_{in} x \leq A_{ub}$$

using a null-space active-set method.

Any bound ($A$, $b$, $lb$, $ub$, $A_in$, $A_lb$, $A_ub$) may be set to the empty matrix (`[]`) if not present. The constraints $A$ and $A_in$ are matrices with each row representing a single constraint. The other bounds are scalars or vectors depending on the number of constraints. The algorithm is faster if the initial guess is feasible.

*options* is a structure specifying additional parameters which control the algorithm. Currently, `qp` recognizes these options: `"MaxIter"`, `"TolX"`, `"AllowSemidefinite"`.

`"MaxIter"` sets the maximum number of algorithm iterations before optimization is halted. The default value is 200. The value must be a positive integer.

`"TolX"` specifies the termination tolerance for the unknown variables $x$. The default is `sqrt (eps)` or approximately 1e-8.

`"AllowSemidefinite"` is a boolean flag: `true` allows positive semi-definite problems (one or more eigenvalues of the Hessian are zero). `false` requires positive definiteness (all eigenvalues positive). The default is `false`.

On return, $x$ is the location of the minimum and *fval* contains the value of the objective function at $x$.

*info* Structure containing run-time information about the algorithm. The following fields are defined:

`solveiter`
: The number of iterations required to find the solution.

`info`
: An integer indicating the status of the solution.

| | |
|---|---|
| 0 | The problem is feasible and convex. Global solution found. |
| 1 | The problem is not convex. Local solution found. |
| 2 | The problem is not convex and unbounded. |
| 3 | Maximum number of iterations reached. |
| 6 | The problem is infeasible. |

**See also:** [sqp], page 821.

```
x = pqpnonneg (c, d)
x = pqpnonneg (c, d, x0)
x = pqpnonneg (c, d, x0, options)
[x, minval] = pqpnonneg (...)
[x, minval, exitflag] = pqpnonneg (...)
[x, minval, exitflag, output] = pqpnonneg (...)
```

`[x, minval, exitflag, output, lambda] = pqpnonneg (...)`

Minimize `(1/2 * x' * c * x + d' * x)` subject to `x >= 0`.

*c* and *d* must be real matrices, and *c* must be symmetric and positive definite.

*x0* is an optional initial guess for the solution *x*.

*options* is an options structure to change the behavior of the algorithm (see [optimset], page 825). `pqpnonneg` recognizes one option: `"MaxIter"`.

Outputs:

*x* The solution matrix

*minval* The minimum attained model value, `1/2*xmin'*c*xmin + d'*xmin`

*exitflag* An indicator of convergence. 0 indicates that the iteration count was exceeded, and therefore convergence was not reached; >0 indicates that the algorithm converged. (The algorithm is stable and will converge given enough iterations.)

*output* A structure with two fields:

- `"algorithm"`: The algorithm used (`"nnls"`)
- `"iterations"`: The number of iterations taken.

*lambda* Lagrange multipliers. If these are nonzero, the corresponding *x* values should be zero, indicating the solution is pressed up against a coordinate plane. The magnitude indicates how much the residual would improve if the `x >= 0` constraints were relaxed in that direction.

**See also:** [lsqnonneg], page 824, [qp], page 819, [optimset], page 825.

## 25.3 Nonlinear Programming

Octave can also perform general nonlinear minimization using a successive quadratic programming solver.

`[x, obj, info, iter, nf, lambda] = sqp (x0, phi)`
`[...] = sqp (x0, phi, g)`
`[...] = sqp (x0, phi, g, h)`
`[...] = sqp (x0, phi, g, h, lb, ub)`
`[...] = sqp (x0, phi, g, h, lb, ub, maxiter)`
`[...] = sqp (x0, phi, g, h, lb, ub, maxiter, tolerance)`

Minimize an objective function using sequential quadratic programming (SQP).

Solve the nonlinear program

$$\min_x \phi(x)$$

subject to

$$g(x) = 0 \qquad h(x) \geq 0 \qquad lb \leq x \leq ub$$

using a sequential quadratic programming method.

The first argument is the initial guess for the vector *x0*.

The second argument is a function handle pointing to the objective function *phi*. The objective function must accept one vector argument and return a scalar.

The second argument may also be a 2- or 3-element cell array of function handles. The first element should point to the objective function, the second should point to a function that computes the gradient of the objective function, and the third should point to a function that computes the Hessian of the objective function. If the gradient function is not supplied, the gradient is computed by finite differences. If the Hessian function is not supplied, a BFGS update formula is used to approximate the Hessian.

When supplied, the gradient function `phi{2}` must accept one vector argument and return a vector. When supplied, the Hessian function `phi{3}` must accept one vector argument and return a matrix.

The third and fourth arguments $g$ and $h$ are function handles pointing to functions that compute the equality constraints and the inequality constraints, respectively. If the problem does not have equality (or inequality) constraints, then use an empty matrix ([]) for $g$ (or $h$). When supplied, these equality and inequality constraint functions must accept one vector argument and return a vector.

The third and fourth arguments may also be 2-element cell arrays of function handles. The first element should point to the constraint function and the second should point to a function that computes the gradient of the constraint function:

$$\left( \frac{\partial f(x)}{\partial x_1}, \frac{\partial f(x)}{\partial x_2}, \ldots, \frac{\partial f(x)}{\partial x_N} \right)^T$$

The fifth and sixth arguments, *lb* and *ub*, contain lower and upper bounds on $x$ and when provided must be vectors of the same size as the vector *x0*. The bounds must be consistent with the equality and inequality constraints $g$ and $h$.

The seventh argument *maxiter* specifies the maximum number of iterations. The default value is 100.

The eighth argument *tolerance* specifies the tolerance for the stopping criteria. The default value is `sqrt (eps)`.

The value returned in *info* may be one of the following:

101 The algorithm terminated normally. All constraints meet the specified tolerance.

102 The BFGS update failed.

103 The maximum number of iterations was reached.

104 The stepsize has become too small, i.e., $\Delta x$, is less than `tol * norm (x)`.

An example of calling `sqp`:

```
function r = g (x)
  r = [ sumsq(x)-10;
        x(2)*x(3)-5*x(4)*x(5);
        x(1)^3+x(2)^3+1 ];
endfunction

function obj = phi (x)
```

```
  obj = exp (prod (x)) - 0.5*(x(1)^3+x(2)^3+1)^2;
endfunction

x0 = [-1.8; 1.7; 1.9; -0.8; -0.8];

[x, obj, info, iter, nf, lambda] = sqp (x0, @phi, @g, [])

x =

  -1.71714
   1.59571
   1.82725
  -0.76364
  -0.76364

obj = 0.053950
info = 101
iter = 8
nf = 10
lambda =

  -0.0401627
   0.0379578
  -0.0052227
```

**See also:** [qp], page 819.

## 25.4 Linear Least Squares

Octave also supports linear least squares minimization. That is, Octave can find the parameter $b$ such that the model $y = xb$ fits data $(x, y)$ as well as possible, assuming zero-mean Gaussian noise. If the noise is assumed to be isotropic the problem can be solved using the '\' or '/' operators, or the `ols` function. In the general case where the noise is assumed to be anisotropic the `gls` is needed.

`[beta, sigma, r] = ols (y, x)`

Ordinary least squares (OLS) estimation.

OLS applies to the multivariate model $y = x\,b + e$ where $y$ is a $t \times p$ matrix, $x$ is a $t \times k$ matrix, $b$ is a $k \times p$ matrix, and $e$ is a $t \times p$ matrix.

Each row of $y$ is a $p$-variate observation in which each column represents a variable. Likewise, the rows of $x$ represent $k$-variate observations or possibly designed values. Furthermore, the collection of observations $x$ must be of adequate rank, $k$, otherwise $b$ cannot be uniquely estimated.

The observation errors, $e$, are assumed to originate from an underlying $p$-variate distribution with zero mean and $p$-by-$p$ covariance matrix $S$, both constant conditioned on $x$. Furthermore, the matrix $S$ is constant with respect to each observation such that $\bar{e} = 0$ and $\mathrm{cov}(\mathrm{vec}(e)) = \mathrm{kron}(s,I)$. (For cases that don't meet this criteria,

such as autocorrelated errors, see generalized least squares, gls, for more efficient estimations.)

The return values *beta*, *sigma*, and *r* are defined as follows.

*beta* The OLS estimator for matrix *b*. *beta* is calculated directly via $(x^T x)^{-1} x^T y$ if the matrix $x^T x$ is of full rank. Otherwise, `beta = pinv (x) * y` where `pinv (x)` denotes the pseudoinverse of *x*.

*sigma* The OLS estimator for the matrix *s*,

```
sigma = (y-x*beta)' * (y-x*beta) / (t-rank(x))
```

*r* The matrix of OLS residuals, `r = y - x*beta`.

**See also:** [gls], page 824, [pinv], page 655.

`[beta, v, r] = gls (y, x, o)`

Generalized least squares (GLS) model.

Perform a generalized least squares estimation for the multivariate model $y = x\,b + e$ where $y$ is a $t \times p$ matrix, $x$ is a $t \times k$ matrix, $b$ is a $k \times p$ matrix and $e$ is a $t \times p$ matrix.

Each row of $y$ is a $p$-variate observation in which each column represents a variable. Likewise, the rows of $x$ represent $k$-variate observations or possibly designed values. Furthermore, the collection of observations $x$ must be of adequate rank, $k$, otherwise $b$ cannot be uniquely estimated.

The observation errors, $e$, are assumed to originate from an underlying $p$-variate distribution with zero mean but possibly heteroscedastic observations. That is, in general, $\bar{e} = 0$ and $\mathrm{cov}(\mathrm{vec}(e)) = s^2 o$ in which $s$ is a scalar and $o$ is a $t\,p \times t\,p$ matrix.

The return values *beta*, *v*, and *r* are defined as follows.

*beta* The GLS estimator for matrix *b*.

*v* The GLS estimator for scalar $s^2$.

*r* The matrix of GLS residuals, $r = y - x * beta$.

**See also:** [ols], page 823.

```
x = lsqnonneg (c, d)
x = lsqnonneg (c, d, x0)
x = lsqnonneg (c, d, x0, options)
[x, resnorm] = lsqnonneg (...)
[x, resnorm, residual] = lsqnonneg (...)
[x, resnorm, residual, exitflag] = lsqnonneg (...)
[x, resnorm, residual, exitflag, output] = lsqnonneg (...)
[x, resnorm, residual, exitflag, output, lambda] = lsqnonneg (...)
```

Minimize `norm (c*x - d)` subject to `x >= 0`.

*c* and *d* must be real matrices.

*x0* is an optional initial guess for the solution *x*.

*options* is an options structure to change the behavior of the algorithm (see [optimset], page 825). `lsqnonneg` recognizes these options: `"MaxIter"`, `"TolX"`.

Outputs:

*resnorm* The squared 2-norm of the residual: `norm (c*x-d)^2`

*residual* The residual: `d-c*x`

*exitflag* An indicator of convergence. 0 indicates that the iteration count was exceeded, and therefore convergence was not reached; >0 indicates that the algorithm converged. (The algorithm is stable and will converge given enough iterations.)

*output* A structure with two fields:

- `"algorithm"`: The algorithm used (`"nnls"`)
- `"iterations"`: The number of iterations taken.

*lambda* Lagrange multipliers. If these are nonzero, the corresponding $x$ values should be zero, indicating the solution is pressed up against a coordinate plane. The magnitude indicates how much the residual would improve if the `x >= 0` constraints were relaxed in that direction.

**See also:** [pqpnonneg], page 820, [lscov], page 825, [optimset], page 825.

`x = lscov (A, b)`
`x = lscov (A, b, V)`
`x = lscov (A, b, V, alg)`
`[x, stdx, mse, S] = lscov (...)`

Compute a generalized linear least squares fit.

Estimate $x$ under the model $b = Ax + w$, where the noise $w$ is assumed to follow a normal distribution with covariance matrix $\sigma^2 V$.

If the size of the coefficient matrix $A$ is n-by-p, the size of the vector/array of constant terms $b$ must be n-by-k.

The optional input argument $V$ may be an n-element vector of positive weights (inverse variances), or an n-by-n symmetric positive semi-definite matrix representing the covariance of $b$. If $V$ is not supplied, the ordinary least squares solution is returned.

The *alg* input argument, a guidance on solution method to use, is currently ignored.

Besides the least-squares estimate matrix $x$ (p-by-k), the function also returns *stdx* (p-by-k), the error standard deviation of estimated $x$; *mse* (k-by-1), the estimated data error covariance scale factors ($\sigma^2$); and $S$ (p-by-p, or p-by-p-by-k if k > 1), the error covariance of $x$.

Reference: Golub and Van Loan, *Matrix Computations (3rd Ed.)*, Johns Hopkins, Section 5.6.3, 1996.

**See also:** [ols], page 823, [gls], page 824, [lsqnonneg], page 824.

`optimset ()`
`options = optimset ()`
`options = optimset (par, val, ...)`
`options = optimset (old, par, val, ...)`
`options = optimset (old, new)`

Create options structure for optimization functions.

When called without any input or output arguments, `optimset` prints a list of all valid optimization parameters.

When called with one output and no inputs, return an options structure with all valid option parameters initialized to `[]`.

When called with a list of parameter/value pairs, return an options structure with only the named parameters initialized.

When the first input is an existing options structure *old*, the values are updated from either the *par*/*val* list or from the options structure *new*.

If *par* does not exactly match the name of a standard parameter, `optimset` will attempt to match *par* to a standard parameter and will set the value of that parameter if a match is found. Matching is case insensitive and is based on character matching at the start of the parameter name. `optimset` produces an error if it finds multiple ambiguous matches. If no standard parameter matches are found a warning is issued and the non-standard parameter is created.

Standard list of valid parameters:

AutoScaling
ComplexEqn
Display Request verbose display of results from optimizations. Values are:

`"off"` [default]
No display.

`"iter"` Display intermediate results for every loop iteration.

`"final"` Display the result of the final loop iteration.

`"notify"` Display the result of the final loop iteration if the function has failed to converge.

FinDiffType
FunValCheck
When enabled, display an error if the objective function returns an invalid value (a complex number, NaN, or Inf). Must be set to `"on"` or `"off"` [default]. Note: the functions `fzero` and `fminbnd` correctly handle Inf values and only complex values or NaN will cause an error in this case.

GradObj When set to `"on"`, the function to be minimized must return a second argument which is the gradient, or first derivative, of the function at the point *x*. If set to `"off"` [default], the gradient is computed via finite differences.

Jacobian When set to `"on"`, the function to be minimized must return a second argument which is the Jacobian, or first derivative, of the function at the point *x*. If set to `"off"` [default], the Jacobian is computed via finite differences.

MaxFunEvals
Maximum number of function evaluations before optimization stops. Must be a positive integer.

MaxIter Maximum number of algorithm iterations before optimization stops. Must be a positive integer.

OutputFcn
A user-defined function executed once per algorithm iteration.

TolFun Termination criterion for the function output. If the difference in the calculated objective function between one algorithm iteration and the next is less than `TolFun` the optimization stops. Must be a positive scalar.

TolX Termination criterion for the function input. If the difference in *x*, the current search point, between one algorithm iteration and the next is less than `TolX` the optimization stops. Must be a positive scalar.

TypicalX

Updating

This list can be extended by the user or other loaded Octave packages. An updated valid parameters list can be queried using the no-argument form of `optimset`.

Note 1: Only parameter names from the standard list are considered when matching short parameter names, and *par* will always be expanded to match a standard parameter even if an exact non-standard match exists. The value of a non-standard parameter that is ambiguous with one or more standard parameters cannot be set by `optimset` and can only be set using `setfield` or dot notation for structs.

Note 2: The optimization options structure is primarily intended for manipulation of known parameters by `optimset` and `optimget`. Unpredictable behavior on future calls to `optimset` or `optimget` can result from creating non-standard or ambiguous parameters or from loading/unloading packages that change the known parameter list after creation of an optimization options structure.

**See also:** [optimget], page 827.

`val = optimget (options, par)`
`val = optimget (options, par, default)`
Return the value of the specific parameter *par* from the optimization options structure *options* created by `optimset`.

If *par* is not defined then return the *default* value if supplied, otherwise return an empty matrix.

If *par* does not exactly match the name of a standard parameter, `optimget` will attempt to match *par* to a standard parameter and will return that parameter's value if a match is found. Matching is case insensitive and is based on character matching at the start of the parameter name. `optimget` produces an error if it finds multiple ambiguous matches. If no standard parameter matches are found a warning is issued. See `optimset` for information about the standard options list.

Note: Only parameter names from the standard list are considered when matching short parameter names, and *par* will always be expanded to match a standard parameter even if an exact non-standard match exists. The value of a non-standard

parameter that is ambiguous with one or more standard parameters cannot be returned by `optimget` and can only be accessed using `getfield` or dot notation for structs.

**See also:** [optimset], page 825.

# 26 Statistics

Octave has support for various statistical methods, emphasizing basic descriptive statistics. More advanced functions for probability distributions, statistical tests, random number generation, etc. are provided by the add-on Statistics package from Octave Packages.

The functions that analyze data all assume that multi-dimensional data is arranged in a matrix where each row is an observation, and each column is a variable. Thus, the matrix defined by

```
a = [ 0.9, 0.7;
      0.1, 0.1;
      0.5, 0.4 ];
```

contains three observations from a two-dimensional distribution. While this is the default data arrangement, most functions support different arrangements.

It should be noted that the statistics functions don't test for data containing NaN, NA, or Inf. These values need to be detected and dealt with explicitly. See [isnan], page 567, [isna], page 46, [isinf], page 567, and [isfinite], page 567.

## 26.1 Descriptive Statistics

One principal goal of descriptive statistics is to represent the essence of a large data set concisely. Octave provides the mean, median, and mode functions which all summarize a data set with just a single number corresponding to the central tendency of the data.

*m* = **mean** (*x*)
*m* = **mean** (*x*, *dim*)
*m* = **mean** (*x*, *vecdim*)
*m* = **mean** (*x*, "*all*")
*m* = **mean** (..., *nanflag*)
*m* = **mean** (..., *outtype*)
*m* = **mean** (..., 'Weights', *w*)

Compute the mean of the elements of *x*.

The mean is defined as

$$\text{mean}(x) \equiv \bar{x} = \frac{1}{N} \sum_{i=1}^{N} x_i$$

where $N$ is the number of elements of *x*.

The weighted mean is defined as

$$\text{mean}_\text{w}(x) \equiv \bar{x}_w = \frac{\sum_{i=1}^{N} w_i x_i}{\sum_{i=1}^{N} w_i}$$

where $N$ is the number of elements of *x*.

If *x* is a vector, then `mean (x)` returns the mean of the elements in *x*.

If *x* is a matrix, then `mean (x)` returns a row vector with each element containing the mean of the corresponding column in *x*.

If $x$ is an array, then `mean (x)` computes the mean along the first non-singleton dimension of $x$.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension of $x$, including any dimension exceeding `ndims (x)`, will return $x$.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of $x$, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored.

Specifying the dimension as `"all"` will cause `mean` to operate on all elements of $x$, and is equivalent to `mean (x(:))`.

The optional input *outtype* specifies the data type that is returned. *outtype* can take the following values:

`'default'`
: Output is of type double, unless the input is single in which case the output is of type single.

`'double'`
: Output is of type double.

`'native'`
: Output is of the same type as the input as reported by `class (x)`, unless the input is logical in which case the output is of type double.

The optional variable *nanflag* specifies whether to include or exclude NaN values from the calculation using any of the previously specified input argument combinations. The default value for *nanflag* is `"includenan"` which keeps NaN values in the calculation. To exclude NaN values set the value of *nanflag* to `"omitnan"`. The output will still contain NaN values if $x$ consists of all NaN values in the operating dimension.

The optional argument pair `"Weights", w` specifies a weighting scheme $w$, which is applied on input $x$, so that `mean` computes the weighted mean. When operating along a single dimension, $w$ must be a vector of the same length as the operating dimension or it must have the same size as $x$. When operating over an array slice defined by *vecdim*, $w$ must have the same size as the operating array slice, i.e., `size (w) == size (x)(vecdim)`, or the same size as $x$.

**See also:** [median], page 830, [mode], page 831, [movmean], page 852.

`m = median (x)`
`m = median (x, dim)`
`m = median (x, vecdim)`
`m = median (x, "all")`
`m = median (..., nanflag)`
`m = median (..., outtype)`

Compute the median value of the elements of $x$.

The median is defined on the sorted data $s$ (`s = sort (x)`) as

$$\text{median}(x) = \begin{cases} s(\lceil N/2 \rceil), & N \text{ odd;} \\ (s(N/2) + s(N/2+1))/2, & N \text{ even.} \end{cases}$$

where $N$ is the number of elements of $x$.

If *x* is a vector, then `median (x)` returns the median of the elements in *x*.

If *x* is a matrix, then `median (x)` returns a row vector with each element containing the median of the corresponding column in *x*.

If *x* is an array, then `median (x)` computes the median along the first non-singleton dimension of *x*.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension of *x*, including any dimension exceeding `ndims (x)`, will return *x*.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored.

Specifying the dimension as `"all"` will cause `median` to operate on all elements of *x*, and is equivalent to `median (x(:))`.

`median (..., outtype)` returns the median with a specified data type, using any of the input arguments in the previous syntaxes. *outtype* can take the following values:

`"default"`
: Output is of type double, unless the input is single in which case the output is of type single.

`"double"`
: Output is of type double.

`"native"`.
: Output is of the same type as the input (`class (x)`), unless the input is logical in which case the output is of type double.

The optional variable *nanflag* specifies whether to include or exclude NaN values from the calculation using any of the previously specified input argument combinations. The default value for *nanflag* is `"includenan"` which keeps NaN values in the calculation. To exclude NaN values set the value of *nanflag* to `"omitnan"`. The output will still contain NaN values if *x* consists of all NaN values in the operating dimension.

**See also:** [mean], page 829, [mode], page 831, [movmedian], page 853.

```
m = mode (x)
m = mode (x, dim)
m = mode (x, vecdim)
m = mode (x, "all")
[m, f, c] = mode (...)
```

Compute the most frequently occurring value in the input data *x*.

`mode` determines the frequency of values along the first non-singleton dimension and returns the value with the highest frequency. If two, or more, values have the same frequency `mode` returns the smallest.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension of *x*, including any dimension exceeding `ndims (x)`, will return *x*.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all

dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored. If all dimensions in *vecdim* are greater than `ndims (x)`, then `mode` will return *x*.

Specifying the dimension as `"all"` will cause `mode` to operate on all elements of *x*, and is equivalent to `mode (x(:))`.

The return variable *f* is the number of occurrences of the mode in the dataset.

The cell array *c* contains all of the elements with the maximum frequency.

**See also:** [mean], page 829, [median], page 830.

Using just one number, such as the mean, to represent an entire data set may not give an accurate picture of the data. One way to characterize the fit is to measure the dispersion of the data. Octave provides several functions for measuring dispersion.

`[s, l] = bounds (x)`
`[s, l] = bounds (x, dim)`
`[s, l] = bounds (x, vecdim)`
`[s, l] = bounds (x, "all")`
`[s, l] = bounds (..., nanflag)`

Return the smallest and largest values of the input data *x*.

If *x* is a vector, then `bounds (x)` returns the smallest and largest values of the elements in *x* in *s* and *l*, respectively.

If *x* is a matrix, then `bounds (x)` returns the smallest and largest values for each column of *x* as row vectors *s* and *l*, respectively.

If *x* is an array, then `bounds (x)` computes the smallest and largest values along the first non-singleton dimension of *x*.

The data in *x* must be numeric. By default, any NaN values are ignored. The size of *s* and *l* is equal to the size of *x* except for the operating dimension, which becomes 1.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension of *x*, including any dimension exceeding `ndims (x)`, will return *x*.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored.

Specifying the dimension as `"all"` will cause `bounds` to operate on all elements of *x*, and is equivalent to `bounds (x(:))`.

The optional variable *nanflag* specifies whether to include or exclude NaN values from the calculation using any of the previously specified input argument combinations. The default value for *nanflag* is `"omitnan"` which does not include NaN values in the result. If the argument `"includenan"` is given, and there is a NaN present, then the result for both smallest (*s*) and largest (*l*) elements will be NaN.

Usage Note: The bounds are a quickly computed measure of the dispersion of a data set, but are less accurate than `iqr` if there are outlying data points.

**See also:** [range], page 833, [iqr], page 833, [mad], page 834, [std], page 836.

```
y = range (x)
y = range (x, dim)
y = range (x, vecdim)
y = range (x, "all")
y = range (..., nanflag)
```

Return the difference between the maximum and the minimum values of the input data *x*.

If *x* is a vector, then `range (x)` returns the difference between the maximum and minimum values of the elements in *x*.

If *x* is a matrix, then `range (x)` returns a row vector *y* with the difference between the maximum and minimum values for each column of *x*.

If *x* is an array, then `range (x)` computes the difference between the maximum and minimum values along the first non-singleton dimension of *x*.

The data in *x* must be numeric. By default, any NaN values are ignored. The size of *r* is equal to the size of *x* except for the operating dimension, which becomes 1.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension of *x*, including any dimension exceeding `ndims (x)`, will return *x*.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored.

Specifying the dimension as `"all"` will cause `range` to operate on all elements of *x*, and is equivalent to `range (x(:))`.

The optional variable *nanflag* specifies whether to include or exclude NaN values from the calculation using any of the previously specified input argument combinations. The default value for *nanflag* is `"omitnan"` which does not include NaN values in the result. If the argument `"includenan"` is given, and there is a NaN present, then the corresponding result will be NaN.

Usage Note: The range is a quickly computed measure of the dispersion of a data set, but is less accurate than `iqr` if there are outlying data points.

**See also:** [bounds], page 832, [iqr], page 833, [mad], page 834, [std], page 836.

```
r = iqr (x)
r = iqr (x, dim)
r = iqr (x, vecdim)
r = iqr (x, "all")
[r, q] = iqr (...)
```

Compute the interquartile range of the input data *x*.

The interquartile range is defined as the difference between the 75th and 25th percentile values of *x* calculated using

```
quantile (x, [0.25 , 0.75])
```

If *x* is a vector, then `iqr (x)` computes the interquartile range of the elements in *x*.

If *x* is a matrix, then `iqr (x)` returns a row vector with each element containing the interquartile range of the corresponding column in *x*.

If *x* is an array, then `iqr (x)` computes the interquartile range along the first non-singleton dimension of *x*.

The data in *x* must be numeric and any NaN values are ignored. The size of *r* is equal to the size of *x* except for the operating dimension, which becomes 1.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension of *x*, including any dimension exceeding `ndims (x)`, will return `zeros (size (x))`.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored.

Specifying the dimension as `"all"` will cause `iqr` to operate on all elements of *x*, and is equivalent to `iqr (x(:))`.

The optional output *q* contains the quantiles for the 25th and 75th percentile of the data.

Usage Note: As a measure of dispersion, the interquartile range is less affected by outliers than either `range` or `std`. The interquartile range of a scalar is necessarily `0`.

**See also:** [bounds], page 832, [mad], page 834, [range], page 833, [std], page 836, [prctile], page 844, [quantile], page 842.

`m = mad (x)`
`m = mad (x, opt)`
`m = mad (x, opt, dim)`
`m = mad (x, opt, vecdim)`
`m = mad (x, opt, "all")`

Compute the mean or median absolute deviation (MAD) of the elements of *x*.

The mean absolute deviation is defined as

```
mad = mean (abs (x - mean (x)))
```

The median absolute deviation is defined as

```
mad = median (abs (x - median (x)))
```

`mad` excludes NaN values from calculation similar to using the `omitnan` option in `mean` and `median`.

If *x* is a vector, then `mad (x)` returns the mean absolute deviation of the elements in *x*.

If *x* is a matrix, then `mad (x)` returns a row vector with each element containing the mean absolute deviation of the corresponding column in *x*.

If *x* is an array, then `mad (x)` computes the mean absolute deviation along the first non-singleton dimension of *x*.

The optional argument *opt* specifies whether mean or median absolute deviation is calculated. The default is `0` which corresponds to mean absolute deviation; a value

of `1` corresponds to median absolute deviation. Passing an empty input `[]` defaults to mean absolute deviation.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension of *x*, including any dimension exceeding `ndims (x)`, will return `zeros (size (x))`.

Specifying the dimension as *vecdim*, a vector of non-repeating dimensions, will return the `mad` over the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored.

Specifying the dimension as `"all"` will cause `mad` to operate on all elements of *x*, and is equivalent to `mad (x(:))`.

Usage Note: As a measure of dispersion, `mad` is less affected by outliers than `std`.

**See also:** [bounds], page 832, [range], page 833, [iqr], page 833, [std], page 836, [mean], page 829, [median], page 830.

`y = meansq (x)`
`y = meansq (x, dim)`
`y = meansq (x, vecdim)`
`y = meansq (x, "all")`
`y = meansq (..., nanflag)`

Compute the mean square of the input data *x*.

The mean square is defined as

$$\text{meansq}(x) = \frac{1}{N} \sum_{i=1}^{N} |x_i|^2$$

where $N$ is the number of elements of *x*.

If *x* is a vector, then `meansq (x)` returns the mean square of the elements in *x*.

If *x* is a matrix, then `meansq (x)` returns a row vector with each element containing the mean square of the corresponding column in *x*.

If *x* is an array, then `meansq (x)` computes the mean square along the first non-singleton dimension of *x*.

The data in *x* must be numeric. The size of *y* is equal to the size of *x* except for the operating dimension, which becomes 1.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension of *x*, including any dimension exceeding `ndims (x)`, will return `x.^2`.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored.

Specifying the dimension as `"all"` will cause `meansq` to operate on all elements of *x*, and is equivalent to `meansq (x(:))`.

The optional variable *nanflag* specifies whether to include or exclude NaN values from the calculation using any of the previously specified input argument combinations.

The default value for *nanflag* is `"includenan"` which keeps NaN values in the calculation. To exclude NaN values set the value of *nanflag* to `"omitnan"`. The output will still contain NaN values if *x* consists of all NaN values in the operating dimension.

**See also:** [rms], page 836, [var], page 838, [std], page 836, [moment], page 841.

`y = rms (x)`
`y = rms (x, dim)`
`y = rms (x, vecdim)`
`y = rms (x, "all")`
`y = rms (..., nanflag)`

Compute the root mean square of the input data *x*.

The root mean square is defined as

$$\mathrm{rms}(x) = \sqrt{\frac{1}{N}\sum_{i=1}^{N}|x_i|^2}$$

where $N$ is the number of elements of *x*.

If *x* is a vector, then `rms (x)` returns the root mean square of the elements in *x*.

If *x* is a matrix, then `rms (x)` returns a row vector with each element containing the root mean square of the corresponding column in *x*.

If *x* is an array, then `rms (x)` computes the root mean square along the first non-singleton dimension of *x*.

The data in *x* must be numeric. The size of *y* is equal to the size of *x* except for the operating dimension, which becomes 1.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension of *x*, including any dimension exceeding `ndims (x)`, will return *x*.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored.

Specifying the dimension as `"all"` will cause `rms` to operate on all elements of *x*, and is equivalent to `rms (x(:))`.

The optional variable *nanflag* specifies whether to include or exclude NaN values from the calculation using any of the previously specified input argument combinations. The default value for *nanflag* is `"includenan"` which keeps NaN values in the calculation. To exclude NaN values set the value of *nanflag* to `"omitnan"`. The output will still contain NaN values if *x* consists of all NaN values in the operating dimension.

**See also:** [meansq], page 835, [var], page 838, [std], page 836, [moment], page 841.

`s = std (x)`
`s = std (x, w)`
`s = std (x, w, dim)`
`s = std (x, w, vecdim)`

`s = std (x, w, "all")`
`s = std (..., nanflag)`
`[s, m] = std (...)`

Compute the standard deviation of the elements of *x*.

The standard deviation is defined as

$$\mathrm{std}(x) = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N}(x_i - \bar{x})^2}$$

where $\bar{x}$ is the mean value of *x* and $N$ is the number of elements of *x*.

If *x* is a vector, then `std (x)` returns the standard deviation of the elements in *x*.

If *x* is a matrix, then `std (x)` returns a row vector with each element containing the standard deviation of the corresponding column in *x*.

If *x* is an array, then `std (x)` computes the standard deviation along the first non-singleton dimension of *x*.

The optional argument *w* determines the weighting scheme to use. Valid values are:

0 [default]:
: Normalize with $N-1$ (population standard deviation). This provides the square root of the best unbiased estimator of the standard deviation.

1:
: Normalize with $N$ (sample standard deviation). This provides the square root of the second moment around the mean.

a vector:
: Compute the weighted standard deviation with non-negative weights. The length of *w* must equal the size of *x* in the operating dimension. NaN values are permitted in *w*, will be multiplied with the associated values in *x*, and can be excluded by the *nanflag* option.

an array:
: Similar to vector weights, but *w* must be the same size as *x*. If the operating dimension is supplied as *vecdim* or `"all"` and *w* is not a scalar, *w* must be an same-sized array.

Note: *w* must always be specified before specifying any of the following dimension options. To use the default value for *w* you may pass an empty input argument `[]`.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension of *x*, including any dimension exceeding `ndims (x)`, will return `zeros (size (x))`.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored.

Specifying the dimension as `"all"` will cause `std` to operate on all elements of *x*, and is equivalent to `std (x(:))`.

The optional variable *nanflag* specifies whether to include or exclude NaN values from the calculation using any of the previously specified input argument combinations.

The default value for *nanflag* is `"includenan"` which keeps NaN values in the calculation. To exclude NaN values set the value of *nanflag* to `"omitnan"`. The output will still contain NaN values if *x* consists of all NaN values in the operating dimension.

The optional second output variable *m* contains the mean of the elements of *x* used to calculate the standard deviation. If *v* is the weighted standard deviation, then *m* is also the weighted mean.

**See also:** [var], page 838, [bounds], page 832, [mad], page 834, [range], page 833, [iqr], page 833, [mean], page 829, [median], page 830.

In addition to knowing the size of a dispersion it is useful to know the shape of the data set. For example, are data points massed to the left or right of the mean? Octave provides several common measures to describe the shape of the data set. Octave can also calculate moments allowing arbitrary shape measures to be developed.

```
v = var (x)
v = var (x, w)
v = var (x, w, dim)
v = var (x, w, vecdim)
v = var (x, w, "all")
v = var (..., nanflag)
[v, m] = var (...)
```

Compute the variance of the elements of *x*.

The variance is defined as

$$\mathrm{var}(x) = \frac{1}{N-1}\sum_{i=1}^{N}(x_i - \bar{x})^2$$

where $\bar{x}$ is the mean value of *x* and *N* is the number of elements of *x*.

If *x* is a vector, then `var (x)` returns the variance of the elements in *x*.

If *x* is a matrix, then `var (x)` returns a row vector with each element containing the variance of the corresponding column in *x*.

If *x* is an array, then `var (x)` computes the variance along the first non-singleton dimension of *x*.

The optional argument *w* determines the weighting scheme to use. Valid values are:

0 [default]:
: Normalize with $N-1$ (population variance). This provides the square root of the best unbiased estimator of the variance.

1:
: Normalize with *N* (sample variance). This provides the square root of the second moment around the mean.

a vector:
: Compute the weighted variance with non-negative weights. The length of *w* must equal the size of *x* in the operating dimension. NaN values are permitted in *w*, will be multiplied with the associated values in *x*, and can be excluded by the *nanflag* option.

an array: Similar to vector weights, but *w* must be the same size as *x*. If the operating dimension is supplied as *vecdim* or `"all"` and *w* is not a scalar, then *w* must match the size of the specified array slice.

Note: *w* must always be specified before specifying any of the following dimension options. To use the default value for *w* you may pass an empty input argument `[]`.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension of *x*, including any dimension exceeding `ndims (x)`, will return `zeros (size (x))`.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored.

Specifying the dimension as `"all"` will cause `var` to operate on all elements of *x*, and is equivalent to `var (x(:))`.

The optional variable *nanflag* specifies whether to include or exclude NaN values from the calculation using any of the previously specified input argument combinations. The default value for *nanflag* is `"includenan"` which keeps NaN values in the calculation. To exclude NaN values set the value of *nanflag* to `"omitnan"`. The output will still contain NaN values if *x* consists of all NaN values in the operating dimension.

The optional second output variable *m* contains the mean of the elements of *x* used to calculate the variance. If *v* is the weighted variance, then *m* is also the weighted mean.

**See also:** [std], page 836, [mean], page 829, [cov], page 865, [skewness], page 839, [kurtosis], page 840, [moment], page 841.

```
y = skewness (x)
y = skewness (x, flag)
y = skewness (x, flag, dim)
y = skewness (x, flag, vecdim)
y = skewness (x, flag, "all")
```

Compute the sample skewness of the input data *x*.

The sample skewness is defined as

$$\text{skewness}(x) = \frac{\frac{1}{N}\sum_{i=1}^{N}(x_i - \bar{x})^3}{\sigma^3},$$

where $N$ is the length of *x*, $\bar{x}$ its mean and $\sigma$ its (uncorrected) standard deviation.

The optional argument *flag* controls which normalization is used. If *flag* is equal to 1 (default value, used when *flag* is omitted or empty), return the sample skewness as defined above. If *flag* is equal to 0, return the adjusted skewness coefficient instead:

$$\text{skewness}(x) = \frac{\sqrt{N(N-1)}}{N-2} \times \frac{\frac{1}{N}\sum_{i=1}^{N}(x_i - \bar{x})^3}{\sigma^3}$$

The adjusted skewness coefficient is obtained by replacing the sample second and third central moments by their bias-corrected versions.

If $x$ is a vector, then `skewness (x)` computes the skewness of the data in $x$.

If $x$ is a matrix, then `skewness (x)` returns a row vector with each element containing the skewness of the data of the corresponding column in $x$.

If $x$ is an array, then `skewness (x)` computes the skewness of the data along the first non-singleton dimension of $x$.

The data in $x$ must be numeric and any NaN values are ignored. The size of $y$ is equal to the size of $x$ except for the operating dimension, which becomes 1.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension of $x$, including any dimension exceeding `ndims (x)`, will return $x$.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of $x$, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored.

Specifying the dimension as `"all"` will cause `skewness` to operate on all elements of $x$, and is equivalent to `skewness (x(:))`.

**See also:** [var], page 838, [kurtosis], page 840, [moment], page 841.

`y = kurtosis (x)`
`y = kurtosis (x, flag)`
`y = kurtosis (x, flag, dim)`
`y = kurtosis (x, flag, vecdim)`
`y = kurtosis (x, flag, "all")`

Compute the sample kurtosis of the input data $x$.

The sample kurtosis is defined as

$$\kappa_1 = \frac{\frac{1}{N}\sum_{i=1}^{N}(x_i - \bar{x})^4}{\sigma^4},$$

where $N$ is the length of $x$, $\bar{x}$ its mean, and $\sigma$ its (uncorrected) standard deviation.

The optional argument *flag* controls which normalization is used. If *flag* is equal to 1 (default value, used when *flag* is omitted or empty), return the sample kurtosis as defined above. If *flag* is equal to 0, return the "bias-corrected" kurtosis coefficient instead:

$$\kappa_0 = 3 + \frac{N-1}{(N-2)(N-3)}\left((N+1)\,\kappa_1 - 3(N-1)\right)$$

The bias-corrected kurtosis coefficient is obtained by replacing the sample second and fourth central moments by their unbiased versions. It is an unbiased estimate of the population kurtosis for normal populations.

If $x$ is a vector, then `kurtosis (x)` computes the kurtosis of the data in $x$.

If $x$ is a matrix, then `kurtosis (x)` returns a row vector with each element containing the kurtosis of the data of the corresponding column in $x$.

If $x$ is an array, then `kurtosis (x)` computes the kurtosis of the data along the first non-singleton dimension of $x$.

The data in *x* must be numeric and any NaN values are ignored. The size of *y* is equal to the size of *x* except for the operating dimension, which becomes 1.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension of *x*, including any dimension exceeding `ndims (x)`, will return *x*.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored.

Specifying the dimension as `"all"` will cause `kurtosis` to operate on all elements of *x*, and is equivalent to `kurtosis (x(:))`.

**See also:** [var], page 838, [skewness], page 839, [moment], page 841.

`m = moment (x, p)`
`m = moment (x, p, dim)`
`m = moment (x, p, vecdim)`
`m = moment (x, p, "all")`
`m = moment (x, p, ..., type)`

Compute the *p*-th central moment of the input data *x*.

The *p*-th central moment of *x* is defined as:

$$\frac{1}{N}\sum_{i=1}^{N}(x_i - \bar{x})^p$$

where $\bar{x}$ is the mean value of *x* and *N* is the number of elements of *x*.

If *x* is a vector, then `moment (x)` computes the *p*-th central moment of the data in *x*.

If *x* is a matrix, then `moment (x)` returns a vector with element containing the *p*-th central moment of the corresponding column in *x*.

If *x* is an array, then `moment (x)` computes the *p*-th central moment along the first non-singleton dimension of *x*.

The data in *x* must be a non-empty numeric array and any NaN values along the operating dimension will return NaN for central moment. The size of *m* is equal to the size of *x* except for the operating dimension, which becomes 1.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension of *x*, including any dimension exceeding `ndims (x)`, will return *x*.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored. If all dimensions in *vecdim* are greater than `ndims (x)`, then `moment` will return *x*.

Specifying the dimension as `"all"` will cause `moment` to operate on all elements of *x*, and is equivalent to `moment (x(:))`.

The optional fourth input argument, *type*, is a string specifying the type of moment to be computed. Valid options are:

`"c"` Central Moment (default).

`"a"`

`"ac"` Absolute Central Moment. The moment about the mean ignoring sign defined as

$$\frac{1}{N}\sum_{i=1}^{N}|x_i - \bar{x}|^p$$

`"r"` Raw Moment. The moment about zero defined as

$$\frac{1}{N}\sum_{i=1}^{N}{x_i}^p$$

`"ar"` Absolute Raw Moment. The moment about zero ignoring sign defined as

$$\frac{1}{N}\sum_{i=1}^{N}|x_i|^p$$

**See also:** [var], page 838, [skewness], page 839, [kurtosis], page 840.

`q = quantile (x)`
`q = quantile (x, p)`
`q = quantile (x, n)`
`q = quantile (x, ..., dim)`
`q = quantile (x, ..., vecdim)`
`q = quantile (x, ..., "all")`
`q = quantile (x, p, ..., method)`
`q = quantile (x, n, ..., method)`

Compute the quantiles of the input data *x*.

If *x* is a vector, then `quantile (x)` computes the quantiles specified by *p* of the data in *x*.

If *x* is a matrix, then `quantile (x)` returns a matrix such that the i-th row of *q* contains the *p*(i)th quantiles of each column of *x*.

If *x* is an array, then `quantile (x)` computes the quantiles specified by *p* along the first non-singleton dimension of *x*.

The data in *x* must be numeric and any NaN values are ignored. The size of *q* is equal to the size of *x* except for the operating dimension, which equals to the number of quantiles specified by *p* or *n*.

*p* is a numeric vector specifying the percentiles to be computed, which correspond to the cumulative probabilities of the data . All elements of *p* must be in the range from 0 to 1. If *p* is unspecified, return the percentiles for `[0.00 0.25 0.50 0.75 1.00]`. Alternatively, the second input argument may be specified as a positive integer value *n*, in which case `quantile` returns the quantiles for *n* evenly spaced cumulative probabilities computed as $(1/(n + 1), 2/(n + 1), \ldots, n/(n + 1))$ for `n > 1`.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension of *x*, including any dimension exceeding `ndims (x)`, will return N copies of *x* along the operating dimension, where N is the number of specified quantiles.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored. If all dimensions in *vecdim* are greater than `ndims (x)`, then `quantile` will return N copies of *x* along the smallest dimension in *vecdim*.

Specifying the dimension as `"all"` will cause `iqr` to operate on all elements of *x*, and is equivalent to `iqr (x(:))`.

The fourth input argument, *methods*, determines the method to calculate the quantiles specified by *p* or *n*. The methods available to calculate sample quantiles are the nine methods used by R (https://www.r-project.org/) and can be specified by the corresponding integer value. The default value is *method* = 5.

Discontinuous sample quantile methods 1, 2, and 3

1. Method 1: Inverse of empirical distribution function.
2. Method 2: Similar to method 1 but with averaging at discontinuities.
3. Method 3: SAS definition: nearest even order statistic.

Continuous sample quantile methods 4 through 9, where $p(k)$ is the linear interpolation function respecting each method's representative cdf.

4. Method 4: $p(k) = k/N$. That is, linear interpolation of the empirical cdf, where $N$ is the length of $P$.
5. Method 5: $p(k) = (k - 0.5)/N$. That is, a piecewise linear function where the knots are the values midway through the steps of the empirical cdf.
6. Method 6: $p(k) = k/(N + 1)$.
7. Method 7: $p(k) = (k - 1)/(N - 1)$.
8. Method 8: $p(k) = (k - 1/3)/(N + 1/3)$. The resulting quantile estimates are approximately median-unbiased regardless of the distribution of $x$.
9. Method 9: $p(k) = (k - 3/8)/(N + 1/4)$. The resulting quantile estimates are approximately unbiased for the expected order statistics if $x$ is normally distributed.

Hyndman and Fan (1996) recommend method 8. Maxima, S, and R (versions prior to 2.0.0) use 7 as their default. Minitab and SPSS use method 6. MATLAB uses method 5.

References:


- R. A. Becker, J. M. Chambers, and A. R. Wilks, *The New S Language*, Wadsworth & Brooks/Cole, 1988.
- R. J. Hyndman, and Y. Fan, "Sample quantiles in statistical packages", *American Statistician*, 50, pp. 361–365, 1996.
- *R: A Language and Environment for Statistical Computing*, https://cran.r-project.org/doc/manuals/fullrefman.pdf.


Examples:

```
x = randi (1000, [10, 1]);  # Create empirical data in range 1-1000
q = quantile (x, [0, 1]);   # Return minimum, maximum of distribution
q = quantile (x, [0.25 0.5 0.75]); # Return quartiles of distribution
```

**See also:** [prctile], page 844.

`q = prctile (x)`
`q = prctile (x, p)`
`q = prctile (x, p, dim)`
`q = prctile (x, p, vecdim)`
`q = prctile (x, p, "all")`
`q = prctile (x, p, ..., method)`

Compute the percentiles of the input data *x*.

If *x* is a vector, then `prctile (x)` computes the percentiles specified by *p* of the data in *x*.

If *x* is a matrix, then `prctile (x)` returns a matrix such that the i-th row of *q* contains the *p*(i)th percentiles of each column of *x*.

If *x* is an array, then `prctile (x)` computes the percentiles specified by *p* along the first non-singleton dimension of *x*.

The data in *x* must be numeric and any NaN values are ignored. The size of *q* is equal to the size of *x* except for the operating dimension, which equals to the number of quantiles specified by *p*.

*p* is a numeric vector specifying the percentiles to be computed. All elements of *p* must be in the range from 0 to 100. If *p* is unspecified, return the percentiles for `[0 25 50 75 100]`.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension of *x*, including any dimension exceeding `ndims (x)`, will return N copies of *x* along the operating dimension, where N is the number of specified percentiles.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored. If all dimensions in *vecdim* are greater than `ndims (x)`, then `quantile` will return N copies of *x* along the smallest dimension in *vecdim*.

Specifying the dimension as `"all"` will cause `iqr` to operate on all elements of *x*, and is equivalent to `iqr (x(:))`.

The fourth input argument, *methods*, determines the method to calculate the percentiles specified by *p*. The methods available to calculate sample percentiles are the nine methods used by R (https://www.r-project.org/) and can be specified by the corresponding integer value. The default value is *method* = 5.

Discontinuous sample quantile methods 1, 2, and 3

1. Method 1: Inverse of empirical distribution function.
2. Method 2: Similar to method 1 but with averaging at discontinuities.
3. Method 3: SAS definition: nearest even order statistic.

Continuous sample quantile methods 4 through 9, where $p(k)$ is the linear interpolation function respecting each method's representative cdf.

4. Method 4: $p(k) = k/N$. That is, linear interpolation of the empirical cdf, where $N$ is the length of $P$.

5. Method 5: $p(k) = (k - 0.5)/N$. That is, a piecewise linear function where the knots are the values midway through the steps of the empirical cdf.
6. Method 6: $p(k) = k/(N + 1)$.
7. Method 7: $p(k) = (k - 1)/(N - 1)$.
8. Method 8: $p(k) = (k - 1/3)/(N + 1/3)$. The resulting quantile estimates are approximately median-unbiased regardless of the distribution of $x$.
9. Method 9: $p(k) = (k - 3/8)/(N + 1/4)$. The resulting quantile estimates are approximately unbiased for the expected order statistics if $x$ is normally distributed.

**See also:** [quantile], page 842.

A summary view of a data set can be generated quickly with the `statistics` function.

```
stats = statistics (x)
stats = statistics (x, dim)
stats = statistics (x, vecdim)
stats = statistics (x, "all")
stats = statistics (..., nanflag)
```

Return a vector with statistics parameters over the input data *x*.

`statistics (x` operates along the first non-singleton dimension of *x* and calculates the following statistical parameters:

1. minimum
2. first quartile
3. median
4. third quartile
5. maximum
6. mean
7. standard deviation
8. skewness
9. kurtosis

If *x* is a row vector, then `statistics (x)` returns a row vector with the aforementioned statistical parameters. If *x* is a column vector, then it returns a column vector.

If *x* is a matrix, then `statistics (x)` returns a matrix such that each column contains the statistical parameters calculated over the corresponding column of *x*.

If *x* is an array, then `statistics (x)` computes the statistical parameters along the first non-singleton dimension of *x*.

The data in *x* must be numeric and by default any NaN values are ignored from the computations of statistical parameters except for the mean and the standard deviation. Set the optional argument *nanflag* to `"omitnan"` to exclude the NaN values from the calculation of the mean and standard deviation parameters. Setting *nanflag* to `"includenan"` is ignored and it is equivalent to calling the `statistics` function without the *nanflag* argument.

The size of *stats* is equal to the size of *x* except for the operating dimension, which equals to 9 (i.e., the number of statistical parameters returned).

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension of *x*, including any dimension exceeding `ndims (x)`, will return *x*.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored.

Specifying the dimension as `"all"` will cause `statistics` to operate on all elements of *x*, and is equivalent to `statistics (x(:))`.

**See also:** [min], page 618, [max], page 617, [median], page 830, [mean], page 829, [std], page 836, [skewness], page 839, [kurtosis], page 840.

## 26.2 Statistics on Sliding Windows of Data

It is often useful to calculate descriptive statistics over a subsection (i.e., window) of a full dataset. Octave provides the function `movfun` which will call an arbitrary function handle with windows of data and accumulate the results. Many of the most commonly desired functions, such as the moving average over a window of data (`movmean`), are already provided.

`y = movfun (fcn, x, wlen)`
`y = movfun (fcn, x, [nb, na])`
`y = movfun (..., "property", value)`

Apply function *fcn* to a moving window of length *wlen* on data *x*.

The moving window length input *wlen* can either be a numeric scalar or a 2-element numeric array `[nb, na]`. The elements included in the moving window depend on the size and value of *wlen* as well as whether the `"SamplePoints"` option has been specified. For full details of element inclusion, see [movslice], page 848.

During calculations the data input *x* is reshaped into a 2-dimensional *wlen*-by-*N* matrix and *fcn* is called on this new matrix. Therefore, *fcn* must accept an array input argument and apply the computation along dimension 1, i.e., down the columns of the array.

When applied to an array (possibly multi-dimensional) with *n* columns, *fcn* may return a result in either of two formats: Format 1) an array of size 1-by-*n*-by-*dim3*-by-...-by-*dimN*. This is the typical output format from Octave core functions. Type `demo ("movfun", 5)` for an example of this use case. Format 2) a row vector of length `n * numel_higher_dims` where *numel_higher_dims* is `prod (size (x)(3:end))`. The output of *fcn* for the i-th input column must be found in the output at indices `i:n:(n*numel_higher_dims)`. This format is useful when concatenating functions into arrays, or when using `nthargout`. Type `demo ("movfun", 6)` for an example of this case.

The calculation can be controlled by specifying *property*/*value* pairs. Valid properties are

`"dim"` Operate along the dimension specified, rather than the default of the first non-singleton dimension.

`"SamplePoints"`
This property specifies a sorted, numeric vector of unique coordinate positions of the data points in *x*. The default value is the vector `[1 : numel (x)]`. When a non-default SamplePoints vector is specified, the moving window length *wlen* is measured against the SamplePoints positions to determine which points are included in each window slice. SamplePoints need not be uniformly spaced. This can result in window slices containing different numbers of points.

`"Endpoints"`
This property controls how results are calculated at the boundaries (endpoints) of the window. Possible values are:

`"shrink"` (default)
The window is truncated at the beginning and end of the array to exclude elements for which there is no source data. For example, with a window of length 3, `y(1) = fcn (x(1:2))`, and `y(end) = fcn (x(end-1:end))`.

`"discard"`
Any *y* values that use a window extending beyond the original data array are deleted. For example, with a 10-element data vector and a window of length 3, the output will contain only 8 elements. The first element would require calculating the function over indices `[0, 1, 2]` and is therefore discarded. The last element would require calculating the function over indices `[9, 10, 11]` and is therefore discarded.

`"fill"` Any window elements outside the data array are replaced by `NaN`. For example, with a window of length 3, `y(1) = fcn ([NaN, x(1:2)])`, and `y(end) = fcn ([x(end-1:end), NaN])`. This option usually results in *y* having `NaN` values at the boundaries, although it is influenced by how *fcn* handles `NaN`, and also by the property `"nancond"`.

*user_value* Any window elements outside the data array are replaced by the specified value *user_value* which must be a numeric scalar. For example, with a window of length 3, `y(1) = fcn ([user_value, x(1:2)])`, and `y(end) = fcn ([x(end-1:end), user_value])`. A common choice for *user_value* is 0.

`"same"` Any window elements outside the data array are replaced by the value of *x* at the boundary. For example, with a window of length 3, `y(1) = fcn ([x(1), x(1:2)])`, and `y(end) = fcn ([x(end-1:end), x(end)])`.

`"periodic"`
: The window is wrapped so that any missing data elements are taken from the other side of the data. For example, with a window of length 3, `y(1) = fcn ([x(end), x(1:2)])`, and `y(end) = fcn ([x(end-1:end), x(1)])`.

Note 1: For non-uniform SamplePoint spacing, the only permitted value for `"EndPoints"` is `"shrink"`.

Note 2: For some `"Endpoints"` options, the window size at the boundaries may not be the same as for the central part, and *fcn* must work in these cases.

`"nancond"`
: Controls haw `NaN` and `NA` values affect the output of `"movfun"`. The value `"includenan"` (default) causes `NaN` and `NA` values to be included in the moving window, and any window slice containing `NaN` or `NA` values will return `NaN` for that element. The value `"omitnan"` causes `"movfun"` to ignore any `NaN` or `NA` values resulting in fewer elements being used to calculate the result for that window slice. If `"omitnan"` is specified and a window slice contains all `NaN` or `NA` values, `"movfun"` returns the value specified by the `"nanval"` property for that element. The values `"includemissing"` and `"omitmissing"` may be used synonymously with `"includenan"` and `"omitnan"`, respectively.

`"nanval"`
: Specifies the value to return when `"nancond"` is set to `"omitnan"` or `"omitmissing"` and all elements in a window are `NaN` or `NA`. `"nanval"` must be a numeric scalar value or `NaN` (default).

`"outdim"`
: A row vector that selects which dimensions of the calculation will appear in the output *y*. This is only useful when *fcn* returns an N-dimensional array in Format 1. The default is to return all output dimensions.

Programming Note: The property `"outdim"` can be used to save memory when the output of *fcn* has many dimensions, or when a wrapper to the base function that selects the desired outputs is too costly. When memory is not an issue, the easiest way to select output dimensions is to first calculate the complete result with `movfun` and then filter that result with indexing. If code complexity is not an issue then a wrapper can be created using anonymous functions. For example, if `basefcn` is a function returning a $K$-dimensional row output, and only dimension $D$ is desired, then the following wrapper could be used.

```
fcn = @(x) basefcn (x)(:,columns(x) * (D-1) + (1:columns(x)));
y = movfun (@fcn, ...);
```

**See also:** [movslice], page 848, [prepad], page 579, [postpad], page 579, [permute], page 572, [reshape], page 573.

`slcidx = movslice (N, wlen)`
`slcidx = movslice (N, wlen, samplepoints)`
`[slcidx, C, Cpre, Cpost, win, wlen, scalar_wlen] = movslice (...)`
: Generate indices to slice a vector of length $N$ into windows of length *wlen*.

The input *N* must be a positive integer.

The moving window length input *wlen* can either be a numeric scalar or a 2-element numeric array. The elements included in the moving window will depend on the size and value of *wlen* as well as whether the *samplepoints* input was specified.

The optional input *samplepoints* is a sorted, numeric vector of unique positions of the *N* data points. The default value is the vector `[1 : N]`. When a non-default *samplepoints* vector is specified, the moving window length *wlen* is measured against the *samplepoints* positions to determine which points are included in each window slice. It should be noted that *samplepoints* need not be uniformly spaced which can result in window slices containing different numbers of points. Because of this, as specified below the shape and content of some `movslice` outputs will be different when a non-default *samplepoints* is used.

The moving window size and included elements will be defined as follows:

- If *samplepoints* has the default value of `1:N` (or has not been specified):
  - For integer-valued *wlen*:
    - For odd, integer-valued, scalar *wlen* the window is symmetric and includes `(wlen - 1) / 2` elements on either side of the central element. For example, the window slice at index 5 with a window length of 3 will include the elements `[4, 5, 6]`.
    - For even, integer-valued, scalar *wlen* the window is asymmetric and has `wlen/2` elements to the left of the central element and `wlen/2 - 1` elements to the right of the central element. For example, the window slice at index 5 with a window length of 4 will include the elements `[3, 4, 5, 6]`.
    - For integer-valued vector *wlen* of the form `[nb, na]` where *nb* and *na* are integer valued the window includes *nb* elements to the left of the central element and *na* elements to the right of the central element. For example, given `wlen = [3, 1]`, the window slice at index 5 will include the elements `[2, 3, 4, 5, 6]`.
  - For non-integer-valued scalar *wlen*:
    - Non-integer-valued scalar *wlen* will be converted to two-element vector form with `nb = na = fix (wlen / 2)`, and then processed as stated above for integer-valued vectors. For example, the window slice at index 5 with `wlen = 2.5` will include the elements `[3, 4, 5, 6, 7]`.
    - Non-integer-valued vector *wlen* will be truncated to integer values with `wlen = fix (wlen)` and then processed as stated above for integer-valued vectors. For example, the window slice at index 5 with `wlen = [1.2, 2.3]` will include the elements `[4, 5, 6, 7]`.
- If *samplepoints* has been specified with a non-default vector:
  - For vector *wlen* specified as `[nb, na]`, the window will include all points within a distance less than or equal to *nb* before and *na* after the central element's position, with point positions defined by the elements of *samplepoints*. For example, at index 5 with `wlen = [2, 3]` and the 3rd-8th elements of *samplepoints* being `[1, 3, 5, 7, 8 ,9]`, the window

slice will include the elements `[4, 5, 6, 7]` corresponding to *samplepoints* `[3, 5, 7, 8]`.

- Scalar *wlen* will be converted to two-element vector form with *nb* = *na* = *wlen* / 2. The window will then include all points within a distance of less than or equal to *nb* before and less than, but not equal to, *na* after the central element's position, [*nb*, *na*). For example, at index 5 with *wlen* = `[2, 3]` and the 3rd-8th elements of *samplepoints* being `[1, 3, 5, 7, 8 ,9]`, the window slice will include the elements `[4, 5, 6]` corresponding to *samplepoints* `[3, 5, 7]`.

The output *slcidx* is an array of indices of the slices of the vector.

- If *samplepoints* is default or unspecified, *slcidx* will contain only the indices of the slices that fit fully within the vector. Each column will be the indices of one slice as the window moves from left to right. The slices will have `fix (`*wlen*`)` elements for scalar *wlen*, or *nb* + *na* + 1 elements for array valued *wlen*.
- If a non-default *samplepoints* has been specified, *slcidx* will be a 2x$N$ array with the first and second rows containing the first and last elements of each slice, respectively.

Optional output $C$ is a row vector of window center positions where the window stays fully within the vector.

Optional outputs *Cpre* and *Cpost* contain the vector elements at the start and end of the vector, respectively, that result in the window extending beyond the ends of the vector.

Optional output *win* contains information for creating the moving window.

- If *samplepoints* is default or unspecified, *win* is a column vector with the same number of rows as *slcidx* that contains the moving window defined as a center relative position stencil.
- If a non-default *samplepoints* has been specified, *win* will be a 2x$N$ array with the first and second rows containing the left and right bounds of each window slice, respectively, using the same coordinates as *samplepoints*. These bounds may lie outside of the position vector specified by *samplepoints*.

Optional output *wlen* returns the window length used by `movslice` in two-element [*nb*, *na*] form.

Optional logical output *scalar_wlen* returns the scalar or vector state of the input *wlen* so that calling functions can determinine whether the moving window should be inclusive or exclusive of the right window endpoints. I.e., inclusive [*nb*, *na*] for vector *wlen* or exclusive [*nb*, *na*) for scalar *wlen*.

**See also:** [movfun], page 846.

*y* = movmad (*x*, *wlen*)
*y* = movmad (*x*, [*nb*, *na*])
*y* = movmad (..., *dim*)
*y* = movmad (..., *nancond*)

`y = movmad (...,` *`property`*`,` *`value`*`)`

Calculate the moving median or mean absolute deviation over a sliding window of length *wlen* on data *x*.

The moving window length input *wlen* can either be a numeric scalar or a 2-element numeric array `[`*`nb`*`,` *`na`*`]`. The elements included in the moving window depend on the size and value of *wlen* as well as whether the `"SamplePoints"` option has been specified. For full details of element inclusion, see [movslice], page 848.

If the optional argument *dim* is given, operate along this dimension.

The optional argument *nancond* is a string that controls how `NaN` and `NA` values affect the output of `"movmad"`. The value `"includenan"` causes `NaN` and `NA` values to be included in the moving window, and any window slice containing `NaN` or `NA` values will return `NaN` for that element. The value `"omitnan"` (default) causes `"movmad"` to ignore any `NaN` or `NA` values resulting in fewer elements being used to calculate the mad for that window slice. If `"omitnan"` is specified and a window slice contains all `NaN` or `NA` values, `"movmad"` returns `NaN` for that element. The values `"includemissing"` and `"omitmissing"` may be used synonymously with `"includenan"` and `"omitnan"`, respectively.

The calculation can be controlled by specifying *property*/*value* pairs:

- The `"method"` property can take the value `"median"` (default) or `"mean"` to control whether `"movmad"` performs median or mean absolute deviation calculations on the data.
- Additional valid properties are `"Endpoints"` and `"SamplePoints"`. For full descriptions of these properties and valid options, see [movfun], page 846.

Programming Note: This function is a wrapper which calls `movfun`. For full documentation of inputs and options, see [movfun], page 846.

Compatibility Note: Prior to Octave 10 this function only calculated mean absolute deviation. For MATLAB compatibility, the default has been changed to median absolute deviation. The `"method"` property is now provided to enable access to both `"mad"` calculation methods. This property should not be expected to be functional outside of Octave code.

**See also:** [mad], page 834, [movfun], page 846, [movslice], page 848, [movmax], page 851, [movmean], page 852, [movmedian], page 853, [movmin], page 853, [movprod], page 854, [movstd], page 855, [movsum], page 855, [movvar], page 856.

`y = movmax (`*`x`*`,` *`wlen`*`)`
`y = movmax (`*`x`*`, [`*`nb`*`,` *`na`*`])`
`y = movmax (...,` *`dim`*`)`
`y = movmax (...,` *`nancond`*`)`
`y = movmax (...,` *`property`*`,` *`value`*`)`

Calculate the moving maximum over a sliding window of length *wlen* on data *x*.

The moving window length input *wlen* can either be a numeric scalar or a 2-element numeric array `[`*`nb`*`,` *`na`*`]`. The elements included in the moving window depend on the size and value of *wlen* as well as whether the `"SamplePoints"` option has been specified. For full details of element inclusion, see [movslice], page 848.

If the optional argument *dim* is given, operate along this dimension.

The optional argument *nancond* is a string that controls how `NaN` and `NA` values affect the output of `"movmax"`. The value `"includenan"` causes `NaN` and `NA` values to be included in the moving window, and any window slice containing `NaN` or `NA` values will return `NaN` for that element. The value `"omitnan"` (default) causes `"movmax"` to ignore any `NaN` or `NA` values resulting in fewer elements being used to calculate the maximum for that window slice. If `"omitnan"` is specified and a window slice contains all `NaN` or `NA` values, `"movmax"` returns `NaN` for that element. The values `"includemissing"` and `"omitmissing"` may be used synonymously with `"includenan"` and `"omitnan"`, respectively.

The calculation can be controlled by specifying *property*/*value* pairs. Valid properties are `"Endpoints"` and `"SamplePoints"`. For full descriptions of these properties and valid options, see [movfun], page 846.

Programming Note: This function is a wrapper which calls `movfun`. For full documentation of inputs and options, see [movfun], page 846.

**See also:** [max], page 617, [movfun], page 846, [movslice], page 848, [movmad], page 850, [movmean], page 852, [movmedian], page 853, [movmin], page 853, [movprod], page 854, [movstd], page 855, [movsum], page 855, [movvar], page 856.

`y = movmean (x, wlen)`
`y = movmean (x, [nb, na])`
`y = movmean (..., dim)`
`y = movmean (..., nancond)`
`y = movmean (..., property, value)`

Calculate the moving average over a sliding window of length *wlen* on data *x*.

The moving window length input *wlen* can either be a numeric scalar or a 2-element numeric array `[nb, na]`. The elements included in the moving window depend on the size and value of *wlen* as well as whether the `"SamplePoints"` option has been specified. For full details of element inclusion, see [movslice], page 848.

If the optional argument *dim* is given, operate along this dimension.

The optional argument *nancond* is a string that controls how `NaN` and `NA` values affect the output of `"movmean"`. The value `"includenan"` (default) causes `NaN` and `NA` values to be included in the moving window, and any window slice containing `NaN` or `NA` values will return `NaN` for that element. The value `"omitnan"` causes `"movmean"` to ignore any `NaN` or `NA` values resulting in fewer elements being used to calculate the mean for that window slice. If `"omitnan"` is specified and a window slice contains all `NaN` or `NA` values, `"movmean"` returns `NaN` for that element. The values `"includemissing"` and `"omitmissing"` may be used synonymously with `"includenan"` and `"omitnan"`, respectively.

The calculation can be controlled by specifying *property*/*value* pairs. Valid properties are `"Endpoints"` and `"SamplePoints"`. For full descriptions of these properties and valid options, see [movfun], page 846.

Programming Note: This function is a wrapper which calls `movfun`. For full documentation of inputs and options, see [movfun], page 846.

**See also:** [mean], page 829, [movfun], page 846, [movslice], page 848, [movmad], page 850, [movmax], page 851, [movmedian], page 853, [movmin], page 853, [movprod], page 854, [movstd], page 855, [movsum], page 855, [movvar], page 856.

`y = movmedian (x, wlen)`
`y = movmedian (x, [nb, na])`
`y = movmedian (..., dim)`
`y = movmedian (..., nancond)`
`y = movmedian (..., property, value)`

Calculate the moving median over a sliding window of length *wlen* on data *x*.

The moving window length input *wlen* can either be a numeric scalar or a 2-element numeric array `[nb, na]`. The elements included in the moving window depend on the size and value of *wlen* as well as whether the `"SamplePoints"` option has been specified. For full details of element inclusion, see [movslice], page 848.

If the optional argument *dim* is given, operate along this dimension.

The optional argument *nancond* is a string that controls how `NaN` and `NA` values affect the output of `"movmedian"`. The value `"includenan"` (default) causes `NaN` and `NA` values to be included in the moving window, and any window slice containing `NaN` or `NA` values will return `NaN` for that element. The value `"omitnan"` causes `"movmedian"` to ignore any `NaN` or `NA` values resulting in fewer elements being used to calculate the median for that window slice. If `"omitnan"` is specified and a window slice contains all `NaN` or `NA` values, `"movmedian"` returns `NaN` for that element. The values `"includemissing"` and `"omitmissing"` may be used synonymously with `"includenan"` and `"omitnan"`, respectively.

The calculation can be controlled by specifying *property*/*value* pairs. Valid properties are `"Endpoints"` and `"SamplePoints"`. For full descriptions of these properties and valid options, see [movfun], page 846.

Programming Note: This function is a wrapper which calls `movfun`. For full documentation of inputs and options, see [movfun], page 846.

**See also:** [median], page 830, [movfun], page 846, [movslice], page 848, [movmad], page 850, [movmax], page 851, [movmean], page 852, [movmin], page 853, [movprod], page 854, [movstd], page 855, [movsum], page 855, [movvar], page 856.

`y = movmin (x, wlen)`
`y = movmin (x, [nb, na])`
`y = movmin (..., dim)`
`y = movmin (..., nancond)`
`y = movmin (..., property, value)`

Calculate the moving minimum over a sliding window of length *wlen* on data *x*.

The moving window length input *wlen* can either be a numeric scalar or a 2-element numeric array `[nb, na]`. The elements included in the moving window depend on the size and value of *wlen* as well as whether the `"SamplePoints"` option has been specified. For full details of element inclusion, see [movslice], page 848.

If the optional argument *dim* is given, operate along this dimension.

The optional argument *nancond* is a string that controls how `NaN` and `NA` values affect the output of `"movmin"`. The value `"includenan"` causes `NaN` and `NA` values to be

included in the moving window, and any window slice containing `NaN` or `NA` values will return `NaN` for that element. The value `"omitnan"` (default) causes `"movmin"` to ignore any `NaN` or `NA` values resulting in fewer elements being used to calculate the minimum for that window slice. If `"omitnan"` is specified and a window slice contains all `NaN` or `NA` values, `"movmin"` returns `NaN` for that element. The values `"includemissing"` and `"omitmissing"` may be used synonymously with `"includenan"` and `"omitnan"`, respectively.

The calculation can be controlled by specifying *property*/*value* pairs. Valid properties are `"Endpoints"` and `"SamplePoints"`. For full descriptions of these properties and valid options, see [movfun], page 846.

Programming Note: This function is a wrapper which calls `movfun`. For full documentation of inputs and options, see [movfun], page 846.

**See also:** [min], page 618, [movfun], page 846, [movslice], page 848, [movmad], page 850, [movmax], page 851, [movmean], page 852, [movmedian], page 853, [movprod], page 854, [movstd], page 855, [movsum], page 855, [movvar], page 856.

`y = movprod (x, wlen)`
`y = movprod (x, [nb, na])`
`y = movprod (..., dim)`
`y = movprod (..., nancond)`
`y = movprod (..., property, value)`

Calculate the moving product over a sliding window of length *wlen* on data *x*.

The moving window length input *wlen* can either be a numeric scalar or a 2-element numeric array `[nb, na]`. The elements included in the moving window depend on the size and value of *wlen* as well as whether the `"SamplePoints"` option has been specified. For full details of element inclusion, see [movslice], page 848.

If the optional argument *dim* is given, operate along this dimension.

The optional argument *nancond* is a string that controls how `NaN` and `NA` values affect the output of `"movprod"`. The value `"includenan"` (default) causes `NaN` and `NA` values to be included in the moving window, and any window slice containing `NaN` or `NA` values will return `NaN` for that element. The value `"omitnan"` causes `"movprod"` to ignore any `NaN` or `NA` values resulting in fewer elements being used to calculate the product for that window slice. If `"omitnan"` is specified and a window slice contains all `NaN` or `NA` values, `"movprod"` returns 1 for that element. The values `"includemissing"` and `"omitmissing"` may be used synonymously with `"includenan"` and `"omitnan"`, respectively.

The calculation can be controlled by specifying *property*/*value* pairs. Valid properties are `"Endpoints"` and `"SamplePoints"`. For full descriptions of these properties and valid options, see [movfun], page 846.

Programming Note: This function is a wrapper which calls `movfun`. For full documentation of inputs and options, see [movfun], page 846.

**See also:** [prod], page 612, [movfun], page 846, [movslice], page 848, [movmad], page 850, [movmax], page 851, [movmean], page 852, [movmedian], page 853, [movmin], page 853, [movstd], page 855, [movsum], page 855, [movvar], page 856.

`y = movstd (x, wlen)`
`y = movstd (x, [nb, na])`
`y = movstd (..., opt)`
`y = movstd (..., opt, dim)`
`y = movstd (..., nancond)`
`y = movstd (..., property, value)`

Calculate the moving standard deviation over a sliding window of length *wlen* on data *x*.

The moving window length input *wlen* can either be a numeric scalar or a 2-element numeric array `[nb, na]`. The elements included in the moving window depend on the size and value of *wlen* as well as whether the `"SamplePoints"` option has been specified. For full details of element inclusion, see [movslice], page 848.

The optional argument *opt* determines the type of normalization to use. Valid values are

0: normalize with $N - 1$, provides the square root of the best unbiased estimator of the variance [default]

1: normalize with $N$, this provides the square root of the second moment around the mean

If the optional argument *dim* is given, operate along this dimension. The normalization argument *opt* must be given before the dimension. To use the default value for *opt* you may pass an empty input argument `[]`.

The optional argument *nancond* is a string that controls how `NaN` and `NA` values affect the output of `"movstd"`. The value `"includenan"` (default) causes `NaN` and `NA` values to be included in the moving window, and any window slice containing `NaN` or `NA` values will return `NaN` for that element. The value `"omitnan"` causes `"movstd"` to ignore any `NaN` or `NA` values resulting in fewer elements being used to calculate the standard deviation for that window slice. If `"omitnan"` is specified and a window slice contains all `NaN` or `NA` values, `"movstd"` returns `NaN` for that element. The values `"includemissing"` and `"omitmissing"` may be used synonymously with `"includenan"` and `"omitnan"`, respectively.

The calculation can be controlled by specifying *property*/*value* pairs. Valid properties are `"Endpoints"` and `"SamplePoints"`. For full descriptions of these properties and valid options, see [movfun], page 846.

Programming Note: This function is a wrapper which calls `movfun`. For full documentation of inputs and options, see [movfun], page 846.

**See also:** [std], page 836, [movfun], page 846, [movslice], page 848, [movmad], page 850, [movmax], page 851, [movmean], page 852, [movmedian], page 853, [movmin], page 853, [movprod], page 854, [movsum], page 855, [movvar], page 856.

`y = movsum (x, wlen)`
`y = movsum (x, [nb, na])`
`y = movsum (..., dim)`
`y = movsum (..., nancond)`
`y = movsum (..., property, value)`

Calculate the moving sum over a sliding window of length *wlen* on data *x*.

The moving window length input *wlen* can either be a numeric scalar or a 2-element numeric array `[nb, na]`. The elements included in the moving window depend on the size and value of *wlen* as well as whether the `"SamplePoints"` option has been specified. For full details of element inclusion, see [movslice], page 848.

If the optional argument *dim* is given, operate along this dimension.

The optional argument *nancond* is a string that controls how `NaN` and `NA` values affect the output of `"movsum"`. The value `"includenan"` (default) causes `NaN` and `NA` values to be included in the moving window, and any window slice containing `NaN` or `NA` values will return `NaN` for that element. The value `"omitnan"` causes `"movsum"` to ignore any `NaN` or `NA` values resulting in fewer elements being used to calculate the sum for that window slice. If `"omitnan"` is specified and a window slice contains all `NaN` or `NA` values, `"movsum"` returns 0 for that element. The values `"includemissing"` and `"omitmissing"` may be used synonymously with `"includenan"` and `"omitnan"`, respectively.

The calculation can be controlled by specifying *property*/*value* pairs. Valid properties are `"Endpoints"` and `"SamplePoints"`. For full descriptions of these properties and valid options, see [movfun], page 846.

Programming Note: This function is a wrapper which calls `movfun`. For full documentation of inputs and options, see [movfun], page 846.

**See also:** [sum], page 611, [movfun], page 846, [movslice], page 848, [movmad], page 850, [movmax], page 851, [movmean], page 852, [movmedian], page 853, [movmin], page 853, [movprod], page 854, [movstd], page 855, [movvar], page 856.

`y = movvar (x, wlen)`
`y = movvar (x, [nb, na])`
`y = movvar (..., opt)`
`y = movvar (..., opt, dim)`
`y = movvar (..., nancond)`
`y = movvar (..., property, value)`

Calculate the moving variance over a sliding window of length *wlen* on data *x*.

The moving window length input *wlen* can either be a numeric scalar or a 2-element numeric array `[nb, na]`. The elements included in the moving window depend on the size and value of *wlen* as well as whether the `"SamplePoints"` option has been specified. For full details of element inclusion, see [movslice], page 848.

The optional argument *opt* determines the type of normalization to use. Valid values are:

0: normalize with $N-1$, provides the best unbiased estimator of the variance [default]

1: normalizes with $N$, this provides the second moment around the mean

If the optional argument *dim* is given, operate along this dimension. The normalization argument *opt* must be given before the dimension. To use the default value for *opt* you may pass an empty input argument `[]`.

The optional argument *nancond* is a string that controls how `NaN` and `NA` values affect the output of `"movvar"`. The value `"includenan"` (default) causes `NaN` and `NA` values

to be included in the moving window, and any window slice containing `NaN` or `NA` values will return `NaN` for that element. The value `"omitnan"` causes `"movvar"` to ignore any `NaN` or `NA` values resulting in fewer elements being used to calculate the variance for that window slice. If `"omitnan"` is specified and a window slice contains all `NaN` or `NA` values, `"movvar"` returns `NaN` for that element. The values `"includemissing"` and `"omitmissing"` may be used synonymously with `"includenan"` and `"omitnan"`, respectively.

The calculation can be controlled by specifying *property/value* pairs. Valid properties are `"Endpoints"` and `"SamplePoints"`. For full descriptions of these properties and valid options, see [movfun], page 846.

Programming Note: This function is a wrapper which calls `movfun`. For full documentation of inputs and options, see [movfun], page 846.

**See also:** [var], page 838, [movfun], page 846, [movslice], page 848, [movmad], page 850, [movmax], page 851, [movmean], page 852, [movmedian], page 853, [movmin], page 853, [movprod], page 854, [movstd], page 855, [movsum], page 855.

## 26.3 Basic Statistical Functions

Octave supports various helpful statistical functions. Many are useful as initial steps to prepare a data set for further analysis. Others provide different measures from those of the basic descriptive statistics.

`y = center (x)`
`y = center (x, dim)`
`y = center (x, vecdim)`
`y = center (x, "all")`
`y = center (..., nanflag)`

Center data by subtracting its mean.

If *x* is a vector, then `center (x)` computes the centered data by subtracting the mean of *x* from each element of *x*.

If *x* is a matrix, then `center (x)` returns a row vector with each element containing the centered data for each column of *x*.

If *x* is an array, then `center (x)` centers the data alonng the first non-singleton dimension of *x*.

The data in *x* must be numeric. The size of *y* is equal to the size of *x*.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension of *x*, including any dimension exceeding `ndims (x)`, will return *x*.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored.

Specifying the dimension as `"all"` will cause `center` to compute the center of all elements of *x*, and is equivalent to `center (x(:))`.

The optional variable *nanflag* specifies whether to include or exclude NaN values from the calculation of the mean using any of the previously specified input argument combinations. The default value for *nanflag* is `"includenan"` which keeps NaN values in the calculation. To exclude NaN values set the value of *nanflag* to `"omitnan"`. Any NaN value along the operating dimensions will result in all corresponding element in *y* being NaN.

Programming Note: `center` has obvious application for normalizing statistical data. It is also useful for improving the precision of general numerical calculations. Whenever there is a large value that is common to a batch of data, the mean can be subtracted off, the calculation performed, and then the mean added back to obtain the final answer.

**See also:** [zscore], page 858.

`z = zscore (x)`
`z = zscore (x, opt)`
`z = zscore (x, opt, dim)`
`z = zscore (x, opt, vecdim)`
`z = zscore (x, opt, "all")`
`z = zscore (..., nanflag)`
`[z, mu, sigma] = zscore (...)`

Compute the z-score of *x*.

For a vector *x*, the z-score is calculated by subtracting the mean and dividing by its standard deviation. If the standard deviation is zero, then divide by 1 instead.

If *x* is a vector, then `zscore (x)` returns the z-score of the elements in *x*.

If *x* is a matrix, then `zscore (x)` returns a row vector with each element containing the z-score of the corresponding column in *x*.

If *x* is an array, then `zscore (x)` computes the z-score along the first non-singleton dimension of *x*.

The optional parameter *opt* determines the normalization to use when computing the standard deviation and has the same definition as the corresponding parameter for `std`.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension of *x*, including any dimension exceeding `ndims (x)`, will return `zeros (size (x))`.

Specifying multiple dimensions with input *vecdim*, a vector of non-repeating dimensions, will operate along the array slice defined by *vecdim*. If *vecdim* indexes all dimensions of *x*, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (x)` is ignored.

Specifying the dimension as `"all"` will cause `zscore` to operate on all elements of *x*, and is equivalent to `zscore (x(:))`.

The optional variable *nanflag* specifies whether to include or exclude `NaN` values from the calculation using any of the previously specified input argument combinations. The default value for *nanflag* is `"includenan"` which keeps `NaN` values in the calculation. To exclude `NaN` values, set the value of *nanflag* to `"omitnan"`. The output will still contain `NaN` values at the same locations as in *x*.

The optional outputs *mu* and *sigma* contain the mean and standard deviation.

**See also:** [mean], page 829, [std], page 836, [center], page 857.

```
z = normalize (x)
z = normalize (x, dim)
z = normalize (..., method)
z = normalize (..., method, option)
z = normalize (..., scale, scaleoption, center, centeroption)
[z, c, s] = normalize (...)
```

Return a normalization of the data in $x$ using one of several available scaling and centering methods.

`normalize` by default will return the `zscore` of $x$, defined as the number of standard deviations each element is from the mean of $x$. This is equivalent to centering at the mean of the data and scaling by the standard deviation. $x$ must be a numeric array of double or single floating point numbers.

The returned value $z$ will have the same size as $x$. The optional return variables $c$ and $s$ are the centering and scaling factors used in the normalization such that:

```
z = (x - c) ./ s
```

If $x$ is a vector, `normalize` will operate on the data in $x$.

If $x$ is a matrix, `normalize` will operate independently on each column in $x$.

If $x$ is an N-dimensional array, `normalize` will operate independently on the first non-singleton dimension in $x$.

If the optional second argument *dim* is given, operate along this dimension.

`normalize` ignores NaN values is $x$ similar to the behavior of the omitnan option in `std`, `mean`, and `median`.

The optional inputs *method* and *option* can be used to specify the type of normalization performed on $x$. Note that only the `scale` and `center` options may be specified together using any of the methods defined below. Valid normalization methods are:

`zscore` (Default) Normalizes the elements in $x$ to the scaled distance from a central value. Valid Options:

- `std` (Default) Data is centered at `mean (x)` and scaled by the standard deviation.
- `robust` Data is centered at `median (x)` and scaled by the median absolute deviation.

`norm` $z$ is the general vector norm of $x$, with *option* being the normalization factor $p$ that determines the vector norm type according to:

$$Z = \left( \sum_k |X_k|^P \right)^{1/P}$$

$p$ can be any positive scalar, specific values being:

- `p = 1` $x$ is normalized by `sum (abs (x))`.

`p = 2` (Default) $x$ is normalized by the Euclidian norm, or vector magnitude, of the elements.

`P = Inf` $x$ is normalized by `max (abs (x))`.

`scale` $x$ is scaled by a factor determined by *option*, which can be a numeric scalar or one of the following:

`std` (Default) $x$ is scaled by its standard deviation.

`mad` $x$ is scaled by its median absolute deviation.

`first` $x$ is scaled by its first element.

`iqr` $x$ is scaled by its interquartile range.

`range` $x$ is scaled to fit the range specified by *option* as a two element scalar row vector. The default range is [0, 1].

`center` $x$ is shifted by an amount determined by *option*, which can be a numeric scalar or one of the following:

`mean` (Default) $x$ is shifted by `mean (x)`.

`median` $x$ is shifted by `median (x)`.

`medianiqr` $x$ is shifted by `median (x)` and scaled by the interquartile range.

Known MATLAB incompatibilities:

1. The option `DataVariables` is only available when input $x$ is a table class, which is not yet implemented in core Octave. See the datatypes and tablicious Octave Packages for an available overloaded method.

**See also:** [zscore], page 858, [iqr], page 833, [norm], page 654, [rescale], page 308, [std], page 836, [median], page 830, [mean], page 829, [mad], page 834.

`n = histc (x, edges)`
`n = histc (x, edges, dim)`
`[n, idx] = histc (...)`

Compute histogram counts.

When $x$ is a vector, the function counts the number of elements of $x$ that fall in the histogram bins defined by *edges*. This must be a vector of monotonically increasing values that define the edges of the histogram bins. $n(k)$ contains the number of elements in $x$ for which $edges(k) <= x < edges(k+1)$. The final element of $n$ contains the number of elements of $x$ exactly equal to the last element of *edges*.

When $x$ is an $N$-dimensional array, the computation is carried out along dimension *dim*. If not specified *dim* defaults to the first non-singleton dimension.

When a second output argument is requested an index matrix is also returned. The *idx* matrix has the same size as $x$. Each element of *idx* contains the index of the histogram bin in which the corresponding element of $x$ was counted.

**See also:** [hist], page 341.

`unique` function documented at [unique], page 871, is often useful for statistics.

`c = nchoosek (n, k)`
`c = nchoosek (set, k)`

Compute the binomial coefficient of $n$ or list all possible combinations of a *set* of items.

If $n$ is a scalar then calculate the binomial coefficient of $n$ and $k$ which is defined as

$$\binom{n}{k} = \frac{n(n-1)(n-2)\cdots(n-k+1)}{k!} = \frac{n!}{k!(n-k)!}$$

This is the number of combinations of $n$ items taken in groups of size $k$.

If the first argument is a vector, *set*, then generate all combinations of the elements of *set*, taken $k$ at a time, with one row per combination. The result $c$ has $k$ columns and `nchoosek (length (set), k)` rows.

For example:

How many ways can three items be grouped into pairs?

```
nchoosek (3, 2)
   ⇒ 3
```

What are the possible pairs?

```
nchoosek (1:3, 2)
   ⇒  1   2
      1   3
      2   3
```

Programming Note: When calculating the binomial coefficient `nchoosek` works only for non-negative, integer arguments. Use `bincoeff` for non-integer and negative scalar arguments, or for computing many binomial coefficients at once with vector inputs for $n$ or $k$.

**See also:** [bincoeff], page 633, [perms], page 861.

`P = perms (v)`
`P = perms (v, "unique")`

Generate all permutations of vector $v$ with one row per permutation.

Results are returned in reverse lexicographic order if $v$ is in ascending order. If $v$ is in a different permutation, then the result is permuted that way too. Consequently, an input in descending order yields a result in normal lexicographic order. The result has size `factorial (n) * n`, where $n$ is the length of $v$. Any repeated elements are included in the output.

If the optional argument `"unique"` is given then only unique permutations are returned, using less memory and taking less time than calling `unique (perms (v), "rows")`.

Example 1

```
perms ([1, 2, 3])
⇒
3   2   1
3   1   2
2   3   1
2   1   3
1   3   2
1   2   3
```

Example 2

```
perms ([1, 1, 2, 2], "unique")
⇒
2   2   1   1
2   1   2   1
2   1   1   2
1   2   2   1
1   2   1   2
1   1   2   2
```

Programming Note: If the `"unique"` option is not used, the length of *v* should be no more than 10-12 to limit memory consumption. Even with `"unique"`, there should be no more than 10-12 unique elements in *v*.

**See also:** [permute], page 572, [randperm], page 593, [nchoosek], page 860.

`y = ranks (x)`
`y = ranks (x, dim)`
`y = ranks (x, dim, rtype)`

Return the ranks (in the sense of order statistics) of *x* along the first non-singleton dimension adjusted for ties.

If the optional *dim* argument is given, operate along this dimension.

The optional parameter *rtype* determines how ties are handled. All examples below assume an input of `[ 1, 2, 2, 4 ]`.

0 or `"fractional"` (default) for fractional ranking (1, 2.5, 2.5, 4);

1 or `"competition"` for competition ranking (1, 2, 2, 4);
2 or `"modified"` for modified competition ranking (1, 3, 3, 4);
3 or `"ordinal"` for ordinal ranking (1, 2, 3, 4);
4 or `"dense"` for dense ranking (1, 2, 2, 3).

**See also:** [spearman], page 868, [kendall], page 868.

`cnt = run_count (x, n)`
`cnt = run_count (x, n, dim)`

Count the upward runs along the first non-singleton dimension of *x* of length 1, 2, ..., *n*-1 and greater than or equal to *n*.

If the optional argument *dim* is given then operate along this dimension.

**See also:** [runlength], page 863.

```
count = runlength (x)
[count, value] = runlength (x)
```

Find the lengths of all sequences of common values.

*count* is a vector with the lengths of each repeated value.

The optional output *value* contains the value that was repeated in the sequence.

```
runlength ([2, 2, 0, 4, 4, 4, 0, 1, 1, 1, 1])
⇒   2   1   3   1   4
```

**See also:** [run_count], page 862.

## 26.4 Forecasting Metrics

Octave supports the following functions for calculating forecasting metrics.

```
E = mape (F, A)
E = mape (F, A, dim)
E = mape (F, A, vecdim)
E = mape (F, A, "all")
E = mape (..., nanflag)
E = mape (..., zeroflag)
E = mape (..., 'Weights', W)
```

Compute the mean absolute percentage error between arrays.

The mean absolute percentage error is defined as

$$\text{mape}(F, A) = E = \frac{1}{N} \sum_{i=1}^{N} \left| \frac{A_i - F_i}{A_i} \right| \times 100$$

where $N$ is the number of elements in $F$ and $A$ after broadcasting is applied to the subtraction operation.

The weighted mean absolute percentage error is defined as

$$\text{mape}_\text{w}(F, A) = E_W = \frac{\sum_{i=1}^{N} W_i \left| \frac{A_i - F_i}{A_i} \right| \times 100}{\sum_{i=1}^{N} W_i}$$

where $N$ is the number of elements in $F$ and $A$ after broadcasting is applied to the subtraction operation.

$F$ and $A$ must either be the same size or have compatible sizes.

If $F$ and $A$ are vectors of the same size, then `mape (F, A)` returns a scalar with the MAPE between the elements of $F$ and $A$.

If `A - F` is a matrix, then `mape (F, A)` returns a row vector with each element containing the MAPE between the corresponding columns of `A - F`.

If `A - F` is an array, then `mape (F, A)` computes the MAPE along the first non-singleton dimension of the difference between the input arrays $F$ and $A$. The size of $E$ along the operating dimension is 1, while all other dimensions are the same as in `A - F`.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension, including any dimension exceeding `ndims (A - F)`, will return `abs ((A - F) ./ A)`.

Specifying the dimensions as *vecdim*, a vector of non-repeating dimensions, will return the mape over the array slice defined by *vecdim*. If *vecdim* indexes all dimensions in `A - F`, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (A - F)` is ignored.

Specifying the dimension as `"all"` will cause `mape` to operate on all elements of `A - F`, and is equivalent to `mape ((A - F)(:))`.

The optional variable *nanflag* specifies whether to include or exclude NaN values from the calculation using any of the previously specified input argument combinations. The default value for *nanflag* is `"includenan"` which keeps NaN values in the calculation. To exclude NaN values set the value of *nanflag* to `"omitnan"`. The output will still contain NaN values if `A - F` consists of all NaN values in the operating dimension.

The optional variable *zeroflag* specifies whether to include or omit zero values from the calculation using any of the previously specified input argument combinations. The default value for *zeroflag* is `"includezero"`, in which case the calculated MAPE is Inf, if $A$ contains one or more zeros. To ignore any zero values in $A$, set the value of *zeroflag* to `"omitzero"`. If $A$ consists of all zero values in the operating dimension, then MAPE is NaN.

The optional paired argument `..., "Weights", W` specifies a weighting scheme $W$, which is applied on the difference of the input arrays $F$ and $A$, so that `mape` computes the weighted MAPE. When operating along a single dimension, $W$ must be a vector of the same length as the operating dimension or it must have the same size as $x$. When operating over an array slice defined by *vecdim*, $W$ must have the same size as the operating array slice, i.e. `size (A - F)(vecdim)`, or the same size as `A - F`.

**See also:** [rmse], page 864, [mean], page 829, [abs], page 605.

`E = rmse (F, A)`
`E = rmse (F, A, dim)`
`E = rmse (F, A, vecdim)`
`E = rmse (F, A, "all")`
`E = rmse (..., nanflag)`
`E = rmse (..., 'Weights', W)`

Compute the root mean squared error between arrays.

The root mean squared error is defined as

$$\text{rmse}(F, A) = E = \sqrt{\frac{1}{N} \sum_{i=1}^{N} |A_i - F_i|^2}$$

where $N$ is the number of elements in $F$ and $A$ after broadcasting is applied to the subtraction operation.

The weighted root mean squared error is defined as

$$\text{rmse}(F, A) = E_W = \sqrt{\frac{\sum_{i=1}^{N} W_i |A_i - F_i|^2}{\sum_{i=1}^{N} W_i}}$$

where $N$ is the number of elements in $F$ and $A$ after broadcasting is applied to the subtraction operation.

$F$ and $A$ must either be the same size or have compatible sizes.

If $F$ and $A$ are vectors of the same size, then `rmse (F, A)` returns a scalar with the RMSE between the elements of $F$ and $A$.

If `A - F` is a matrix, then `rmse (F, A)` returns a row vector with each element containing the RMSE between the corresponding columns of `A - F`.

If `A - F` is an array, then `rmse (F, A)` computes the RMSE along the first non-singleton dimension of the difference between the input arrays $F$ and $A$. The size of $E$ along the operating dimension is 1, while all other dimensions are the same as in `A - F`.

The optional input *dim* specifies the dimension to operate on and must be a positive integer. Specifying any singleton dimension, including any dimension exceeding `ndims (A - F)`, will return `abs ((A - F) ./ A)`.

Specifying the dimensions as *vecdim*, a vector of non-repeating dimensions, will return the rmse over the array slice defined by *vecdim*. If *vecdim* indexes all dimensions in `A - F`, then it is equivalent to the option `"all"`. Any dimension in *vecdim* greater than `ndims (A - F)` is ignored.

Specifying the dimension as `"all"` will cause `rmse` to operate on all elements of `A - F`, and is equivalent to `rmse ((A - F)(:))`.

The optional variable *nanflag* specifies whether to include or exclude NaN values from the calculation using any of the previously specified input argument combinations. The default value for *nanflag* is `"includenan"` which keeps NaN values in the calculation. To exclude NaN values set the value of *nanflag* to `"omitnan"`. The output will still contain NaN values if `A - F` consists of all NaN values in the operating dimension.

The optional paired argument `..., "Weights", W` specifies a weighting scheme $W$, which is applied on the difference of the input arrays $F$ and $A$, so that `rmse` computes the weighted RMSE. When operating along a single dimension, $W$ must be a vector of the same length as the operating dimension or it must have the same size as $x$. When operating over an array slice defined by *vecdim*, $W$ must have the same size as the operating array slice, i.e. `size (A - F)(vecdim)`, or the same size as `A - F`.

**See also:** [mape], page 863, [meansq], page 835, [rms], page 836.

## 26.5 Correlation and Regression Analysis

Octave supports several functions for investigating the relationship between sets of quantitative variables.

`c = cov (x)`
`c = cov (x, y)`
`c = cov (..., opt)`
`c = cov (..., nanflag)`
Compute the covariance matrix.

The covariance between two variable vectors $A$ and $B$ is calculated as:

$$\sigma_{ij} = \frac{1}{N-1} \sum_{i=1}^{N} (a_i - \bar{a})(b_i - \bar{b})$$

where $\bar{a}$ and $\bar{b}$ are the mean values of $a$ and $b$ and $N$ is the length of the vectors $a$ and $b$.

If called with one argument, compute `cov (x, x)`. If $x$ is a vector, this is the scalar variance of $x$. If $x$ is a matrix, each row of $x$ is treated as an observation, and each column as a variable, and the ($i$, $j$)-th entry of `cov (x)` is the covariance between the $i$-th and $j$-th columns in $x$. If $x$ has dimensions n x m, the output $c$ will be a m x m square covariance matrix.

If called with two arguments, compute `cov (x, y)`, the covariance between two random variables $x$ and $y$. $x$ and $y$ must have the same number of elements, and will be treated as vectors with the covariance computed as `cov (x(:), y(:))`. The output will be a 2 x 2 covariance matrix.

The optional argument *opt* determines the type of normalization to use. Valid values are

0 [default]:
: Normalize with $N - 1$. This provides the best unbiased estimator of the covariance.

1:
: Normalize with $N$. This provides the second moment around the mean. *opt* is set to 1 for N = 1.

The optional argument *nanflag* must appear last in the argument list and controls how NaN values are handled by `cov`. The three valid values are:

includenan [default]:
: Leave NaN values in $x$ and $y$. Output will follow the normal rules for handling NaN values in arithmetic operations.

omitrows:
: Rows containing NaN values are trimmed from both $x$ and $y$ prior to calculating the covariance. A NaN in one variable will remove that row from both $x$ and $y$.

partialrows:
: Rows containing NaN values are ignored from both $x$ and $y$ independently for each $i$-th and $j$-th covariance calculation. This may result in a different number of observations, $N$, being used to calculated each element of the covariance matrix.

Compatibility Note: Before Octave v9.1.0, `cov` treated rows $x$ and $y$ as multivariate random variables. Newer versions attempt to maintain full compatibility with MATLAB by treating $x$ and $y$ as two univariate distributions regardless of shape, resulting in a 2x2 output matrix. Code relying on Octave's previous definition will need to be modified when running this newer version of `cov`. The previous behavior can be obtained by using the NaN package's `covm` function as `covm (x, y, "D")`.

**See also:** [corr], page 867.

`r = corr (x)`
`r = corr (x, y)`

Compute matrix of correlation coefficients.

If each row of $x$ and $y$ is an observation and each column is a variable, then the $(i, j)$-th entry of `corr (x, y)` is the correlation between the $i$-th variable in $x$ and the $j$-th variable in $y$. $x$ and $y$ must have the same number of rows (observations). The correlation coefficient is calculated for two variable vectors $A$ and $B$ (columns of $x$ and $y$) as:

$$\operatorname{corr}(A, B) = \frac{\operatorname{cov}(A, B)}{\operatorname{std}(A)\operatorname{std}(B)}$$

If called with one argument, compute `corr (x, x)`, the correlation between the each pair of columns of $x$.

**See also:** [cov], page 865, [corrcoef], page 867.

`r = corrcoef (x)`
`r = corrcoef (x, y)`
`r = corrcoef (..., param, value, ...)`
`[r, p] = corrcoef (...)`
`[r, p, lci, hci] = corrcoef (...)`

Compute a matrix of correlation coefficients.

$x$ is an array where each column contains a variable and each row is an observation.

If a second input $y$ (of the same size as $x$) is given then calculate the correlation coefficients between $x$ and $y$.

*param*, *value* are optional pairs of parameters and values which modify the calculation. Valid options are:

`"alpha"` Confidence level used for the bounds of the confidence interval, *lci* and *hci*. Default is 0.05, i.e., 95% confidence interval.

`"rows"` Determine processing of NaN values. Acceptable values are `"all"`, `"complete"`, and `"pairwise"`. Default is `"all"`. With `"complete"`, only the rows without NaN values are considered. With `"pairwise"`, the selection of NaN-free rows is made for each pair of variables.

Output $r$ is a matrix of Pearson's product moment correlation coefficients for each pair of variables.

Output $p$ is a matrix of pair-wise p-values testing for the null hypothesis of a correlation coefficient of zero.

Outputs *lci* and *hci* are matrices containing, respectively, the lower and higher bounds of the 95% confidence interval of each correlation coefficient.

**See also:** [corr], page 867, [cov], page 865, [std], page 836.

`r = corrcov (c)`
`[r, s] = corrcov (c)`

Convert matrix of covariance coefficients to a matrix of correlation coefficients.

Given a numeric, square, symmetric, and positive semi-definite matrix of covariance coefficients, $c$, calculate the corresponding linear correlation coefficient matrix, $r$, that would be produced from the same variables that produced $c$.

The correlation and covariance coefficients for two vectors, $A$ and $B$, are related by:

$$\operatorname{corr}(A, B) = \frac{\operatorname{cov}(A, B)}{\sqrt{\operatorname{cov}(A, A) \cdot \operatorname{cov}(B, B)}}$$

The output $r$ will be the same shape and type as the input, $c$, where each element $r(i,j)$ contains the correlation coefficient calculated using $c(i,j)$ as described above.

If the optional output $s$ is requested, it will contain a column vector of the standard deviations of the variables that produced the covariance matrix $c$. Noting that the variances of the variables are contained in the main diagonal of $c$, this is equivalent output to `sqrt (diag ( c))`.

Implementation Note: A proper covariance matrix as produced by `cov` is square, symmetric, and positive semi-definite. `corrcov` only uses an approximate check for the latter that may pass improper inputs with small negative eigenvalues.

**See also:** [cov], page 865, [corr], page 867, [corrcoef], page 867.

*rho* = **spearman** (*x*)
*rho* = **spearman** (*x*, *y*)

Compute Spearman's rank correlation coefficient $\rho$.

For two data vectors $x$ and $y$, Spearman's $\rho$ is the correlation coefficient of the ranks of $x$ and $y$.

If $x$ and $y$ are drawn from independent distributions, $\rho$ has zero mean and variance $1/(N-1)$, where $N$ is the length of the $x$ and $y$ vectors, and is asymptotically normally distributed.

`spearman (x)` is equivalent to `spearman (x, x)`.

**See also:** [ranks], page 862, [kendall], page 868.

*tau* = **kendall** (*x*)
*tau* = **kendall** (*x*, *y*)

Compute Kendall's $\tau$.

For two data vectors $x$, $y$ of common length $N$, Kendall's $\tau$ is the correlation of the signs of all rank differences of $x$ and $y$; i.e., if both $x$ and $y$ have distinct entries, then

$$\tau = \frac{1}{N(N-1)} \sum_{i,j} \operatorname{sign}(q_i - q_j) \operatorname{sign}(r_i - r_j)$$

in which the $q_i$ and $r_i$ are the ranks of $x$ and $y$, respectively.

If $x$ and $y$ are drawn from independent distributions, Kendall's $\tau$ is asymptotically normal with mean 0 and variance $\frac{2(2N+5)}{9N(N-1)}$.

`kendall (x)` is equivalent to `kendall (x, x)`.

**See also:** [ranks], page 862, [spearman], page 868.

## 26.6 Distributions

Octave has functions for computing the Probability Density Function (PDF), the Cumulative Distribution function (CDF), and the quantile (the inverse of the CDF) for arbitrary user-defined distributions (discrete) and for experimental data (empirical).

The following table summarizes the supported distributions (in alphabetical order).

| Distribution | PDF | CDF | Quantile |
|---|---|---|---|
| Univariate Discrete | discrete_pdf | discrete_cdf | discrete_inv |
| Empirical | empirical_pdf | empirical_cdf | empirical_inv |

*pdf* = `discrete_pdf` (*x*, *v*, *p*)
: For each element of *x*, compute the probability density function (PDF) at *x* of a univariate discrete distribution which assumes the values in *v* with probabilities *p*.

*cdf* = `discrete_cdf` (*x*, *v*, *p*)
: For each element of *x*, compute the cumulative distribution function (CDF) at *x* of a univariate discrete distribution which assumes the values in *v* with probabilities *p*.

*q* = `discrete_inv` (*x*, *v*, *p*)
: For each element of *x*, compute the quantile (the inverse of the CDF) at *x* of the univariate distribution which assumes the values in *v* with probabilities *p*.

*pdf* = `empirical_pdf` (*x*, *data*)
: For each element of *x*, compute the probability density function (PDF) at *x* of the empirical distribution obtained from the univariate sample *data*.

*cdf* = `empirical_cdf` (*x*, *data*)
: For each element of *x*, compute the cumulative distribution function (CDF) at *x* of the empirical distribution obtained from the univariate sample *data*.

*q* = `empirical_inv` (*x*, *data*)
: For each element of *x*, compute the quantile (the inverse of the CDF) at *x* of the empirical distribution obtained from the univariate sample *data*.

## 26.7 Random Number Generation

Octave can generate random numbers from a large number of distributions. The random number generators are based on the random number generators described in Section 16.3 [Special Utility Matrices], page 580.

The following table summarizes the available random number generators (in alphabetical order).

| Distribution | Function |
|---|---|
| Univariate Discrete Distribution | discrete_rnd |
| Empirical Distribution | empirical_rnd |
| Exponential Distribution | rande |
| Gamma Distribution | randg |
| Poisson Distribution | randp |
| Standard Normal Distribution | randn |
| Uniform Distribution | rand |
| Uniform Distribution (integers) | randi |

*rnd* = discrete_rnd (*v*, *p*)
*rnd* = discrete_rnd (*v*, *p*, *r*)
*rnd* = discrete_rnd (*v*, *p*, *r*, *c*, ...)
*rnd* = discrete_rnd (*v*, *p*, [*sz*])

Return a matrix of random samples from the univariate distribution which assumes the values in *v* with probabilities *p*.

When called with a single size argument, return a square matrix with the dimension specified. When called with more than one scalar argument the first two arguments are taken as the number of rows and columns and any further arguments specify additional matrix dimensions. The size may also be specified with a vector of dimensions *sz*.

If no size arguments are given then the result matrix is the common size of *v* and *p*.

*rnd* = empirical_rnd (*data*)
*rnd* = empirical_rnd (*data*, *r*)
*rnd* = empirical_rnd (*data*, *r*, *c*, ...)
*rnd* = empirical_rnd (*data*, [*sz*])

Return a matrix of random samples from the empirical distribution obtained from the univariate sample *data*.

When called with a single size argument, return a square matrix with the dimension specified. When called with more than one scalar argument the first two arguments are taken as the number of rows and columns and any further arguments specify additional matrix dimensions. The size may also be specified with a vector of dimensions *sz*.

If no size arguments are given then the result matrix is a random ordering of the sample *data*.

# 27 Sets

Octave has a number of functions for managing sets of data. A set is defined as a collection of unique elements and is typically represented by a vector of numbers sorted in ascending order. Any vector or matrix can be converted to a set by removing duplicates through the use of the `unique` function. However, it isn't necessary to explicitly create a set as all of the functions which operate on sets will convert their input to a set before proceeding.

`y = unique (x)`
`y = unique (x, "rows")`
`y = unique (..., "sorted")`
`y = unique (..., "stable")`
`[y, i, j] = unique (...)`
`[y, i, j] = unique (..., "first")`
`[y, i, j] = unique (..., "last")`
`[y, i, j] = unique (..., "legacy")`

Return the unique elements of *x*.

If the input *x* is a column vector then return a column vector; Otherwise, return a row vector. *x* may also be a cell array of strings.

If the optional argument `"rows"` is given then return the unique rows of *x*. The input must be a 2-D numeric matrix to use this option.

The optional argument `"sorted"`/`"stable"` controls the order in which unique values appear in the output. The default is `"sorted"` and values in the output are placed in ascending order. The alternative `"stable"` preserves the order found in the input *x*.

If requested, return column index vectors *i* and *j* such that `y = x(i)` and `x = y(j)`.

Additionally, if *i* is a requested output then one of the flags `"first"` or `"last"` may be given. If `"last"` is specified, return the highest possible indices in *i*, otherwise, if `"first"` is specified, return the lowest. The default is `"first"`.

Example 1 : sort order

```
unique ([3, 1, 1, 2])
⇒ [1, 2, 3]
unique ([3, 1, 1, 2], "stable")
⇒ [3, 1, 2]
```

Example 2 : index selection

```
[~, i] = unique ([3, 1, 1, 2], "first")
⇒ i = [2; 4; 1]
[~, i] = unique ([3, 1, 1, 2], "last")
⇒ i = [3; 4; 1]
```

Programming Notes: The input flag `"legacy"` changes the algorithm to be compatible with MATLAB releases prior to R2012b. Specifically, The index ordering flag is changed to `"last"`, and the shape of the outputs *i*, *j* will follow the shape of the input *x* rather than always being column vectors.

**See also:** [uniquetol], page 872, [union], page 873, [intersect], page 873, [setdiff], page 874, [setxor], page 874, [ismember], page 875.

`c = uniquetol (A)`
`c = uniquetol (A, tol)`
`c = uniquetol (..., property, value)`
`[c, ia, ic] = uniquetol (...)`

Return the unique elements of $A$ within tolerance *tol*.

Two values, $x$ and $y$, are within relative tolerance if `abs (x - y) <= tol * max (abs (A(:)))`.

The input $A$ must be a real (non-complex) floating point type (double or single).

If *tol* is unspecified, the default tolerance is 1e-12 for double precision input or 1e-6 for single precision input.

The function may also be called with the following optional property/value pairs. Property/value pairs must be passed after other input arguments:

`"ByRows"` (default: `false`)
: When true, return the unique rows of $A$. $A$ must be a 2-D array to use this option. For rows, the criteria for uniqueness is changed to `all (abs (x - y) <= tol*max (abs (A),[],1))` which compares each column component of a row against a column-specific tolerance.

`"DataScale"`
: The tolerance test is changed to `abs (x - y) <= tol*DS` where *DS* is a scalar unless the property `"ByRows"` is true. In that case, *DS* can either be a scalar or a vector with a length equal to the number of columns in $A$. Using a value of `1.0` for *DS* will change the tolerance from a relative one to an absolute tolerance. Using a value of `Inf` will disable testing.

`"OutputAllIndices"` (default: `false`)
: When true, *ia* is a cell array (not a vector) that contains the indices for *all* elements in $A$ that are within tolerance of a value in $C$. That is, each cell in *ia* corresponds to a single unique value in $C$, and the values in each cell correspond to locations in $A$.

The output $c$ is a row vector if the input $A$ is a row vector. For all other cases, a column vector is returned.

The optional output *ia* is a column index vector such that `c = A(ia)`. If the `"ByRows"` property is true, the condition is `c = A(ia, :)`. If the `"OutputAllIndices"` property is true, then the values `A(ia{i})` are all within tolerance of the unique value `c(i)`.

The optional output *ic* is a column index vector such that `A = c(ic)` when $A$ is a vector. When $A$ is a matrix, `A(:) = c(ic)`. If the `"ByRows"` property is true then `A = c(ic,:)`.

Example: small round-off errors require `uniquetol`, not `unique`

```
x = [1:5];
## Inverse_Function (Function (x)) should return exactly x
y = exp (log (x));
D = unique ([x, y])
⇒ [1   2   3   3   4   5   5]
C = uniquetol ([x, y])
⇒ [1   2   3   4   5]
```

**See also:** [unique], page 871, [union], page 873, [intersect], page 873, [setdiff], page 874, [setxor], page 874, [ismember], page 875.

## 27.1 Set Operations

Octave supports several basic set operations. Octave can compute the union, intersection, and difference of two sets. Octave also supports the *Exclusive Or* set operation.

The functions for set operations all work in the same way by accepting two input sets and returning a third set. As an example, assume that `a` and `b` contains two sets, then

```
union (a, b)
```

computes the union of the two sets.

Finally, determining whether elements belong to a set can be done with the `ismember` function. Because sets are ordered this operation is very efficient and is of order O(log2(n)) which is preferable to the `find` function which is of order O(n).

`c = intersect (a, b)`
`c = intersect (a, b, "rows")`
`c = intersect (..., "sorted")`
`c = intersect (..., "stable")`
`c = intersect (..., "legacy")`
`[c, ia, ib] = intersect (...)`

Return the unique elements common to both *a* and *b*.

If *a* and *b* are both row vectors then return a row vector; Otherwise, return a column vector. The inputs may also be cell arrays of strings.

If the optional input `"rows"` is given then return the common rows of *a* and *b*. The inputs must be 2-D numeric matrices to use this option.

The optional argument `"sorted"`/`"stable"` controls the order in which unique values appear in the output. The default is `"sorted"` and values in the output are placed in ascending order. The alternative `"stable"` preserves the order found in the input.

If requested, return column index vectors *ia* and *ib* such that `c = a(ia)` and `c = b(ib)`.

Programming Note: The input flag `"legacy"` changes the algorithm to be compatible with MATLAB releases prior to R2012b.

**See also:** [unique], page 871, [union], page 873, [setdiff], page 874, [setxor], page 874, [ismember], page 875.

`c = union (a, b)`
`c = union (a, b, "rows")`
`c = union (..., "sorted")`
`c = union (..., "stable")`
`c = union (..., "legacy")`
`[c, ia, ib] = union (...)`

Return the unique elements that are in either *a* or *b*.

If *a* and *b* are both row vectors then return a row vector; Otherwise, return a column vector. The inputs may also be cell arrays of strings.

If the optional input `"rows"` is given then return rows that are in either *a* or *b*. The inputs must be 2-D numeric matrices to use this option.

The optional argument `"sorted"`/`"stable"` controls the order in which unique values appear in the output. The default is `"sorted"` and values in the output are placed in ascending order. The alternative `"stable"` preserves the order found in the input.

The optional outputs *ia* and *ib* are column index vectors such that `a(ia)` and `b(ib)` are disjoint sets whose union is *c*.

Programming Note: The input flag `"legacy"` changes the algorithm to be compatible with MATLAB releases prior to R2012b.

**See also:** [unique], page 871, [intersect], page 873, [setdiff], page 874, [setxor], page 874, [ismember], page 875.

`c = setdiff (a, b)`
`c = setdiff (a, b, "rows")`
`c = setdiff (..., "sorted")`
`c = setdiff (..., "stable")`
`c = setdiff (..., "legacy")`
`[c, ia] = setdiff (...)`

Return the unique elements in *a* that are not in *b*.

If *a* is a row vector return a row vector; Otherwise, return a column vector. The inputs may also be cell arrays of strings.

If the optional input `"rows"` is given then return the rows in *a* that are not in *b*. The inputs must be 2-D numeric matrices to use this option.

The optional argument `"sorted"`/`"stable"` controls the order in which unique values appear in the output. The default is `"sorted"` and values in the output are placed in ascending order. The alternative `"stable"` preserves the order found in the input.

If requested, return the index vector *ia* such that `c = a(ia)`.

Programming Note: The input flag `"legacy"` changes the algorithm to be compatible with MATLAB releases prior to R2012b.

**See also:** [unique], page 871, [union], page 873, [intersect], page 873, [setxor], page 874, [ismember], page 875.

`c = setxor (a, b)`
`c = setxor (a, b, "rows")`
`c = setxor (..., "sorted")`
`c = setxor (..., "stable")`
`c = setxor (..., "legacy")`
`[c, ia, ib] = setxor (...)`

Return the unique elements exclusive to sets *a* or *b*.

If *a* and *b* are both row vectors then return a row vector; Otherwise, return a column vector. The inputs may also be cell arrays of strings.

If the optional input `"rows"` is given then return the rows exclusive to sets *a* and *b*. The inputs must be 2-D numeric matrices to use this option.

The optional argument `"sorted"`/`"stable"` controls the order in which unique values appear in the output. The default is `"sorted"` and values in the output are placed in ascending order. The alternative `"stable"` preserves the order found in the input.

The optional outputs *ia* and *ib* are column index vectors such that `a(ia)` and `b(ib)` are disjoint sets whose union is *c*.

Programming Note: The input flag `"legacy"` changes the algorithm to be compatible with MATLAB releases prior to R2012b.

**See also:** [unique], page 871, [union], page 873, [intersect], page 873, [setdiff], page 874, [ismember], page 875.

*tf* = ismember (*a*, *s*)
*tf* = ismember (*a*, *s*, "*rows*")
[*tf*, *s_idx*] = ismember (...)

Return a logical matrix *tf* with the same shape as *a* which is true (1) if the element in *a* is found in *s* and false (0) if it is not.

If a second output argument is requested then the index into *s* of each matching element is also returned.

```
a = [3, 10, 1];
s = [0:9];
[tf, s_idx] = ismember (a, s)
     ⇒ tf = [1, 0, 1]
     ⇒ s_idx = [4, 0, 2]
```

The inputs *a* and *s* may also be cell arrays.

```
a = {"abc"};
s = {"abc", "def"};
[tf, s_idx] = ismember (a, s)
     ⇒ tf = 1
     ⇒ s_idx = 1
```

If the optional third argument `"rows"` is given then compare rows in *a* with rows in *s*. The inputs must be 2-D matrices with the same number of columns to use this option.

```
a = [1:3; 5:7; 4:6];
s = [0:2; 1:3; 2:4; 3:5; 4:6];
[tf, s_idx] = ismember (a, s, "rows")
     ⇒ tf = logical ([1; 0; 1])
     ⇒ s_idx = [2; 0; 5];
```

**See also:** [lookup], page 569, [unique], page 871, [union], page 873, [intersect], page 873, [setdiff], page 874, [setxor], page 874, [ismembertol], page 875.

*tf* = ismembertol (*a*, *s*)
*tf* = ismembertol (*a*, *s*, *tol*)
*tf* = ismembertol (*a*, *s*, *name*, *value*)
[*tf*, *s_idx*] = ismembertol (...)

Check if values are members of a set within a tolerance.

This functions returns a logical matrix *tf* with the same shape as *a* which is true (1) where the element in *a* is close to *s* within a tolerance *tol* and false (0) if it is not. If *tol* is not specified, a default tolerance of `1e-6` is used.

If a second output argument is requested then the index into *s* of each matching element is also returned.

The inputs *a* and *s* must be numeric values.

```
a = [3, 10, 1];
s = [0:9];
[tf, s_idx] = ismembertol (a, s)
     ⇒ tf = [1, 0, 1]
     ⇒ s_idx = [4, 0, 2]
```

Optional property/value pairs may be given to change the function's behavior. The property may be one of following strings:

`"ByRows"` If set to `false` (default), all elements in *a* and *s* are treated separately. If set to `true`, *tf* will be `true` for each row in *a* that matches a row in *s* within the given tolerance. Two rows, *u* and *v*, are within tolerance if they fulfill the condition `all (abs (u-v) <= tol*max (abs ([a;s])))`.

`"OutputAllIndices"`
If set to `false` (default), *s_idx* contains indices for one of the matches. If set to `true`, *s_idx* is a cell array containing the indices for all elements in *s* that are within tolerance of the corresponding value in *a*.

`"DataScale"`
The provided value *DS* is used to change the scale factor in the tolerance test to `abs (u-v) <= tol*`*DS*. By default, the maximum absolute value in *a* and *s* is used as the scale factor.

Example:

```
s = [1:6].' * pi;
a = 10.^log10 (x);
[tf, s_idx] = ismembertol (a, s);
```

**See also:** [ismember], page 875, [lookup], page 569, [unique], page 871, [union], page 873, [intersect], page 873, [setdiff], page 874, [setxor], page 874.

`p = powerset (a)`
`p = powerset (a, "`*rows*`")`
Compute the powerset (all subsets) of the set *a*.

The set *a* must be a numerical matrix or a cell array of strings. The output will always be a cell array of either vectors or strings.

With the optional argument `"rows"`, each row of the set *a* is considered one element of the set. The input must be a 2-D numeric matrix to use this argument.

**See also:** [unique], page 871, [union], page 873, [intersect], page 873, [setdiff], page 874, [setxor], page 874, [ismember], page 875.

# 28 Polynomial Manipulations

In Octave, a polynomial is represented by its coefficients (arranged in descending order). For example, a vector $c$ of length $N+1$ corresponds to the following polynomial of order $N$

$$p(x) = c_1 x^N + \ldots + c_N x + c_{N+1}.$$

## 28.1 Evaluating Polynomials

The value of a polynomial represented by the vector $c$ can be evaluated at the point $x$ very easily, as the following example shows:

```
N = length (c) - 1;
val = dot (x.^(N:-1:0), c);
```

While the above example shows how easy it is to compute the value of a polynomial, it isn't the most stable algorithm. With larger polynomials you should use more elegant algorithms, such as Horner's Method, which is exactly what the Octave function `polyval` does.

In the case where $x$ is a square matrix, the polynomial given by $c$ is still well-defined. As when $x$ is a scalar the obvious implementation is easily expressed in Octave, but also in this case more elegant algorithms perform better. The `polyvalm` function provides such an algorithm.

`y = polyval (p, x)`
`y = polyval (p, x, [], mu)`
`[y, dy] = polyval (p, x, s)`
`[y, dy] = polyval (p, x, s, mu)`

Evaluate the polynomial $p$ at the specified values of $x$.

If $x$ is a vector or matrix, the polynomial is evaluated for each of the elements of $x$.

When *mu* is present, evaluate the polynomial for ($x$ - *mu*(1)) / *mu*(2).

In addition to evaluating the polynomial, the second output represents the prediction interval, $y$ +/- $dy$, which contains at least 50% of the future predictions. To calculate the prediction interval, the structured variable $s$, originating from `polyfit`, must be supplied.

**See also:** [polyvalm], page 877, [polyaffine], page 883, [polyfit], page 883, [roots], page 878, [poly], page 893.

`y = polyvalm (c, x)`

Evaluate a polynomial in the matrix sense.

`polyvalm (c, x)` will evaluate the polynomial in the matrix sense, i.e., matrix multiplication is used instead of element by element multiplication as used in `polyval`.

The argument $x$ must be a square matrix.

**See also:** [polyval], page 877, [roots], page 878, [poly], page 893.

## 28.2 Finding Roots

Octave can find the roots of a given polynomial. This is done by computing the companion matrix of the polynomial (see the `compan` function for a definition), and then finding its eigenvalues.

`r = roots (c)`

Compute the roots of the polynomial $c$.

For a vector $c$ with $N$ components, return the roots of the polynomial

$$c_1 x^{N-1} + \cdots + c_{N-1} x + c_N.$$

As an example, the following code finds the roots of the quadratic polynomial

$$p(x) = x^2 - 5.$$

```
c = [1, 0, -5];
roots (c)
⇒  2.2361
⇒ -2.2361
```

Note that the true result is $\pm\sqrt{5}$ which is roughly $\pm 2.2361$.

**See also:** [poly], page 893, [compan], page 878, [fzero], page 710.

`z = polyeig (C0, C1, ..., Cl)`
`[v, z] = polyeig (C0, C1, ..., Cl)`

Solve the polynomial eigenvalue problem of degree $l$.

Given an $n$x$n$ matrix polynomial

`C(s) = C0 + C1 s + ... + Cl s^l`

`polyeig` solves the eigenvalue problem

`(C0 + C1 z + ... + Cl z^l) v = 0`.

Note that the eigenvalues $z$ are the zeros of the matrix polynomial. $z$ is a row vector with `n*l` elements. $v$ is a matrix ($n$ x $n$*$l$) with columns that correspond to the eigenvectors.

**See also:** [eig], page 649, [eigs], page 754, [compan], page 878.

`A = compan (c)`

Compute the companion matrix corresponding to polynomial coefficient vector $c$.

The companion matrix is

$$A = \begin{bmatrix} -c_2/c_1 & -c_3/c_1 & \cdots & -c_N/c_1 & -c_{N+1}/c_1 \\ 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & 0 \end{bmatrix}.$$

The eigenvalues of the companion matrix are equal to the roots of the polynomial.

**See also:** [roots], page 878, [poly], page 893, [eig], page 649.

`[multp, idxp] = mpoles (p)`
`[multp, idxp] = mpoles (p, tol)`
`[multp, idxp] = mpoles (p, tol, reorder)`

Identify unique poles in *p* and their associated multiplicity.

By default, the output is ordered from the pole with the largest magnitude to the smallest magnitude.

Two poles are considered to be multiples if the difference between them is less than the relative tolerance *tol*.

```
abs (p1 - p0) / abs (p0) < tol
```

If the pole is 0 then no scaling is done and *tol* is interpreted as an absolute tolerance. The default value for *tol* is 0.001.

If the optional parameter *reorder* is false/zero, poles are not sorted.

The output *multp* is a vector specifying the multiplicity of the poles. `multp(n)` refers to the multiplicity of the Nth pole `p(idxp(n))`.

For example:

```
p = [2 3 1 1 2];
[m, n] = mpoles (p)
   ⇒ m = [1; 1; 2; 1; 2]
   ⇒ n = [2; 5; 1; 4; 3]
   ⇒ p(n) = [3, 2, 2, 1, 1]
```

**See also:** [residue], page 881, [poly], page 893, [roots], page 878, [conv], page 879, [deconv], page 880.

## 28.3 Products of Polynomials

`y = conv (a, b)`
`y = conv (a, b, shape)`

Convolve two vectors *a* and *b*.

When *a* and *b* are the coefficient vectors of two polynomials, the convolution represents the coefficient vector of the product polynomial.

The size of the result is determined by the optional *shape* argument which takes the following values

*shape* = `"full"`
: Return the full convolution. (default) The result is a vector with length equal to `length (a) + length (b) - 1`.

*shape* = `"same"`
: Return the central part of the convolution with the same size as *a*.

*shape* = `"valid"`
: Return only the parts which do not include zero-padded edges. The size of the result is `max (size (a) - size (b) + 1, 0)`.

**See also:** [deconv], page 880, [conv2], page 880, [convn], page 880, [fftconv], page 928.

`C = convn (A, B)`
`C = convn (A, B, shape)`

Return the N-D convolution of $A$ and $B$.

The size of the result is determined by the optional *shape* argument which takes the following values

`"full"` Return the full convolution. (default)

`"same"` Return central part of the convolution with the same size as $A$. The central part of the convolution begins at the indices `floor ([size(B)/2] + 1)`.

`"valid"` Return only the parts which do not include zero-padded edges. The size of the result is `max (size (A) - size (B) + 1, 0)`.

**See also:** [conv2], page 880, [conv], page 879.

`b = deconv (y, a)`
`[b, r] = deconv (y, a)`

Deconvolve two vectors (polynomial division).

`[b, r] = deconv (y, a)` solves for $b$ and $r$ such that `y = conv (a, b) + r`.

If $y$ and $a$ are polynomial coefficient vectors, $b$ will contain the coefficients of the polynomial quotient and $r$ will be a remainder polynomial of lowest order.

**See also:** [conv], page 879, [residue], page 881.

`C = conv2 (A, B)`
`C = conv2 (v1, v2, m)`
`C = conv2 (..., shape)`

Return the 2-D convolution of $A$ and $B$.

The size of the result is determined by the optional *shape* argument which takes the following values

`"full"` Return the full convolution. (default)

`"same"` Return the central part of the convolution with the same size as $A$. The central part of the convolution begins at the indices `floor ([size(B)/2] + 1)`.

`"valid"` Return only the parts which do not include zero-padded edges. The size of the result is `max (size (A) - size (B) + 1, 0)`.

When the third argument is a matrix, return the convolution of the matrix $m$ by the vector *v1* in the column direction and by the vector *v2* in the row direction.

**See also:** [conv], page 879, [convn], page 880.

`q = polygcd (b, a)`
`q = polygcd (b, a, tol)`

Find the greatest common divisor of two polynomials.

This is equivalent to the polynomial found by multiplying together all the common roots. Together with deconv, you can reduce a ratio of two polynomials.

The tolerance *tol* defaults to `sqrt (eps)`.

**Caution:** This is a numerically unstable algorithm and should not be used on large polynomials.

Example code:

```
polygcd (poly (1:8), poly (3:12)) - poly (3:8)
⇒ [ 0, 0, 0, 0, 0, 0, 0 ]
deconv (poly (1:8), polygcd (poly (1:8), poly (3:12))) - poly (1:2)
⇒ [ 0, 0, 0 ]
```

**See also:** [poly], page 893, [roots], page 878, [conv], page 879, [deconv], page 880, [residue], page 881.

`[r, p, k, e] = residue (b, a)`
`[b, a] = residue (r, p, k)`
`[b, a] = residue (r, p, k, e)`

The first calling form computes the partial fraction expansion for the quotient of the polynomials, $b$ and $a$.

The quotient is defined as

$$\frac{B(s)}{A(s)} = \sum_{m=1}^{M} \frac{r_m}{(s - p_m)_m^e} + \sum_{i=1}^{N} k_i s^{N-i}.$$

where $M$ is the number of poles (the length of the $r$, $p$, and $e$), the $k$ vector is a polynomial of order $N - 1$ representing the direct contribution, and the $e$ vector specifies the multiplicity of the m-th residue's pole.

For example,

```
b = [1, 1, 1];
a = [1, -5, 8, -4];
[r, p, k, e] = residue (b, a)
   ⇒ r = [-2; 7; 3]
   ⇒ p = [2; 2; 1]
   ⇒ k = [](0x0)
   ⇒ e = [1; 2; 1]
```

which represents the following partial fraction expansion

$$\frac{s^2 + s + 1}{s^3 - 5s^2 + 8s - 4} = \frac{-2}{s - 2} + \frac{7}{(s - 2)^2} + \frac{3}{s - 1}$$

The second calling form performs the inverse operation and computes the reconstituted quotient of polynomials, b(s)/a(s), from the partial fraction expansion; represented by the residues, poles, and a direct polynomial specified by $r$, $p$ and $k$, and the pole multiplicity $e$.

If the multiplicity, $e$, is not explicitly specified the multiplicity is determined by the function `mpoles`.

For example:

```
r = [-2; 7; 3];
p = [2; 2; 1];
k = [1, 0];
[b, a] = residue (r, p, k)
   ⇒ b = [1, -5, 9, -3, 1]
   ⇒ a = [1, -5, 8, -4]

where mpoles is used to determine e = [1; 2; 1]
```

Alternatively the multiplicity may be defined explicitly, for example,

```
r = [7; 3; -2];
p = [2; 1; 2];
k = [1, 0];
e = [2; 1; 1];
[b, a] = residue (r, p, k, e)
   ⇒ b = [1, -5, 9, -3, 1]
   ⇒ a = [1, -5, 8, -4]
```

which represents the following partial fraction expansion

$$\frac{-2}{s-2} + \frac{7}{(s-2)^2} + \frac{3}{s-1} + s = \frac{s^4 - 5s^3 + 9s^2 - 3s + 1}{s^3 - 5s^2 + 8s - 4}$$

**See also:** [mpoles], page 879, [poly], page 893, [roots], page 878, [conv], page 879, [deconv], page 880.

## 28.4 Derivatives / Integrals / Transforms

Octave comes with functions for computing the derivative and the integral of a polynomial. The functions `polyder` and `polyint` both return new polynomials describing the result. As an example we'll compute the definite integral of $p(x) = x^2 + 1$ from 0 to 3.

```
c = [1, 0, 1];
integral = polyint (c);
area = polyval (integral, 3) - polyval (integral, 0)
⇒ 12
```

*k* = polyder (*p*)
*k* = polyder (*a*, *b*)
[*q*, *d*] = polyder (*b*, *a*)

Return the coefficients of the derivative of the polynomial whose coefficients are given by the vector *p*.

If a pair of polynomials is given, return the derivative of the product $a * b$.

If two inputs and two outputs are given, return the derivative of the polynomial quotient $b/a$. The quotient numerator is in $q$ and the denominator in $d$.

**See also:** [polyint], page 883, [polyval], page 877, [polyreduce], page 893.

`q = polyint (p)`
`q = polyint (p, k)`
Return the coefficients of the integral of the polynomial whose coefficients are represented by the vector *p*.

The variable *k* is the constant of integration, which by default is set to zero.

**See also:** [polyder], page 882, [polyval], page 877.

`g = polyaffine (f, mu)`
Return the coefficients of the polynomial vector *f* after an affine transformation.

If *f* is the vector representing the polynomial f(x), then `g = polyaffine (f, mu)` is the vector representing:

```
g(x) = f( (x - mu(1)) / mu(2) )
```

**See also:** [polyval], page 877, [polyfit], page 883.

## 28.5 Polynomial Interpolation

Octave comes with good support for various kinds of interpolation, most of which are described in Chapter 29 [Interpolation], page 895. One simple alternative to the functions described in the aforementioned chapter, is to fit a single polynomial, or a piecewise polynomial (spline) to some given data points. To avoid a highly fluctuating polynomial, one most often wants to fit a low-order polynomial to data. This usually means that it is necessary to fit the polynomial in a least-squares sense, which just is what the `polyfit` function does.

`p = polyfit (x, y, n)`
`[p, s] = polyfit (x, y, n)`
`[p, s, mu] = polyfit (x, y, n)`
Return the coefficients of a polynomial $p(x)$ of degree $n$ that minimizes the least-squares-error of the fit to the points `[x(:), y(:)]`.

$n$ is typically an integer $\geq 0$ specifying the degree of the approximating polynomial. If $n$ is a logical vector, it is used as a mask to selectively force the corresponding polynomial coefficients to be used or ignored.

The polynomial coefficients are returned in the row vector *p*. The output *p* may be directly used with `polyval` to estimate values using the fitted polynomial.

The optional output *s* is a structure containing the following fields:

'`yf`' The values of the polynomial for each value of *x*.

'`V`' The Vandermonde matrix used to compute the polynomial coefficients.

'`X (deprecated, will be removed in Octave 13)`'
The Vandermonde matrix used to compute the polynomial coefficients.

'`R`' Triangular factor R from the QR decomposition.

'`C`' The unscaled covariance matrix, formally equal to the inverse of *v*'**v*, but computed in a way minimizing roundoff error propagation.

'`df`' The degrees of freedom.

'normr' The norm of the residuals.

'rsquared'
Coefficient of determination, or R-squared.

The second output may be used by `polyval` to calculate the statistical error limits of the predicted values. In particular, the standard deviation of $p$ coefficients is given by

`sqrt (diag (s.C)/s.df) * s.normr`.

When the third output, *mu*, is present the original data is centered and scaled which can improve the numerical stability of the fit. The coefficients $p$ are associated with a polynomial in

`xhat = (x - mu(1)) / mu(2)`
where $mu(1)$ = mean ($x$), and $mu(2)$ = std ($x$).

Example 1 : logical $n$ and integer $n$

```
f = @(x) x.^2 + 5;   # data-generating function
x = 0:5;
y = f (x);
## Fit data to polynomial A*x^3 + B*x^1
p = polyfit (x, y, logical ([1, 0, 1, 0]))
⇒ p = [ 0.0680, 0, 4.2444, 0 ]
## Fit data to polynomial using all terms up to x^3
p = polyfit (x, y, 3)
⇒ p = [ -4.9608e-17, 1.0000e+00, -1.6906e-15, 5.0000e+00 ]
```

Programming Note: If the desired polynomial degree $n \geq$ number of data points, the solution is not unique. Instead, the fitting calculates `m = numel (x) - 1` terms. The return polynomial has coefficients for $x$^$n$, $x$^$n$-1, . . . , $x$^$n$-$m$, and the constant term $x$^0, with all other coefficients set to 0.

**See also:** [polyval], page 877, [polyaffine], page 883, [roots], page 878, [vander], page 600, [zscore], page 858.

In situations where a single polynomial isn't good enough, a solution is to use several polynomials pieced together. The function `splinefit` fits a piecewise polynomial (spline) to a set of data.

`pp = splinefit (x, y, breaks)`
`pp = splinefit (x, y, p)`
`pp = splinefit (. . . , "periodic", periodic)`
`pp = splinefit (. . . , "robust", robust)`
`pp = splinefit (. . . , "beta", beta)`
`pp = splinefit (. . . , "order", order)`
`pp = splinefit (. . . , "constraints", constraints)`

Fit a piecewise cubic spline with breaks (knots) *breaks* to the noisy data, $x$ and $y$.

$x$ is a vector, and $y$ is a vector or N-D array. If $y$ is an N-D array, then $x$(j) is matched to $y$(:,. . .,:,j).

$p$ is a positive integer defining the number of intervals along $x$, and $p$+1 is the number of breaks. The number of points in each interval differ by no more than 1.

The optional property *periodic* is a logical value which specifies whether a periodic boundary condition is applied to the spline. The length of the period is `max (breaks) - min (breaks)`. The default value is `false`.

The optional property *robust* is a logical value which specifies if robust fitting is to be applied to reduce the influence of outlying data points. Three iterations of weighted least squares are performed. Weights are computed from previous residuals. The sensitivity of outlier identification is controlled by the property *beta*. The value of *beta* is restricted to the range, 0 < *beta* < 1. The default value is *beta* = 1/2. Values close to 0 give all data equal weighting. Increasing values of *beta* reduce the influence of outlying data. Values close to unity may cause instability or rank deficiency.

The fitted spline is returned as a piecewise polynomial, *pp*, and may be evaluated using `ppval`.

The splines are constructed of polynomials with degree *order*. The default is a cubic, *order*=3. A spline with P pieces has P+*order* degrees of freedom. With periodic boundary conditions the degrees of freedom are reduced to P.

The optional property, *constraints*, is a structure specifying linear constraints on the fit. The structure has three fields, `"xc"`, `"yc"`, and `"cc"`.

`"xc"` Vector of the x-locations of the constraints.

`"yc"` Constraining values at the locations *xc*. The default is an array of zeros.

`"cc"` Coefficients (matrix). The default is an array of ones. The number of rows is limited to the order of the piecewise polynomials, *order*.

Constraints are linear combinations of derivatives of order 0 to *order*-1 according to

$$cc(1,j) \cdot y(xc(j)) + cc(2,j) \cdot y\prime(xc(j)) + \ldots = yc(:,\ldots,:,j).$$

**See also:** [interp1], page 895, [unmkpp], page 891, [ppval], page 892, [spline], page 898, [pchip], page 934, [ppder], page 892, [ppint], page 892, [ppjumps], page 892.

The number of *breaks* (or knots) used to construct the piecewise polynomial is a significant factor in suppressing the noise present in the input data, *x* and *y*. This is demonstrated by the example below.

```
x = 2 * pi * rand (1, 200);
y = sin (x) + sin (2 * x) + 0.2 * randn (size (x));
## Uniform breaks
breaks = linspace (0, 2 * pi, 41); % 41 breaks, 40 pieces
pp1 = splinefit (x, y, breaks);
## Breaks interpolated from data
pp2 = splinefit (x, y, 10);  % 11 breaks, 10 pieces
## Plot
xx = linspace (0, 2 * pi, 400);
y1 = ppval (pp1, xx);
y2 = ppval (pp2, xx);
plot (x, y, ".", xx, [y1; y2])
axis tight
ylim auto
legend ({"data", "41 breaks, 40 pieces", "11 breaks, 10 pieces"})
```

The result of which can be seen in Figure 28.1.

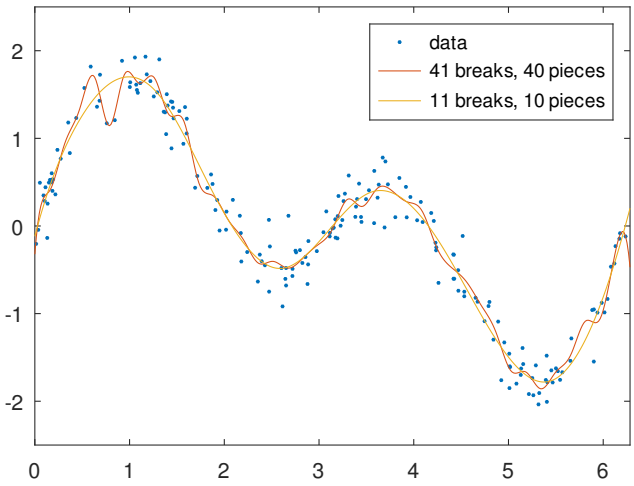



Figure 28.1: Comparison of a fitting a piecewise polynomial with 41 breaks to one with 11 breaks. The fit with the large number of breaks exhibits a fast ripple that is not present in the underlying function.

The piecewise polynomial fit, provided by `splinefit`, has continuous derivatives up to the *order*-1. For example, a cubic fit has continuous first and second derivatives. This is demonstrated by the code

```
## Data (200 points)
x = 2 * pi * rand (1, 200);
y = sin (x) + sin (2 * x) + 0.1 * randn (size (x));
## Piecewise constant
pp1 = splinefit (x, y, 8, "order", 0);
## Piecewise linear
pp2 = splinefit (x, y, 8, "order", 1);
## Piecewise quadratic
pp3 = splinefit (x, y, 8, "order", 2);
## Piecewise cubic
pp4 = splinefit (x, y, 8, "order", 3);
## Piecewise quartic
pp5 = splinefit (x, y, 8, "order", 4);
## Plot
xx = linspace (0, 2 * pi, 400);
y1 = ppval (pp1, xx);
y2 = ppval (pp2, xx);
y3 = ppval (pp3, xx);
y4 = ppval (pp4, xx);
y5 = ppval (pp5, xx);
plot (x, y, ".", xx, [y1; y2; y3; y4; y5])
```

```
axis tight
ylim auto
legend ({"data", "order 0", "order 1", "order 2", "order 3", "order 4"})
```

The result of which can be seen in Figure 28.2.

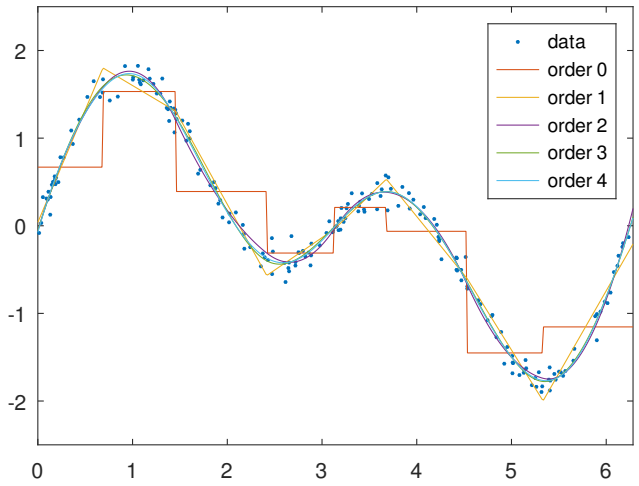



Figure 28.2: Comparison of a piecewise constant, linear, quadratic, cubic, and quartic polynomials with 8 breaks to noisy data. The higher order solutions more accurately represent the underlying function, but come with the expense of computational complexity.

When the underlying function to provide a fit to is periodic, `splinefit` is able to apply the boundary conditions needed to manifest a periodic fit. This is demonstrated by the code below.

```
## Data (100 points)
x = 2 * pi * [0, (rand (1, 98)), 1];
y = sin (x) - cos (2 * x) + 0.2 * randn (size (x));
## No constraints
pp1 = splinefit (x, y, 10, "order", 5);
## Periodic boundaries
pp2 = splinefit (x, y, 10, "order", 5, "periodic", true);
## Plot
xx = linspace (0, 2 * pi, 400);
y1 = ppval (pp1, xx);
y2 = ppval (pp2, xx);
plot (x, y, ".", xx, [y1; y2])
axis tight
ylim auto
legend ({"data", "no constraints", "periodic"})
```

The result of which can be seen in Figure 28.3.

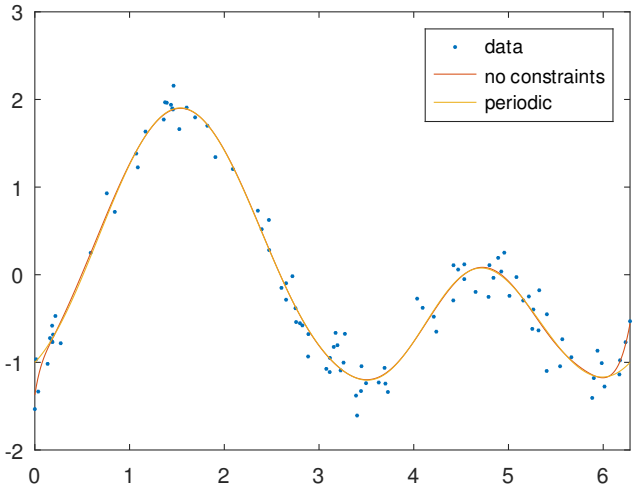



Figure 28.3: Comparison of piecewise polynomial fits to a noisy periodic function with, and without, periodic boundary conditions.

More complex constraints may be added as well. For example, the code below illustrates a periodic fit with values that have been clamped at the endpoints, and a second periodic fit which is hinged at the endpoints.

```
## Data (200 points)
x = 2 * pi * rand (1, 200);
y = sin (2 * x) + 0.1 * randn (size (x));
## Breaks
breaks = linspace (0, 2 * pi, 10);
## Clamped endpoints, y = y' = 0
xc = [0, 0, 2*pi, 2*pi];
cc = [(eye (2)), (eye (2))];
con = struct ("xc", xc, "cc", cc);
pp1 = splinefit (x, y, breaks, "constraints", con);
## Hinged periodic endpoints, y = 0
con = struct ("xc", 0);
pp2 = splinefit (x, y, breaks, "constraints", con, "periodic", true);
## Plot
xx = linspace (0, 2 * pi, 400);
y1 = ppval (pp1, xx);
y2 = ppval (pp2, xx);
plot (x, y, ".", xx, [y1; y2])
axis tight
ylim auto
legend ({"data", "clamped", "hinged periodic"})
```

The result of which can be seen in Figure 28.4.

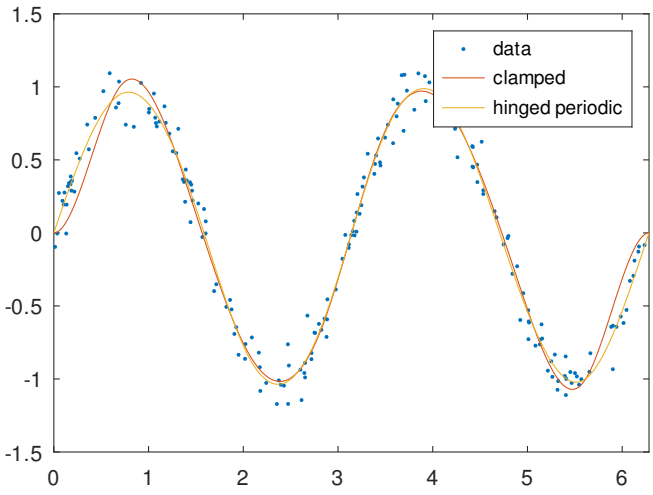



Figure 28.4: Comparison of two periodic piecewise cubic fits to a noisy periodic signal. One fit has its endpoints clamped and the second has its endpoints hinged.

The `splinefit` function also provides the convenience of a *robust* fitting, where the effect of outlying data is reduced. In the example below, three different fits are provided. Two with differing levels of outlier suppression and a third illustrating the non-robust solution.

```
## Data
x = linspace (0, 2*pi, 200);
y = sin (x) + sin (2 * x) + 0.05 * randn (size (x));
## Add outliers
x = [x, linspace(0,2*pi,60)];
y = [y, -ones(1,60)];
## Fit splines with hinged conditions
con = struct ("xc", [0, 2*pi]);
## Robust fitting, beta = 0.25
pp1 = splinefit (x, y, 8, "constraints", con, "beta", 0.25);
## Robust fitting, beta = 0.75
pp2 = splinefit (x, y, 8, "constraints", con, "beta", 0.75);
## No robust fitting
pp3 = splinefit (x, y, 8, "constraints", con);
## Plot
xx = linspace (0, 2*pi, 400);
y1 = ppval (pp1, xx);
y2 = ppval (pp2, xx);
y3 = ppval (pp3, xx);
plot (x, y, ".", xx, [y1; y2; y3])
legend ({"data with outliers","robust, beta = 0.25", ...
         "robust, beta = 0.75", "no robust fitting"})
axis tight
```

```
ylim auto
```

The result of which can be seen in Figure 28.5.

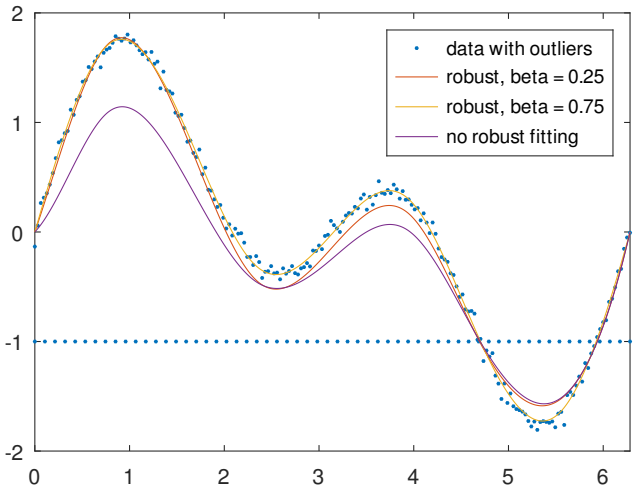



Figure 28.5: Comparison of two different levels of robust fitting ($beta = 0.25$ and 0.75) to noisy data combined with outlying data. A conventional fit, without robust fitting ($beta = 0$) is also included.

A very specific form of polynomial interpretation is the Padé approximant. For control systems, a continuous-time delay can be modeled very simply with the approximant.

`[num, den] = padecoef (T)`
`[num, den] = padecoef (T, N)`

Compute the $N$th-order Padé approximant of the continuous-time delay $T$ in transfer function form. The Padé approximant of $e^{-sT}$ is defined by the following equation

$$e^{-sT} \approx \frac{P_n(s)}{Q_n(s)}$$

where both $P_n(s)$ and $Q_n(s)$ are $N^{th}$-order rational functions defined by the following expressions

$$P_n(s) = \sum_{k=0}^{N} \frac{(2N-k)!N!}{(2N)!k!(N-k)!}(-sT)^k$$

$$Q_n(s) = P_n(-s)$$

The inputs $T$ and $N$ must be non-negative numeric scalars. If $N$ is unspecified it defaults to 1.

The output row vectors *num* and *den* contain the numerator and denominator coefficients in descending powers of s. Both are $N$th-order polynomials.

For example:

```
t = 0.1;
n = 4;
[num, den] = padecoef (t, n)
⇒ num =

      1.0000e-04  -2.0000e-02   1.8000e+00  -8.4000e+01   1.6800e+03

⇒ den =

      1.0000e-04   2.0000e-02   1.8000e+00   8.4000e+01   1.6800e+03
```

The function, `ppval`, evaluates the piecewise polynomials, created by `mkpp` or other means, and `unmkpp` returns detailed information about the piecewise polynomial.

The following example shows how to combine two linear functions and a quadratic into one function. Each of these functions is expressed on adjoined intervals.

```
x = [-2, -1, 1, 2];
p = [ 0,  1, 0;
      1, -2, 1;
      0, -1, 1 ];
pp = mkpp (x, p);
xi = linspace (-2, 2, 50);
yi = ppval (pp, xi);
plot (xi, yi);
```

*pp* = mkpp (*breaks*, *coefs*)
*pp* = mkpp (*breaks*, *coefs*, *d*)

Construct a piecewise polynomial (pp) structure from sample points *breaks* and coefficients *coefs*.

*breaks* must be a vector of strictly increasing values. The number of intervals is given by *ni* = `length (`*breaks*`) - 1`.

When *m* is the polynomial order *coefs* must be of size: *ni*-by-(*m* + 1).

The i-th row of *coefs*, `coefs(`*i*`,:)`, contains the coefficients for the polynomial over the *i*-th interval, ordered from highest (*m*) to lowest (*0*) degree.

*coefs* may also be a multi-dimensional array, specifying a vector-valued or array-valued polynomial. In that case the polynomial order *m* is defined by the length of the last dimension of *coefs*. The size of first dimension(s) are given by the scalar or vector *d*. If *d* is not given it is set to `1`. In this case `p(`*r*`, `*i*`, :)` contains the coefficients for the *r*-th polynomial defined on interval *i*. In any case *coefs* is reshaped to a 2-D matrix of size `[`*ni*`*prod(`*d*`) `*m*`]`.

Programming Note: `ppval` evaluates polynomials at *xi* `-` *breaks*`(i)`, i.e., it subtracts the lower endpoint of the current interval from *xi*. This must be taken into account when creating piecewise polynomials objects with `mkpp`.

**See also:** [unmkpp], page 891, [ppval], page 892, [spline], page 898, [pchip], page 934, [ppder], page 892, [ppint], page 892, [ppjumps], page 892.

[*x*, *p*, *n*, *k*, *d*] = unmkpp (*pp*)

Extract the components of a piecewise polynomial structure *pp*.

This function is the inverse of `mkpp`: it extracts the inputs to `mkpp` needed to create the piecewise polynomial structure *pp*. The code below makes this relation explicit:

```
[breaks, coefs, numinter, order, dim] = unmkpp (pp);
pp2  = mkpp (breaks, coefs, dim);
```

The piecewise polynomial structure `pp2` obtained in this way, is identical to the original `pp`. The same can be obtained by directly accessing the fields of the structure `pp`.

The components are:

*x* Sample points or breaks.

*p* Polynomial coefficients for points in sample interval. `p(i, :)` contains the coefficients for the polynomial over interval *i* ordered from highest to lowest degree. If `d > 1`, then *p* is a matrix of size `[n*prod(d) m]`, where the `i + (1:d)` rows are the coefficients of all the *d* polynomials in the interval *i*.

*n* Number of polynomial pieces or intervals, `n = length (x) - 1`.

*k* Order of the polynomial plus 1.

*d* Number of polynomials defined for each interval.

**See also:** [mkpp], page 891, [ppval], page 892, [spline], page 898, [pchip], page 934.

`yi = ppval (pp, xi)`

Evaluate the piecewise polynomial structure *pp* at the points *xi*.

If *pp* describes a scalar polynomial function, the result is an array of the same shape as *xi*. Otherwise, the size of the result is `[pp.dim, length(xi)]` if *xi* is a vector, or `[pp.dim, size(xi)]` if it is a multi-dimensional array.

**See also:** [mkpp], page 891, [unmkpp], page 891, [spline], page 898, [pchip], page 934.

`ppd = ppder (pp)`
`ppd = ppder (pp, m)`

Compute the piecewise *m*-th derivative of a piecewise polynomial struct *pp*.

If *m* is omitted the first derivative is calculated.

**See also:** [mkpp], page 891, [ppval], page 892, [ppint], page 892.

`ppi = ppint (pp)`
`ppi = ppint (pp, c)`

Compute the integral of the piecewise polynomial struct *pp*.

*c*, if given, is the constant of integration.

**See also:** [mkpp], page 891, [ppval], page 892, [ppder], page 892.

`jumps = ppjumps (pp)`

Evaluate the boundary jumps of a piecewise polynomial.

If there are *n* intervals, and the dimensionality of *pp* is *d*, the resulting array has dimensions `[d, n-1]`.

**See also:** [mkpp], page 891.

## 28.6 Miscellaneous Functions

`y = poly (A)`
`y = poly (x)`

If $A$ is a square $N$-by-$N$ matrix, `poly (A)` is the row vector of the coefficients of `det (z * eye (N) - A)`, the characteristic polynomial of $A$.

For example, the following code finds the eigenvalues of $A$ which are the roots of `poly (A)`.

```
roots (poly (eye (3)))
    ⇒ 1.00001 + 0.00001i
      1.00001 - 0.00001i
      0.99999 + 0.00000i
```

In fact, all three eigenvalues are exactly 1 which emphasizes that for numerical performance the `eig` function should be used to compute eigenvalues.

If $x$ is a vector, `poly (x)` is a vector of the coefficients of the polynomial whose roots are the elements of $x$. That is, if $c$ is a polynomial, then the elements of `d = roots (poly (c))` are contained in $c$. The vectors $c$ and $d$ are not identical, however, due to sorting and numerical errors.

**See also:** [roots], page 878, [eig], page 649.

`polyout (c)`
`polyout (c, x)`
`str = polyout (...)`

Display a formatted version of the polynomial $c$.

The formatted polynomial

$$c(x) = c_1 x^n + \ldots + c_n x + c_{n+1}$$

is returned as a string or written to the screen if `nargout` is zero.

The second argument $x$ specifies the variable name to use for each term and defaults to the string `"s"`.

**See also:** [polyreduce], page 893.

`p = polyreduce (c)`

Reduce a polynomial coefficient vector to a minimum number of terms by stripping off any leading zeros.

**See also:** [polyout], page 893.

# 29 Interpolation

## 29.1 One-dimensional Interpolation

Octave supports several methods for one-dimensional interpolation, most of which are described in this section. Section 28.5 [Polynomial Interpolation], page 883, and Section 30.4 [Interpolation on Scattered Data], page 918, describe additional methods.

*yi* = `interp1` (*x*, *y*, *xi*)
*yi* = `interp1` (*y*, *xi*)
*yi* = `interp1` (..., *method*)
*yi* = `interp1` (..., *extrap*)
*yi* = `interp1` (..., "*left*")
*yi* = `interp1` (..., "*right*")
*pp* = `interp1` (..., "*pp*")

One-dimensional interpolation.

Interpolate input data to determine the value of *yi* at the points *xi*. If not specified, *x* is taken to be the indices of *y* (`1:length (y)`). If *y* is a matrix or an N-dimensional array, the interpolation is performed on each column of *y*.

The interpolation *method* is one of:

`"nearest"`
: Return the nearest neighbor.

`"previous"`
: Return the previous neighbor.

`"next"`
: Return the next neighbor.

`"linear"` (default)
: Linear interpolation from nearest neighbors.

`"pchip"`
: Piecewise cubic Hermite interpolating polynomial—shape-preserving interpolation with smooth first derivative.

`"cubic"`
: Cubic interpolation (same as `"pchip"`).

`"spline"`
: Cubic spline interpolation—smooth first and second derivatives throughout the curve.

Adding '*' to the start of any method above forces `interp1` to assume that *x* is uniformly spaced, and only `x(1)` and `x(2)` are referenced. This is usually faster, and is never slower. The default method is `"linear"`.

If *extrap* is the string `"extrap"`, then extrapolate values beyond the endpoints using the current *method*. If *extrap* is a number, then replace values beyond the endpoints with that number. When unspecified, *extrap* defaults to `NA`.

If the string argument `"pp"` is specified, then *xi* should not be supplied and `interp1` returns a piecewise polynomial object. This object can later be used with `ppval` to evaluate the interpolation. There is an equivalence, such that `ppval (interp1 (x, y, method, "pp"), xi) == interp1 (x, y, xi, method, "extrap")`.

Duplicate points in $x$ specify a discontinuous interpolant. There may be at most 2 consecutive points with the same value. If $x$ is increasing, the default discontinuous interpolant is right-continuous. If $x$ is decreasing, the default discontinuous interpolant is left-continuous. The continuity condition of the interpolant may be specified by using the options `"left"` or `"right"` to select a left-continuous or right-continuous interpolant, respectively. Discontinuous interpolation is only allowed for `"nearest"` and `"linear"` methods; in all other cases, the $x$-values must be unique.

An example of the use of `interp1` is

```
xf = [0:0.05:10];
yf = sin (2*pi*xf/5);
xp = [0:10];
yp = sin (2*pi*xp/5);
lin = interp1 (xp, yp, xf);
near = interp1 (xp, yp, xf, "nearest");
pch = interp1 (xp, yp, xf, "pchip");
spl = interp1 (xp, yp, xf, "spline");
plot (xf,yf,"r", xf,near,"g", xf,lin,"b", xf,pch,"c", xf,spl,"m",
      xp,yp,"r*");
legend ("original", "nearest", "linear", "pchip", "spline");
```

**See also:** [pchip], page 934, [spline], page 898, [interpft], page 897, [interp2], page 899, [interp3], page 900, [interpn], page 901.

There are some important differences between the various interpolation methods. The `"spline"` method enforces that both the first and second derivatives of the interpolated values have a continuous derivative, whereas the other methods do not. This means that the results of the `"spline"` method are generally smoother. If the function to be interpolated is in fact smooth, then `"spline"` will give excellent results. However, if the function to be evaluated is in some manner discontinuous, then `"pchip"` interpolation might give better results.

This can be demonstrated by the code

```
t = -2:2;
dt = 1;
ti =-2:0.025:2;
dti = 0.025;
y = sign (t);
ys = interp1 (t,y,ti,"spline");
yp = interp1 (t,y,ti,"pchip");
ddys = diff (diff (ys)./dti) ./ dti;
ddyp = diff (diff (yp)./dti) ./ dti;
figure (1);
plot (ti,ys,"r-", ti,yp,"g-");
legend ("spline", "pchip", 4);
figure (2);
plot (ti,ddys,"r+", ti,ddyp,"g*");
legend ("spline", "pchip");
```

The result of which can be seen in Figure 29.1 and Figure 29.2.

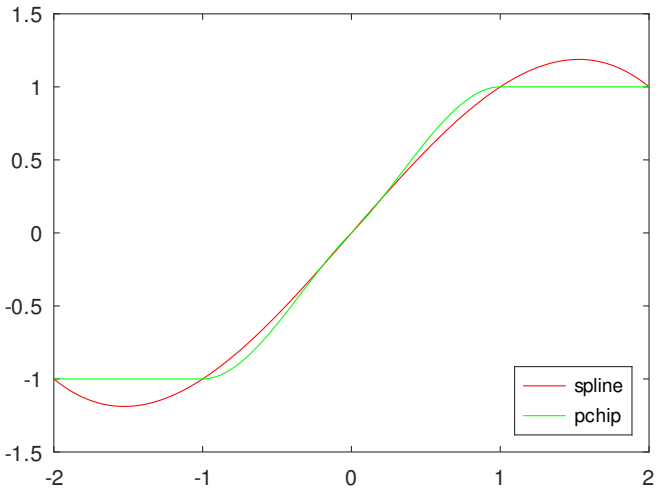



Figure 29.1: Comparison of `"pchip"` and `"spline"` interpolation methods for a step function

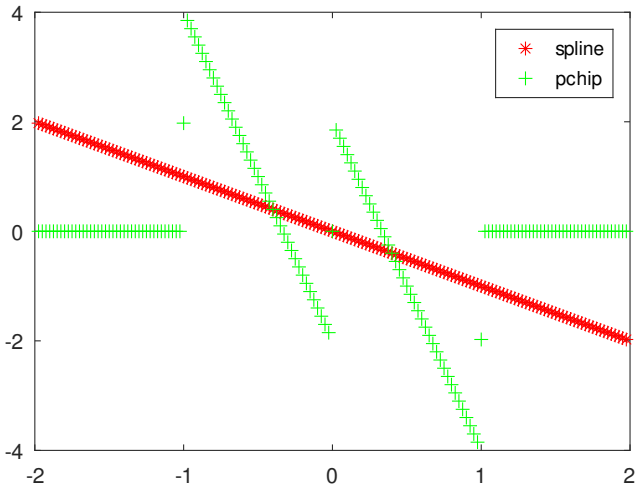



Figure 29.2: Comparison of the second derivative of the `"pchip"` and `"spline"` interpolation methods for a step function

Fourier interpolation, is a resampling technique where a signal is converted to the frequency domain, padded with zeros and then reconverted to the time domain.

`y = interpft (x, n)`
`y = interpft (x, n, dim)`
Fourier interpolation.

If $x$ is a vector then $x$ is resampled with $n$ points. The data in $x$ is assumed to be equispaced. If $x$ is a matrix or an N-dimensional array, the interpolation is performed on each column of $x$.

If *dim* is specified, then interpolate along the dimension *dim*.

`interpft` assumes that the interpolated function is periodic, and so assumptions are made about the endpoints of the interpolation.

**See also:** [interp1], page 895.

There are two significant limitations on Fourier interpolation. First, the function signal is assumed to be periodic, and so non-periodic signals will be poorly represented at the edges. Second, both the signal and its interpolation are required to be sampled at equispaced points. An example of the use of `interpft` is

```
t = 0 : 0.3 : pi; dt = t(2)-t(1);
n = length (t); k = 100;
ti = t(1) + [0 : k-1]*dt*n/k;
y = sin (4*t + 0.3) .* cos (3*t - 0.1);
yp = sin (4*ti + 0.3) .* cos (3*ti - 0.1);
plot (ti, yp, "g", ti, interp1 (t, y, ti, "spline"), "b", ...
      ti, interpft (y, k), "c", t, y, "r+");
legend ("sin(4t+0.3)cos(3t-0.1)", "spline", "interpft", "data");
```

which demonstrates the poor behavior of Fourier interpolation for non-periodic functions, as can be seen in Figure 29.3.

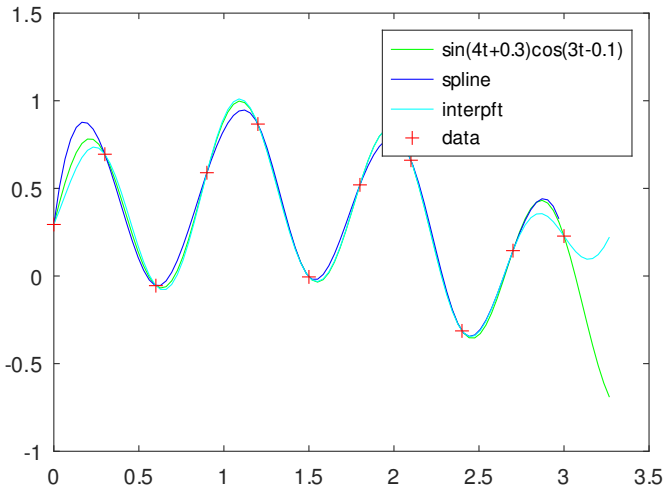



Figure 29.3: Comparison of `interp1` and `interpft` for non-periodic data

In addition, the support functions `spline` and `lookup` that underlie the `interp1` function can be called directly.

`pp = spline (x, y)`
`yi = spline (x, y, xi)`

Return the cubic spline interpolant of points $x$ and $y$.

When called with two arguments, return the piecewise polynomial $pp$ that may be used with `ppval` to evaluate the polynomial at specific points.

When called with a third input argument, `spline` evaluates the spline at the points $xi$. The third calling form `spline (x, y, xi)` is equivalent to `ppval (spline (x, y), xi)`.

The variable $x$ must be a vector of length $n$.

$y$ can be either a vector or array. If $y$ is a vector it must have a length of either $n$ or `n + 2`. If the length of $y$ is $n$, then the `"not-a-knot"` end condition is used. If the length of $y$ is `n + 2`, then the first and last values of the vector $y$ are the values of the first derivative of the cubic spline at the endpoints.

If $y$ is an array, then the size of $y$ must have the form

$$[s_1, s_2, \cdots, s_k, n]$$

or

$$[s_1, s_2, \cdots, s_k, n + 2].$$

The array is reshaped internally to a matrix where the leading dimension is given by

$$s_1 s_2 \cdots s_k$$

and each row of this matrix is then treated separately. Note that this is exactly the opposite of `interp1` but is done for MATLAB compatibility.

**See also:** [pchip], page 934, [ppval], page 892, [mkpp], page 891, [unmkpp], page 891.

## 29.2 Multi-dimensional Interpolation

There are three multi-dimensional interpolation functions in Octave, with similar capabilities. Methods using Delaunay tessellation are described in Section 30.4 [Interpolation on Scattered Data], page 918.

`zi = interp2 (x, y, z, xi, yi)`
`zi = interp2 (z, xi, yi)`
`zi = interp2 (z, n)`
`zi = interp2 (z)`
`zi = interp2 (..., method)`
`zi = interp2 (..., method, extrap)`

Two-dimensional interpolation.

Interpolate numeric reference data $x$, $y$, $z$ to determine $zi$ at the coordinates $xi$, $yi$. The reference data $x$, $y$ can be matrices, as returned by `meshgrid`, in which case the sizes of $x$, $y$, and $z$ must be equal. If $x$, $y$ are vectors describing a grid then `length (x) == columns (z)` and `length (y) == rows (z)`. In either case the input data must be strictly monotonic.

If called without *x*, *y*, and just a single reference data matrix *z*, the 2-D region `x = 1:columns (z), y = 1:rows (z)` is assumed. This saves memory if the grid is regular and the distance between points is not important.

If called with a single reference data matrix *z* and a refinement value *n*, then perform interpolation over a grid where each original interval has been recursively subdivided *n* times. This results in `2^n-1` additional points for every interval in the original grid. If *n* is omitted a value of 1 is used. As an example, the interval [0,1] with `n==2` results in a refined interval with points at [0, 1/4, 1/2, 3/4, 1].

The interpolation *method* is one of:

`"nearest"`
: Return the nearest neighbor.

`"linear"` (default)
: Linear interpolation from nearest neighbors.

`"pchip"`
: Piecewise cubic Hermite interpolating polynomial—shape-preserving interpolation with smooth first derivative.

`"cubic"`
: Cubic interpolation using a convolution kernel function—third order method with smooth first derivative.

`"spline"`
: Cubic spline interpolation—smooth first and second derivatives throughout the curve.

*extrap* is a scalar number. It replaces values beyond the endpoints with *extrap*. Note that if *extrap* is used, *method* must be specified as well. If *extrap* is omitted and the *method* is `"spline"`, then the extrapolated values of the `"spline"` are used. Otherwise the default *extrap* value for any other *method* is `"NA"`.

**See also:** [interp1], page 895, [interp3], page 900, [interpn], page 901, [meshgrid], page 398.

```
vi = interp3 (x, y, z, v, xi, yi, zi)
vi = interp3 (v, xi, yi, zi)
vi = interp3 (v, n)
vi = interp3 (v)
vi = interp3 (..., method)
vi = interp3 (..., method, extrapval)
```

Three-dimensional interpolation.

Interpolate numeric reference data *x*, *y*, *z*, *v* to determine *vi* at the coordinates *xi*, *yi*, *zi*. The reference data *x*, *y*, *z* can be matrices, as returned by `meshgrid`, in which case the sizes of *x*, *y*, *z*, and *v* must be equal. If *x*, *y*, *z* are vectors describing a cubic grid then `length (x) == columns (v)`, `length (y) == rows (v)`, and `length (z) == size (v, 3)`. In either case the input data must be strictly monotonic.

If called without *x*, *y*, *z*, and just a single reference data matrix *v*, the 3-D region `x = 1:columns (v), y = 1:rows (v), z = 1:size (v, 3)` is assumed. This saves memory if the grid is regular and the distance between points is not important.

If called with a single reference data matrix *v* and a refinement value *n*, then perform interpolation over a 3-D grid where each original interval has been recursively

subdivided *n* times. This results in `2^n-1` additional points for every interval in the original grid. If *n* is omitted a value of 1 is used. As an example, the interval [0,1] with `n==2` results in a refined interval with points at [0, 1/4, 1/2, 3/4, 1].

The interpolation *method* is one of:

`"nearest"`
: Return the nearest neighbor.

`"linear"` (default)
: Linear interpolation from nearest neighbors.

`"cubic"`
: Piecewise cubic Hermite interpolating polynomial—shape-preserving interpolation with smooth first derivative (not implemented yet).

`"spline"`
: Cubic spline interpolation—smooth first and second derivatives throughout the curve.

*extrapval* is a scalar number. It replaces values beyond the endpoints with *extrapval*. Note that if *extrapval* is used, *method* must be specified as well. If *extrapval* is omitted and the *method* is `"spline"`, then the extrapolated values of the `"spline"` are used. Otherwise the default *extrapval* value for any other *method* is `"NA"`.

**See also:** [interp1], page 895, [interp2], page 899, [interpn], page 901, [meshgrid], page 398.

`vi = interpn (x1, x2, ..., v, y1, y2, ...)`
`vi = interpn (v, y1, y2, ...)`
`vi = interpn (v, m)`
`vi = interpn (v)`
`vi = interpn (..., method)`
`vi = interpn (..., method, extrapval)`

Perform *n*-dimensional interpolation, where *n* is at least two.

Each element of the *n*-dimensional numeric array *v* represents a value at a location given by the parameters *x1*, *x2*, ..., *xn*. The parameters *x1*, *x2*, ..., *xn* are either *n*-dimensional arrays of the same size as the array *v* in the `"ndgrid"` format or vectors.

The parameters *y1*, *y2*, ..., *yn* represent the points at which the array *vi* is interpolated. They can be vectors of the same length and orientation in which case they are interpreted as coordinates of scattered points. If they are vectors of differing orientation or length, they are used to form a grid in `"ndgrid"` format. They can also be *n*-dimensional arrays of equal size.

If *x1*, ..., *xn* are omitted, they are assumed to be `x1 = 1 : size (v, 1)`, etc. If *m* is specified, then the interpolation adds a point half way between each of the interpolation points. This process is performed *m* times. If only *v* is specified, then *m* is assumed to be `1`.

The interpolation *method* is one of:

`"nearest"`
: Return the nearest neighbor.

`"linear"` (default)
: Linear interpolation from nearest neighbors.

`"pchip"` Piecewise cubic Hermite interpolating polynomial—shape-preserving interpolation with smooth first derivative (not implemented yet).

`"cubic"` Cubic interpolation (same as `"pchip"` [not implemented yet]).

`"spline"` Cubic spline interpolation—smooth first and second derivatives throughout the curve.

The default method is `"linear"`.

*extrapval* is a scalar number. It replaces values beyond the endpoints with *extrapval*. Note that if *extrapval* is used, *method* must be specified as well. If *extrapval* is omitted and the *method* is `"spline"`, then the extrapolated values of the `"spline"` are used. Otherwise the default *extrapval* value for any other *method* is `NA`.

**See also:** [interp1], page 895, [interp2], page 899, [interp3], page 900, [spline], page 898, [ndgrid], page 399.

A significant difference between `interpn` and the other two multi-dimensional interpolation functions is the fashion in which the dimensions are treated. For `interp2` and `interp3`, the y-axis is considered to be the columns of the matrix, whereas the x-axis corresponds to the rows of the array. As Octave indexes arrays in column major order, the first dimension of any array is the columns, and so `interpn` effectively reverses the 'x' and 'y' dimensions. Consider the example,

```
x = y = z = -1:1;
f = @(x,y,z) x.^2 - y - z.^2;
[xx, yy, zz] = meshgrid (x, y, z);
v = f (xx,yy,zz);
xi = yi = zi = -1:0.1:1;
[xxi, yyi, zzi] = meshgrid (xi, yi, zi);
vi = interp3 (x, y, z, v, xxi, yyi, zzi, "spline");
[xxi, yyi, zzi] = ndgrid (xi, yi, zi);
vi2 = interpn (x, y, z, v, xxi, yyi, zzi, "spline");
mesh (zi, yi, squeeze (vi2(1,:,:)));
```

where `vi` and `vi2` are identical. The reversal of the dimensions is treated in the `meshgrid` and `ndgrid` functions respectively. The result of this code can be seen in Figure 29.4.

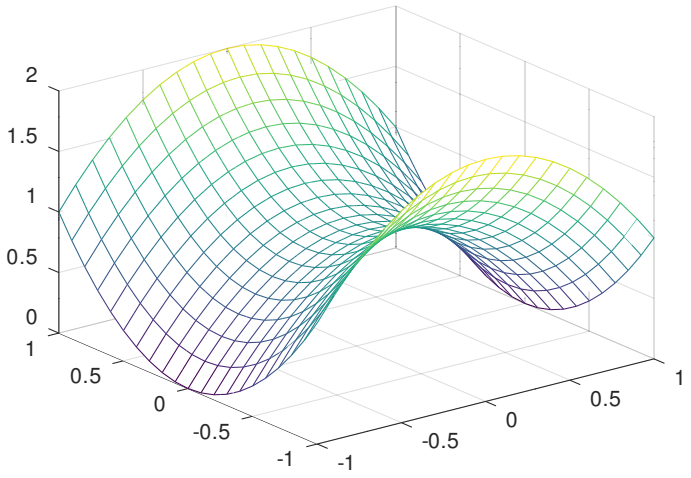



Figure 29.4: Demonstration of the use of `interpn`

# 30 Geometry

Much of the geometry code in Octave is based on the Qhull library[1]. Some of the documentation for Qhull, particularly for the options that can be passed to `delaunay`, `voronoi` and `convhull`, etc., is relevant to Octave users.

## 30.1 Delaunay Triangulation

The Delaunay triangulation is constructed from a set of circum-circles. These circum-circles are chosen so that there are at least three of the points in the set to triangulation on the circumference of the circum-circle. None of the points in the set of points falls within any of the circum-circles.

In general there are only three points on the circumference of any circum-circle. However, in some cases, and in particular for the case of a regular grid, 4 or more points can be on a single circum-circle. In this case the Delaunay triangulation is not unique.

*tri* = delaunay (*x*, *y*)
*tetr* = delaunay (*x*, *y*, *z*)
*tri* = delaunay (*x*)
*tri* = delaunay (*...*, *options*)

Compute the Delaunay triangulation for a 2-D or 3-D set of points.

For 2-D sets, the return value *tri* is a set of triangles which satisfies the Delaunay circum-circle criterion, i.e., no data point from [*x*, *y*] is within the circum-circle of the defining triangle. The set of triangles *tri* is a matrix of size [n, 3]. Each row defines a triangle and the three columns are the three vertices of the triangle. The value of `tri(i,j)` is an index into *x* and *y* for the location of the j-th vertex of the i-th triangle.

For 3-D sets, the return value *tetr* is a set of tetrahedrons which satisfies the Delaunay circum-circle criterion, i.e., no data point from [*x*, *y*, *z*] is within the circum-circle of the defining tetrahedron. The set of tetrahedrons is a matrix of size [n, 4]. Each row defines a tetrahedron and the four columns are the four vertices of the tetrahedron. The value of `tetr(i,j)` is an index into *x*, *y*, *z* for the location of the j-th vertex of the i-th tetrahedron.

The input *x* may also be a matrix with two or three columns where the first column contains x-data, the second y-data, and the optional third column contains z-data.

An optional final argument, which must be a string or cell array of strings, contains options passed to the underlying qhull command. See the documentation for the Qhull library for details http://www.qhull.org/html/qh-quick.htm#options. The default options are `{"Qt", "Qbb", "Qc"}`. If Qhull fails for 2-D input the triangulation is attempted again with the options `{"Qt", "Qbb", "Qc", "Qz"}` which may result in reduced accuracy.

If *options* is not present or `[]` then the default arguments are used. Otherwise, *options* replaces the default argument list. To append user options to the defaults it

[1] C.B. Barber, D.P. Dobkin, H.T. Huhdanpaa, "The Quickhull Algorithm for Convex Hulls", *ACM Trans. on Mathematical Software*, 22(4), pp. 469–483, Dec 1996, http://www.qhull.org

is necessary to repeat the default arguments in *options*. Use a null string to pass no arguments.

```
x = rand (1, 10);
y = rand (1, 10);
tri = delaunay (x, y);
triplot (tri, x, y);
hold on;
plot (x, y, "r*");
axis ([0,1,0,1]);
```

**See also:** [delaunayn], page 906, [convhull], page 916, [voronoi], page 912, [triplot], page 907, [trimesh], page 908, [tetramesh], page 909, [trisurf], page 908.

For 3-D inputs `delaunay` returns a set of tetrahedra that satisfy the Delaunay circum-circle criteria. Similarly, `delaunayn` returns the N-dimensional simplex satisfying the Delaunay circum-circle criteria. The N-dimensional extension of a triangulation is called a tessellation.

*T* = `delaunayn` (*pts*)
*T* = `delaunayn` (*pts*, *options*)

Compute the Delaunay triangulation for an N-dimensional set of points.

The Delaunay triangulation is a tessellation of the convex hull of a set of points such that no N-sphere defined by the N-triangles contains any other points from the set.

The input matrix *pts* of size [n, dim] contains n points in a space of dimension dim. The return matrix *T* has size [m, dim+1]. Each row of *T* contains a set of indices back into the original set of points *pts* which describes a simplex of dimension dim. For example, a 2-D simplex is a triangle and 3-D simplex is a tetrahedron.

An optional second argument, which must be a string or cell array of strings, contains options passed to the underlying qhull command. See the documentation for the Qhull library for details `http://www.qhull.org/html/qh-quick.htm#options`. The default options depend on the dimension of the input:

- 2-D and 3-D: *options* = `{"Qt", "Qbb", "Qc"}`
- 4-D and higher: *options* = `{"Qt", "Qbb", "Qc", "Qx"}`

If Qhull fails for 2-D input the triangulation is attempted again with the options `{"Qt", "Qbb", "Qc", "Qz"}` which may result in reduced accuracy.

If *options* is not present or `[]` then the default arguments are used. Otherwise, *options* replaces the default argument list. To append user options to the defaults it is necessary to repeat the default arguments in *options*. Use a null string to pass no arguments.

**See also:** [delaunay], page 905, [convhulln], page 917, [voronoin], page 913, [trimesh], page 908, [tetramesh], page 909.

An example of a Delaunay triangulation of a set of points is

```
rand ("state", 1);
x = rand (1, 10);
y = rand (1, 10);
T = delaunay (x, y);
X = [ x(T(:,1)); x(T(:,2)); x(T(:,3)); x(T(:,1)) ];
Y = [ y(T(:,1)); y(T(:,2)); y(T(:,3)); y(T(:,1)) ];
axis ([0, 1, 0, 1]);
plot (X, Y, "b", x, y, "r*");
```

The result of which can be seen in Figure 30.1.

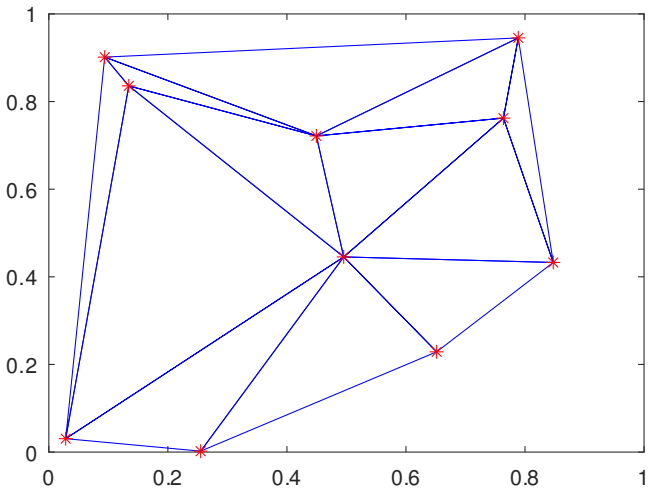



Figure 30.1: Delaunay triangulation of a random set of points

### 30.1.1 Plotting the Triangulation

Octave has the functions `triplot`, `trimesh`, and `trisurf` to plot the Delaunay triangulation of a 2-dimensional set of points. `tetramesh` will plot the triangulation of a 3-dimensional set of points.

`triplot (`*tri*, *x*, *y*`)`
`triplot (`*tri*, *x*, *y*, *linespec*`)`
*h* `= triplot (...)`

Plot a 2-D triangular mesh.

*tri* is typically the output of a Delaunay triangulation over the grid of *x*, *y*. Every row of *tri* represents one triangle and contains three indices into [*x*, *y*] which are the vertices of the triangles in the x-y plane.

The linestyle to use for the plot can be defined with the argument *linespec* of the same format as the `plot` command.

The optional return value *h* is a graphics handle to the created patch object.

**See also:** [plot], page 334, [trimesh], page 908, [trisurf], page 908, [delaunay], page 905.

`trimesh` (*tri*, *x*, *y*, *z*, *c*)
`trimesh` (*tri*, *x*, *y*, *z*)
`trimesh` (*tri*, *x*, *y*)
`trimesh` (..., *prop*, *val*, ...)
*h* = `trimesh` (...)

Plot a 3-D triangular wireframe mesh.

In contrast to `mesh`, which plots a mesh using rectangles, `trimesh` plots the mesh using triangles.

*tri* is typically the output of a Delaunay triangulation over the grid of *x*, *y*. Every row of *tri* represents one triangle and contains three indices into [*x*, *y*] which are the vertices of the triangles in the x-y plane. *z* determines the height above the plane of each vertex. If no *z* input is given then the triangles are plotted as a 2-D figure.

The color of the trimesh is computed by linearly scaling the *z* values to fit the range of the current colormap. Use `clim` and/or change the colormap to control the appearance.

Optionally, the color of the mesh can be specified independently of *z* by supplying *c*, which is a vector for colormap data, or a matrix with three columns for RGB data. The number of colors specified in *c* must either equal the number of vertices in *z* or the number of triangles in *tri*.

Any property/value pairs are passed directly to the underlying patch object. The full list of properties is documented at Section 15.3.3.8 [Patch Properties], page 495.

The optional return value *h* is a graphics handle to the created patch object.

**See also:** [mesh], page 383, [tetramesh], page 909, [triplot], page 907, [trisurf], page 908, [delaunay], page 905, [patch], page 451, [hidden], page 385.

`trisurf` (*tri*, *x*, *y*, *z*, *c*)
`trisurf` (*tri*, *x*, *y*, *z*)
`trisurf` (..., *prop*, *val*, ...)
*h* = `trisurf` (...)

Plot a 3-D triangular surface.

In contrast to `surf`, which plots a surface mesh using rectangles, `trisurf` plots the mesh using triangles.

*tri* is typically the output of a Delaunay triangulation over the grid of *x*, *y*. Every row of *tri* represents one triangle and contains three indices into [*x*, *y*] which are the vertices of the triangles in the x-y plane. *z* determines the height above the plane of each vertex.

The color of the trisurf is computed by linearly scaling the *z* values to fit the range of the current colormap. Use `clim` and/or change the colormap to control the appearance.

Optionally, the color of the mesh can be specified independently of *z* by supplying *c*, which is a vector for colormap data, or a matrix with three columns for RGB data. The number of colors specified in *c* must either equal the number of vertices in *z* or the number of triangles in *tri*. When specifying the color at each vertex the triangle will be colored according to the color of the first vertex only (see patch documentation and the `"FaceColor"` property when set to `"flat"`).

Any property/value pairs are passed directly to the underlying patch object. The full list of properties is documented at Section 15.3.3.8 [Patch Properties], page 495.

The optional return value *h* is a graphics handle to the created patch object.

**See also:** [surf], page 385, [triplot], page 907, [trimesh], page 908, [delaunay], page 905, [patch], page 451, [shading], page 407.

`tetramesh (`*T*`, `*X*`)`
`tetramesh (`*T*`, `*X*`, `*C*`)`
`tetramesh (...,` *property*`,` *val*`, ...)`
*h* `= tetramesh (...)`

Display the tetrahedrons defined in the m-by-4 matrix *T* as 3-D patches.

*T* is typically the output of a Delaunay triangulation of a 3-D set of points. Every row of *T* contains four indices into the n-by-3 matrix *X* of the vertices of a tetrahedron. Every row in *X* represents one point in 3-D space.

The vector *C* specifies the color of each tetrahedron as an index into the current colormap. The default value is 1:m where m is the number of tetrahedrons; the indices are scaled to map to the full range of the colormap. If there are more tetrahedrons than colors in the colormap then the values in *C* are cyclically repeated.

Calling `tetramesh (..., "property", "value", ...)` passes all property/value pairs directly to the patch function as additional arguments. The full list of properties is documented at Section 15.3.3.8 [Patch Properties], page 495.

The optional return value *h* is a vector of patch handles where each handle represents one tetrahedron in the order given by *T*. A typical use case for *h* is to turn the respective patch `"visible"` property `"on"` or `"off"`.

Type `demo tetramesh` to see examples on using `tetramesh`.

**See also:** [trimesh], page 908, [delaunay], page 905, [delaunayn], page 906, [patch], page 451.

The difference between `triplot`, and `trimesh` or `trisurf`, is that the former only plots the 2-dimensional triangulation itself, whereas the second two plot the value of a function `f (x, y)`. An example of the use of the `triplot` function is

```
rand ("state", 2)
x = rand (20, 1);
y = rand (20, 1);
tri = delaunay (x, y);
triplot (tri, x, y);
```

which plots the Delaunay triangulation of a set of random points in 2-dimensions. The output of the above can be seen in Figure 30.2.

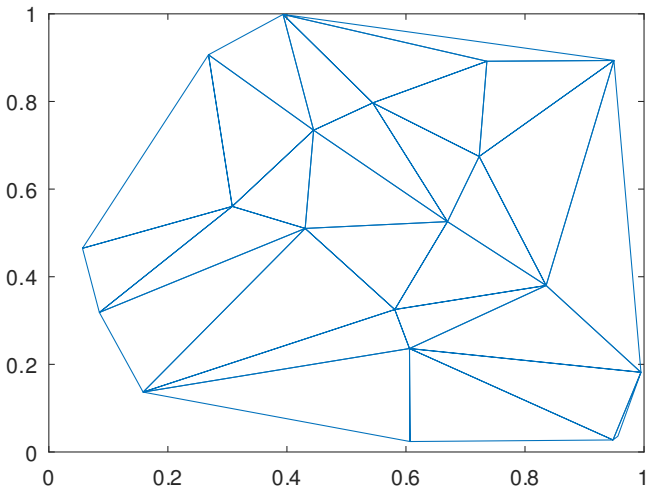



Figure 30.2: Delaunay triangulation of a random set of points

### 30.1.2 Identifying Points in Triangulation

It is often necessary to identify whether a particular point in the N-dimensional space is within the Delaunay tessellation of a set of points in this N-dimensional space, and if so which N-simplex contains the point and which point in the tessellation is closest to the desired point. The function `tsearch` performs this function in a triangulation, and the functions `tsearchn` and `dsearchn` in an N-dimensional tessellation.

To identify whether a particular point represented by a vector $p$ falls within one of the simplices of an N-simplex, we can write the Cartesian coordinates of the point in a parametric form with respect to the N-simplex. This parametric form is called the Barycentric Coordinates of the point. If the points defining the N-simplex are given by $N + 1$ vectors `t(i,:)`, then the Barycentric coordinates defining the point $p$ are given by

```
p = beta * t
```

where *beta* contains $N + 1$ values that together as a vector represent the Barycentric coordinates of the point $p$. To ensure a unique solution for the values of *beta* an additional criteria of

```
sum (beta) == 1
```

is imposed, and we can therefore write the above as

```
p - t(end, :) = beta(1:end-1) * (t(1:end-1, :)
                - ones (N, 1) * t(end, :)
```

Solving for *beta* we can then write

```
beta(1:end-1) = (p - t(end, :)) /
                (t(1:end-1, :) - ones (N, 1) * t(end, :))
beta(end) = sum (beta(1:end-1))
```

which gives the formula for the conversion of the Cartesian coordinates of the point *p* to the Barycentric coordinates *beta*. An important property of the Barycentric coordinates is that for all points in the N-simplex

```
0 <= beta(i) <= 1
```

Therefore, the test in `tsearch` and `tsearchn` essentially only needs to express each point in terms of the Barycentric coordinates of each of the simplices of the N-simplex and test the values of *beta*. This is exactly the implementation used in `tsearchn`. `tsearch` is optimized for 2-dimensions and the Barycentric coordinates are not explicitly formed.

*idx* = **tsearch** (*x*, *y*, *t*, *xi*, *yi*)

Search for the enclosing Delaunay convex hull.

For `t = delaunay (x, y)`, finds the index in *t* containing the points `(xi, yi)`. For points outside the convex hull, *idx* is NaN.

Programming Note: The algorithm is `O`(*M***N*) for locating *M* points in *N* triangles. Performance is typically much faster if the points to be located are in a single continuous path; a point is first checked against the region its predecessor was found in, speeding up lookups for points along a continuous path.

**See also:** [delaunay], page 905, [delaunayn], page 906.

*idx* = **tsearchn** (*x*, *t*, *xi*)
[*idx*, *p*] = **tsearchn** (*x*, *t*, *xi*)

Find the simplexes enclosing the given points.

`tsearchn` is typically used with `delaunayn`: `t = delaunayn (x)` returns a set of simplexes `t`, then `tsearchn` returns the row index of *t* containing each point of *xi*. For points outside the convex hull, *idx* is NaN.

If requested, `tsearchn` also returns the barycentric coordinates *p* of the enclosing simplexes.

**See also:** [tsearch], page 911, [dsearchn], page 911, [delaunayn], page 906.

An example of the use of `tsearch` can be seen with the simple triangulation

```
x = [-1; -1; 1; 1];
y = [-1; 1; -1; 1];
tri = [1, 2, 3; 2, 3, 4];
```

consisting of two triangles defined by *tri*. We can then identify which triangle a point falls in like

```
tsearch (x, y, tri, -0.5, -0.5)
⇒ 1
tsearch (x, y, tri, 0.5, 0.5)
⇒ 2
```

and we can confirm that a point doesn't lie within one of the triangles like

```
tsearch (x, y, tri, 2, 2)
⇒ NaN
```

The `dsearchn` function finds the closest point in a tessellation to the desired point. The desired point does not necessarily have to be in the tessellation, and even if it the returned point of the tessellation does not have to be one of the vertices of the N-simplex within which the desired point is found.

`idx = dsearchn (x, tri, xi)`
`idx = dsearchn (x, tri, xi, outval)`
`idx = dsearchn (x, xi)`
`[idx, d] = dsearchn (...)`

Return the index *idx* of the closest point in *x* to the elements *xi*.

If *outval* is supplied, then the values of *xi* that are not contained within one of the simplices *tri* are set to *outval*. Generally, *tri* is returned from `delaunayn (x)`.

The optional output *d* contains a column vector of distances between the query points *xi* and the nearest simplex points *x*.

Compatibility note: The `dsearchn` algorithm only uses the input *tri* when *outdim* is specified to determine if any points lie outside of the triangulation region. For compatibility, *tri* is accepted as an input even when *outdim* is not specified, but it is not used or checked to be a valid triangulation, and providing it will not affect either the output *idx* or the calculation efficiency.

**See also:** [tsearchn], page 911, [delaunayn], page 906.

An example of the use of `dsearchn`, using the above values of *x*, *y* and *tri* is

```
dsearchn ([x, y], tri, [-2, -2])
⇒ 1
```

If you wish the points that are outside the tessellation to be flagged, then `dsearchn` can be used as

```
dsearchn ([x, y], tri, [-2, -2], NaN)
⇒ NaN
dsearchn ([x, y], tri, [-0.5, -0.5], NaN)
⇒ 1
```

where the point outside the tessellation are then flagged with `NaN`.

## 30.2 Voronoi Diagrams

A Voronoi diagram or Voronoi tessellation of a set of points *s* in an N-dimensional space, is the tessellation of the N-dimensional space such that all points in `v(p)`, a partitions of the tessellation where *p* is a member of *s*, are closer to *p* than any other point in *s*. The Voronoi diagram is related to the Delaunay triangulation of a set of points, in that the vertices of the Voronoi tessellation are the centers of the circum-circles of the simplices of the Delaunay tessellation.

`voronoi (x, y)`
`voronoi (x, y, options)`
`voronoi (..., "linespec")`
`voronoi (hax, ...)`
`h = voronoi (...)`
`[vx, vy] = voronoi (...)`

Plot the Voronoi diagram of points `(x, y)`.

The Voronoi facets with points at infinity are not drawn.

The *options* argument, which must be a string or cell array of strings, contains options passed to the underlying qhull command. See the documentation for the Qhull library for details http://www.qhull.org/html/qh-quick.htm#options.

If `"linespec"` is given it is used to set the color and line style of the plot.

If an axes graphics handle *hax* is supplied then the Voronoi diagram is drawn on the specified axes rather than in a new figure.

If a single output argument is requested then the Voronoi diagram will be plotted and a graphics handle *h* to the plot is returned.

[*vx*, *vy*] = voronoi (. . .) returns the Voronoi vertices instead of plotting the diagram.

```
x = rand (10, 1);
y = rand (size (x));
h = convhull (x, y);
[vx, vy] = voronoi (x, y);
plot (vx, vy, "-b", x, y, "o", x(h), y(h), "-g");
legend ("", "points", "hull");
```

**See also:** [voronoin], page 913, [delaunay], page 905, [convhull], page 916.

`[`*C*`, `*F*`] = voronoin (`*pts*`)`
`[`*C*`, `*F*`] = voronoin (`*pts*`, `*options*`)`

Compute N-dimensional Voronoi facets.

The input matrix *pts* of size [n, dim] contains n points in a space of dimension dim.

*C* contains the points of the Voronoi facets. The list *F* contains, for each facet, the indices of the Voronoi points.

An optional second argument, which must be a string or cell array of strings, contains options passed to the underlying qhull command. See the documentation for the Qhull library for details http://www.qhull.org/html/qh-quick.htm#options.

The default options depend on the dimension of the input:

- 2-D and 3-D: *options* = `{"Qbb"}`
- 4-D and higher: *options* = `{"Qbb", "Qx"}`

If *options* is not present or `[]` then the default arguments are used. Otherwise, *options* replaces the default argument list. To append user options to the defaults it is necessary to repeat the default arguments in *options*. Use a null string to pass no arguments.

**See also:** [voronoi], page 912, [convhulln], page 917, [delaunayn], page 906.

An example of the use of `voronoi` is

```
rand ("state",9);
x = rand (10,1);
y = rand (10,1);
tri = delaunay (x, y);
[vx, vy] = voronoi (x, y, tri);
triplot (tri, x, y, "b");
hold on;
plot (vx, vy, "r");
```

The result of which can be seen in Figure 30.3. Note that the circum-circle of one of the triangles has been added to this figure, to make the relationship between the Delaunay tessellation and the Voronoi diagram clearer.

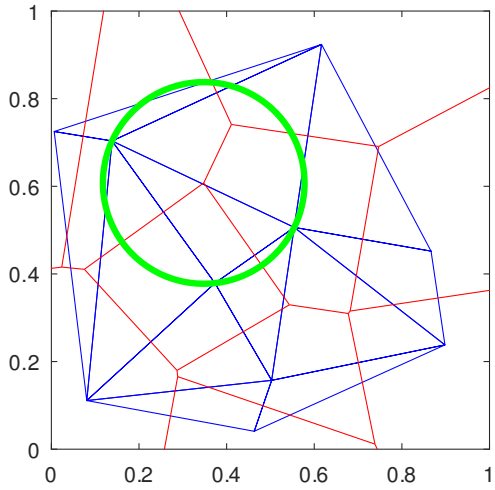



Figure 30.3: Delaunay triangulation (blue lines) and Voronoi diagram (red lines) of a random set of points

Additional information about the size of the facets of a Voronoi diagram, and which points of a set of points is in a polygon can be had with the `polyarea` and `inpolygon` functions respectively.

`a = polyarea (`*x*, *y*`)`
`a = polyarea (`*x*, *y*, *dim*`)`

Determine area of a polygon by triangle method.

The variables *x* and *y* define the vertex pairs, and must therefore have the same shape. They can be either vectors or arrays. If they are arrays then the columns of *x* and *y* are treated separately and an area returned for each.

If the optional *dim* argument is given, then `polyarea` works along this dimension of the arrays *x* and *y*.

An example of the use of `polyarea` might be

```
rand ("state", 2);
x = rand (10, 1);
y = rand (10, 1);
[c, f] = voronoin ([x, y]);
af = zeros (size (f));
for i = 1 : length (f)
  af(i) = polyarea (c (f {i, :}, 1), c (f {i, :}, 2));
endfor
```

Facets of the Voronoi diagram with a vertex at infinity have infinity area. A simplified version of `polyarea` for rectangles is available with `rectint`

*area* = `rectint` (*a*, *b*)
: Compute area or volume of intersection of rectangles or N-D boxes.

  Compute the area of intersection of rectangles in *a* and rectangles in *b*. N-dimensional boxes are supported in which case the volume, or hypervolume is computed according to the number of dimensions.

  2-dimensional rectangles are defined as `[xpos ypos width height]` where xpos and ypos are the position of the bottom left corner. Higher dimensions are supported where the coordinates for the minimum value of each dimension follow the length of the box in that dimension, e.g., `[xpos ypos zpos kpos ... width height depth k_length ...]`.

  Each row of *a* and *b* define a rectangle, and if both define multiple rectangles, then the output, *area*, is a matrix where the i-th row corresponds to the i-th row of a and the j-th column corresponds to the j-th row of b.

  **See also:** [polyarea], page 914.

*in* = `inpolygon` (*x*, *y*, *xv*, *yv*)
[*in*, *on*] = `inpolygon` (*x*, *y*, *xv*, *yv*)
: For a polygon defined by vertex points `(`*xv*`, `*yv*`)`, return true if the points `(`*x*`, `*y*`)` are inside (or on the boundary) of the polygon; Otherwise, return false.

  The input variables *x* and *y*, must have the same dimension.

  The optional output *on* returns true if the points are exactly on the polygon edge, and false otherwise.

  **See also:** [delaunay], page 905.

An example of the use of `inpolygon` might be

```
randn ("state", 2);
x = randn (100, 1);
y = randn (100, 1);
vx = cos (pi * [-1 : 0.1: 1]);
vy = sin (pi * [-1 : 0.1 : 1]);
in = inpolygon (x, y, vx, vy);
plot (vx, vy, x(in), y(in), "r+", x(!in), y(!in), "bo");
axis ([-2, 2, -2, 2]);
```

The result of which can be seen in Figure 30.4.

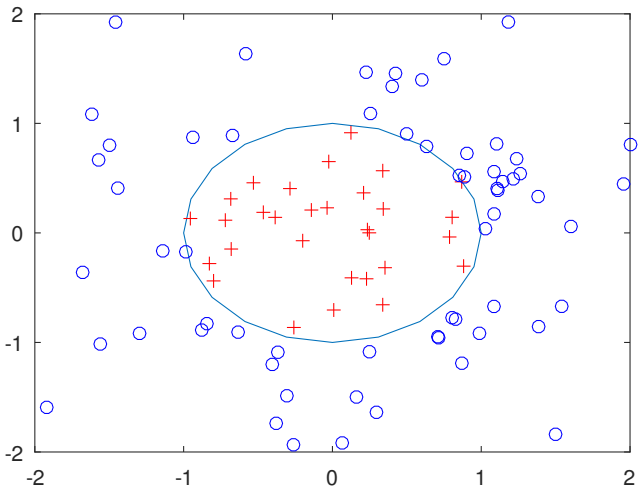



Figure 30.4: Demonstration of the `inpolygon` function to determine the points inside a polygon

## 30.3 Convex Hull

The convex hull of a set of points is the minimum convex envelope containing all of the points. Octave has the functions `convhull` and `convhulln` to calculate the convex hull of 2-dimensional and N-dimensional sets of points.

`H = convhull (x, y)`
`H = convhull (x, y, z)`
`H = convhull (x)`
`H = convhull (..., options)`
`[H, V] = convhull (...)`

Compute the convex hull of a 2-D or 3-D set of points.

The hull $H$ is a linear index vector into the original set of points that specifies which points form the enclosing hull. For 2-D inputs only, the output is ordered in a counterclockwise manner around the hull.

The input $x$ may also be a matrix with two or three columns where the first column contains x-data, the second y-data, and the optional third column contains z-data.

An optional final argument, which must be a string or cell array of strings, contains options passed to the underlying qhull command. See the documentation for the Qhull library for details http://www.qhull.org/html/qh-quick.htm#options. The default option is `{"Qt"}`.

If *options* is not present or `[]` then the default arguments are used. Otherwise, *options* replaces the default argument list. To append user options to the defaults it is necessary to repeat the default arguments in *options*. Use a null string to pass no arguments.

If the second output *V* is requested the volume of the enclosing convex hull is calculated.

**See also:** [convhulln], page 917, [delaunay], page 905, [voronoi], page 912.

`h = convhulln (pts)`
`h = convhulln (pts, options)`
`[h, v] = convhulln (...)`
Compute the convex hull of the set of points *pts*.

*pts* is a matrix of size [n, dim] containing n points in a space of dimension dim.

The hull *h* is an index vector into the set of points and specifies which points form the enclosing hull.

An optional second argument, which must be a string or cell array of strings, contains options passed to the underlying qhull command. See the documentation for the Qhull library for details http://www.qhull.org/html/qh-quick.htm#options. The default options depend on the dimension of the input:

- 2-D, 3-D, 4-D: *options* = `{"Qt"}`
- 5-D and higher: *options* = `{"Qt", "Qx"}`

If *options* is not present or `[]` then the default arguments are used. Otherwise, *options* replaces the default argument list. To append user options to the defaults it is necessary to repeat the default arguments in *options*. Use a null string to pass no arguments.

If the second output *v* is requested the volume of the enclosing convex hull is calculated.

**See also:** [convhull], page 916, [delaunayn], page 906, [voronoin], page 913.

An example of the use of `convhull` is

```
x = -3:0.05:3;
y = abs (sin (x));
k = convhull (x, y);
plot (x(k), y(k), "r-", x, y, "b+");
axis ([-3.05, 3.05, -0.05, 1.05]);
```

The output of the above can be seen in Figure 30.5.

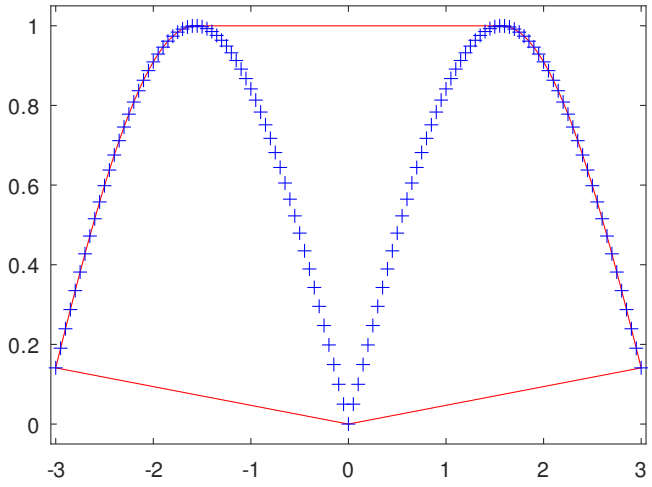



Figure 30.5: The convex hull of a simple set of points

## 30.4 Interpolation on Scattered Data

An important use of the Delaunay tessellation is that it can be used to interpolate from scattered data to an arbitrary set of points. To do this the N-simplex of the known set of points is calculated with `delaunay` or `delaunayn`. Then the simplices in to which the desired points are found are identified. Finally the vertices of the simplices are used to interpolate to the desired points. The functions that perform this interpolation are `griddata`, `griddata3` and `griddatan`.

*zi* = griddata (*x*, *y*, *z*, *xi*, *yi*)
*zi* = griddata (*x*, *y*, *z*, *xi*, *yi*, *method*)
[*xi*, *yi*, *zi*] = griddata (...)
*vi* = griddata (*x*, *y*, *z*, *v*, *xi*, *yi*, *zi*)
*vi* = griddata (*x*, *y*, *z*, *v*, *xi*, *yi*, *zi*, *method*)
*vi* = griddata (*x*, *y*, *z*, *v*, *xi*, *yi*, *zi*, *method*, *options*)

Interpolate irregular 2-D and 3-D source data at specified points.

For 2-D interpolation, the inputs *x* and *y* define the points where the function *z* = f (*x*, *y*) is evaluated. The inputs *x*, *y*, *z* are either vectors of the same length, or the unequal vectors *x*, *y* are expanded to a 2-D grid with `meshgrid` and *z* is a 2-D matrix matching the resulting size of the X-Y grid.

The interpolation points are (*xi*, *yi*). If either *xi* or *y* is a row vector and the other is a column vector, then `meshgrid (xi, yi)` will be used to create a mesh of interpolation points.

For 3-D interpolation, the inputs *x*, *y*, and *z* define the points where the function *v* = f (*x*, *y*, *z*) is evaluated. The inputs *x*, *y*, *z* are either vectors of the same length, or if they are of unequal length, then they are expanded to a 3-D grid with `meshgrid`.

The size of the input *v* must match the size of the original data, either as a vector or a matrix.

The outputs *zi* (for 2-D) or *vi* (for 3-D) will contain the interpolated values of *z* or *v*, respectively, with the output size matching that of the interpolation points.

The optional input interpolation *method* can be `"nearest"`, `"linear"`, or for 2-D data `"v4"`. When the method is `"nearest"`, the output *vi* will be the closest point in the original data (*x*, *y*, *z*) to the query point (*xi*, *yi*, *zi*). When the method is `"linear"`, the output *vi* will be a linear interpolation between the two closest points in the original source data in each dimension. For 2-D cases only, the `"v4"` method is also available which implements a biharmonic spline interpolation. If *method* is omitted or empty, it defaults to `"linear"`.

For 3-D interpolation, the optional argument *options* is passed directly to Qhull when computing the Delaunay triangulation used for interpolation. For more information on the defaults and how to pass different values, see [delaunayn], page 906.

Programming Notes: If the input is complex the real and imaginary parts are interpolated separately. Interpolation is normally based on a Delaunay triangulation. Any query values outside the convex hull of the input points will return `NaN`. However, the `"v4"` method does not use the triangulation and will return values outside the original data (extrapolation).

**See also:** [griddata3], page 919, [griddatan], page 919, [delaunay], page 905.

`vi = griddata3 (x, y, z, v, xi, yi, zi)`
`vi = griddata3 (x, y, z, v, xi, yi, zi, method)`
`vi = griddata3 (x, y, z, v, xi, yi, zi, method, options)`

Interpolate irregular 3-D source data at specified points.

The inputs *x*, *y*, and *z* define the points where the function `v = f (x, y, z)` is evaluated. The inputs *x*, *y*, *z* are either vectors of the same length, or if they are of unequal length, then they are expanded to a 3-D grid with `meshgrid`. The size of the input *v* must match the size of the original data, either as a vector or a matrix.

The interpolation points are specified by *xi*, *yi*, *zi*.

The optional input interpolation *method* can be `"nearest"` or `"linear"`. When the method is `"nearest"`, the output *vi* will be the closest point in the original data (*x*, *y*, *z*) to the query point (*xi*, *yi*, *zi*). When the method is `"linear"`, the output *vi* will be a linear interpolation between the two closest points in the original source data in each dimension. If *method* is omitted or empty, it defaults to `"linear"`.

The optional argument *options* is passed directly to Qhull when computing the Delaunay triangulation used for interpolation. See `delaunayn` for information on the defaults and how to pass different values.

Programming Notes: If the input is complex the real and imaginary parts are interpolated separately. Interpolation is based on a Delaunay triangulation and any query values outside the convex hull of the input points will return `NaN`.

**See also:** [griddata], page 918, [griddatan], page 919, [delaunayn], page 906.

`yi = griddatan (x, y, xi)`
`yi = griddatan (x, y, xi, method)`

*yi* = griddatan (*x*, *y*, *xi*, *method*, *options*)

Interpolate irregular source data *x*, *y* at points specified by *xi*.

The input *x* is an MxN matrix representing M points in an N-dimensional space. The input *y* is a single-valued column vector (Mx1) representing a function evaluated at the points *x*, i.e., *y* = `fcn (x)`. The input *xi* is a list of points for which the function output *yi* should be approximated through interpolation. *xi* must have the same number of columns (*N*) as *x* so that the dimensionality matches.

The optional input interpolation *method* can be `"nearest"` or `"linear"`. When the method is `"nearest"`, the output *yi* will be the closest point in the original data *x* to the query point *xi*. When the method is `"linear"`, the output *yi* will be a linear interpolation between the two closest points in the original source data. If *method* is omitted or empty, it defaults to `"linear"`.

The optional argument *options* is passed directly to Qhull when computing the Delaunay triangulation used for interpolation. See `delaunayn` for information on the defaults and how to pass different values.

Example

```
## Evaluate sombrero() function at irregular data points
x = 16*gallery ("uniformdata", [200,1], 1) - 8;
y = 16*gallery ("uniformdata", [200,1], 11) - 8;
z = sin (sqrt (x.^2 + y.^2)) ./ sqrt (x.^2 + y.^2);
## Create a regular grid and interpolate data
[xi, yi] = ndgrid (linspace (-8, 8, 50));
zi = griddatan ([x, y], z, [xi(:), yi(:)]);
zi = reshape (zi, size (xi));
## Plot results
clf ();
plot3 (x, y, z, "or");
hold on
surf (xi, yi, zi);
legend ("Original Data", "Interpolated Data");
```

Programming Notes: If the input is complex the real and imaginary parts are interpolated separately. Interpolation is based on a Delaunay triangulation and any query values outside the convex hull of the input points will return `NaN`. For 2-D and 3-D data additional interpolation methods are available by using the `griddata` function.

**See also:** [griddata], page 918, [griddata3], page 919, [delaunayn], page 906.

An example of the use of the `griddata` function is

```
rand ("state", 1);
x = 2*rand (1000,1) - 1;
y = 2*rand (size (x)) - 1;
z = sin (2*(x.^2+y.^2));
[xx,yy] = meshgrid (linspace (-1,1,32));
zz = griddata (x, y, z, xx, yy);
mesh (xx, yy, zz);
```

that interpolates from a random scattering of points, to a uniform grid. The output of the above can be seen in Figure 30.6.

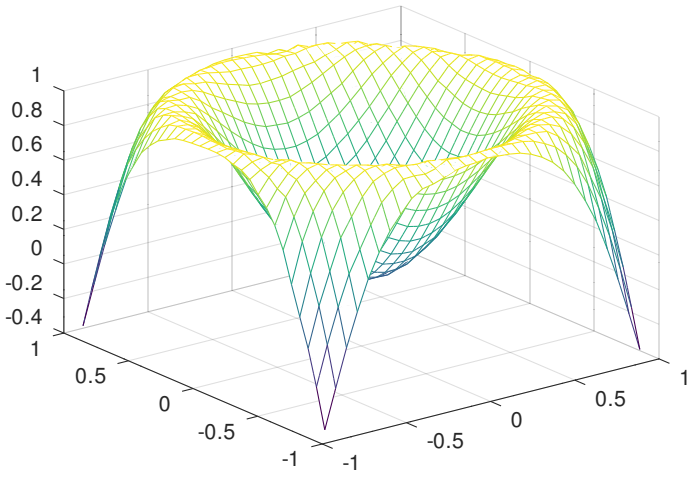



Figure 30.6: Interpolation from a scattered data to a regular grid

## 30.5 Vector Rotation Matrices

Also included in Octave's geometry functions are primitive functions to enable vector rotations in 3-dimensional space. Separate functions are provided for rotation about each of the principle axes, $x$, $y$, and $z$. According to Euler's rotation theorem, any arbitrary rotation, $R$, of any vector, $p$, can be expressed as a product of the three principle rotations:

$$p' = R \cdot p = R_z \cdot R_y \cdot R_x \cdot p$$

`T = rotx (angle)`

`rotx` returns the 3x3 transformation matrix corresponding to an active rotation of a vector about the x-axis by the specified *angle*, given in degrees, where a positive angle corresponds to a counterclockwise rotation when viewing the y-z plane from the positive x side.

The form of the transformation matrix is:

$$T = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(angle) & -\sin(angle) \\ 0 & \sin(angle) & \cos(angle) \end{bmatrix}.$$

This rotation matrix is intended to be used as a left-multiplying matrix when acting on a column vector, using the notation `v = T*u`. For example, a vector, $u$, pointing along the positive y-axis, rotated 90-degrees about the x-axis, will result in a vector pointing along the positive z-axis:

```
>> u = [0 1 0]'
u =
   0
   1
   0

>> T = rotx (90)
T =
   1.00000   0.00000   0.00000
   0.00000   0.00000  -1.00000
   0.00000   1.00000   0.00000

>> v = T*u
v =
   0.00000
   0.00000
   1.00000
```

**See also:** [roty], page 922, [rotz], page 923.

`T = roty (angle)`

`roty` returns the 3x3 transformation matrix corresponding to an active rotation of a vector about the y-axis by the specified *angle*, given in degrees, where a positive angle corresponds to a counterclockwise rotation when viewing the z-x plane from the positive y side.

The form of the transformation matrix is:

$$T = \begin{bmatrix} \cos(angle) & 0 & \sin(angle) \\ 0 & 1 & 0 \\ -\sin(angle) & 0 & \cos(angle) \end{bmatrix}.$$

This rotation matrix is intended to be used as a left-multiplying matrix when acting on a column vector, using the notation `v = T*u`. For example, a vector, $u$, pointing along the positive z-axis, rotated 90-degrees about the y-axis, will result in a vector pointing along the positive x-axis:

```
>> u = [0 0 1]'
 u =
    0
    0
    1

 >> T = roty (90)
 T =
    0.00000   0.00000   1.00000
    0.00000   1.00000   0.00000
   -1.00000   0.00000   0.00000

 >> v = T*u
 v =
    1.00000
    0.00000
    0.00000
```

**See also:** [rotx], page 921, [rotz], page 923.

*T* = rotz (*angle*)

`rotz` returns the 3x3 transformation matrix corresponding to an active rotation of a vector about the z-axis by the specified *angle*, given in degrees, where a positive angle corresponds to a counterclockwise rotation when viewing the x-y plane from the positive z side.

The form of the transformation matrix is:

$$T = \begin{bmatrix} \cos(angle) & -\sin(angle) & 0 \\ \sin(angle) & \cos(angle) & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

This rotation matrix is intended to be used as a left-multiplying matrix when acting on a column vector, using the notation `v = T*u`. For example, a vector, $u$, pointing along the positive x-axis, rotated 90-degrees about the z-axis, will result in a vector pointing along the positive y-axis:

```
>> u = [1 0 0]'
 u =
    1
    0
    0

 >> T = rotz (90)
 T =
    0.00000  -1.00000   0.00000
    1.00000   0.00000   0.00000
    0.00000   0.00000   1.00000

 >> v = T*u
 v =
    0.00000
    1.00000
    0.00000
```

**See also:** [rotx], page 921, [roty], page 922.

# 31 Signal Processing

This chapter describes the signal processing and fast Fourier transform functions available in Octave. Fast Fourier transforms are computed with the FFTW or FFTPACK libraries depending on how Octave is built.

`y = fft (x)`
`y = fft (x, n)`
`y = fft (x, n, dim)`

Compute the discrete Fourier transform of $x$ using a Fast Fourier Transform (FFT) algorithm.

The FFT is calculated along the first non-singleton dimension of the array. Thus if $x$ is a matrix, `fft (x)` computes the FFT for each column of $x$.

If called with two arguments, $n$ is expected to be an integer specifying the number of elements of $x$ to use, or an empty matrix to specify that its value should be ignored. If $n$ is larger than the dimension along which the FFT is calculated, then $x$ is resized and padded with zeros. Otherwise, if $n$ is smaller than the dimension along which the FFT is calculated, then $x$ is truncated.

If called with three arguments, *dim* is an integer specifying the dimension of the matrix along which the FFT is performed.

**See also:** [ifft], page 925, [fft2], page 925, [fftn], page 926, [fftw], page 926.

`x = ifft (y)`
`x = ifft (y, n)`
`x = ifft (y, n, dim)`

Compute the inverse discrete Fourier transform of $y$ using a Fast Fourier Transform (FFT) algorithm.

The inverse FFT is calculated along the first non-singleton dimension of the array. Thus if $y$ is a matrix, `ifft (y)` computes the inverse FFT for each column of $y$.

If called with two arguments, $n$ is expected to be an integer specifying the number of elements of $y$ to use, or an empty matrix to specify that its value should be ignored. If $n$ is larger than the dimension along which the inverse FFT is calculated, then $y$ is resized and padded with zeros. Otherwise, if $n$ is smaller than the dimension along which the inverse FFT is calculated, then $y$ is truncated.

If called with three arguments, *dim* is an integer specifying the dimension of the matrix along which the inverse FFT is performed.

**See also:** [fft], page 925, [ifft2], page 926, [ifftn], page 926, [fftw], page 926.

`B = fft2 (A)`
`B = fft2 (A, m, n)`

Compute the two-dimensional discrete Fourier transform of $A$ using a Fast Fourier Transform (FFT) algorithm.

The optional arguments $m$ and $n$ may be used specify the number of rows and columns of $A$ to use. If either of these is larger than the size of $A$, $A$ is resized and padded with zeros.

If $A$ is a multi-dimensional matrix, each two-dimensional sub-matrix of $A$ is treated separately.

**See also:** [ifft2], page 926, [fft], page 925, [fftn], page 926, [fftw], page 926.

`A = ifft2 (B)`
`A = ifft2 (B, m, n)`

Compute the inverse two-dimensional discrete Fourier transform of $B$ using a Fast Fourier Transform (FFT) algorithm.

The optional arguments $m$ and $n$ may be used specify the number of rows and columns of $B$ to use. If either of these is larger than the size of $B$, $B$ is resized and padded with zeros.

If $B$ is a multi-dimensional matrix, each two-dimensional sub-matrix of $B$ is treated separately.

**See also:** [fft2], page 925, [ifft], page 925, [ifftn], page 926, [fftw], page 926.

`B = fftn (A)`
`B = fftn (A, size)`

Compute the N-dimensional discrete Fourier transform of $A$ using a Fast Fourier Transform (FFT) algorithm.

The optional vector argument *size* may be used specify the dimensions of the array to be used. If an element of *size* is smaller than the corresponding dimension of $A$, then the dimension of $A$ is truncated prior to performing the FFT. Otherwise, if an element of *size* is larger than the corresponding dimension then $A$ is resized and padded with zeros.

**See also:** [ifftn], page 926, [fft], page 925, [fft2], page 925, [fftw], page 926.

`A = ifftn (B)`
`A = ifftn (B, size)`

Compute the inverse N-dimensional discrete Fourier transform of $B$ using a Fast Fourier Transform (FFT) algorithm.

The optional vector argument *size* may be used specify the dimensions of the array to be used. If an element of *size* is smaller than the corresponding dimension of $B$, then the dimension of $B$ is truncated prior to performing the inverse FFT. Otherwise, if an element of *size* is larger than the corresponding dimension then $B$ is resized and padded with zeros.

**See also:** [fftn], page 926, [ifft], page 925, [ifft2], page 926, [fftw], page 926.

Octave uses the FFTW libraries to perform FFT computations. When Octave starts up and initializes the FFTW libraries, they read a system wide file (on a Unix system, it is typically `/etc/fftw/wisdom`) that contains information useful to speed up FFT computations. This information is called the *wisdom*. The system-wide file allows wisdom to be shared between all applications using the FFTW libraries.

Use the `fftw` function to generate and save wisdom. Using the utilities provided together with the FFTW libraries (`fftw-wisdom` on Unix systems), you can even add wisdom generated by Octave to the system-wide wisdom file.

*method* = fftw ("*planner*")
fftw ("*planner*", *method*)
*wisdom* = fftw ("*dwisdom*")
fftw ("*dwisdom*", *wisdom*)
*nthreads* = fftw ("*threads*")
fftw ("*threads*", *nthreads*)

Manage FFTW wisdom data.

Wisdom data can be used to significantly accelerate the calculation of the FFTs, but implies an initial cost in its calculation. When the FFTW libraries are initialized, they read a system wide wisdom file (typically in `/etc/fftw/wisdom`), allowing wisdom to be shared between applications other than Octave. Alternatively, the `fftw` function can be used to import wisdom. For example,

```
wisdom = fftw ("dwisdom")
```

will save the existing wisdom used by Octave to the string *wisdom*. This string can then be saved to a file and restored using the `save` and `load` commands respectively. This existing wisdom can be re-imported as follows

```
fftw ("dwisdom", wisdom)
```

If *wisdom* is an empty string, then the wisdom used is cleared.

During the calculation of Fourier transforms further wisdom is generated. The fashion in which this wisdom is generated is also controlled by the `fftw` function. There are five different manners in which the wisdom can be treated:

`"estimate"`
: Specifies that no run-time measurement of the optimal means of calculating a particular is performed, and a simple heuristic is used to pick a (probably sub-optimal) plan. The advantage of this method is that there is little or no overhead in the generation of the plan, which is appropriate for a Fourier transform that will be calculated once.

`"measure"`
: In this case a range of algorithms to perform the transform is considered and the best is selected based on their execution time.

`"patient"`
: Similar to `"measure"`, but a wider range of algorithms is considered.

`"exhaustive"`
: Like `"measure"`, but all possible algorithms that may be used to treat the transform are considered.

`"hybrid"`
: As run-time measurement of the algorithm can be expensive, this is a compromise where `"measure"` is used for transforms up to the size of 8192 and beyond that the `"estimate"` method is used.

The default method is `"estimate"`. The current method can be queried with

```
method = fftw ("planner")
```

or set by using

```
fftw ("planner", method)
```

Note that calculated wisdom will be lost when restarting Octave. However, the wisdom data can be reloaded if it is saved to a file as described above. Saved wisdom files should not be used on different platforms since they will not be efficient and the point of calculating the wisdom is lost.

The number of threads used for computing the plans and executing the transforms can be set with

```
fftw ("threads", NTHREADS)
```

Note that Octave must be compiled with multi-threaded FFTW support for this feature. By default, the number of (logical) processors available to the current process or *3* is used (whichever is smaller).

**See also:** [fft], page 925, [ifft], page 925, [fft2], page 925, [ifft2], page 926, [fftn], page 926, [ifftn], page 926.

```
c = fftconv (x, y)
c = fftconv (x, y, n)
```

Convolve two vectors using the FFT for computation.

`c = fftconv (x, y)` returns a vector of length equal to `length (x) + length (y) - 1`. If $x$ and $y$ are the coefficient vectors of two polynomials, the returned value is the coefficient vector of the product polynomial.

The computation uses the FFT by calling the function `fftfilt`. If the optional argument $n$ is specified, an N-point FFT is used.

**See also:** [deconv], page 880, [conv], page 879, [conv2], page 880.

```
y = fftfilt (b, x)
y = fftfilt (b, x, n)
```

Filter $x$ with the FIR filter $b$ using the FFT.

If $x$ is a matrix, filter each column of the matrix.

Given the optional third argument, $n$, `fftfilt` uses the overlap-add method to filter $x$ with $b$ using an N-point FFT. The FFT size must be an even power of 2 and must be greater than or equal to the length of $b$. If the specified $n$ does not meet these criteria, it is automatically adjusted to the nearest value that does.

**See also:** [filter], page 928, [filter2], page 929.

```
y = filter (b, a, x)
[y, sf] = filter (b, a, x, si)
[y, sf] = filter (b, a, x, [], dim)
[y, sf] = filter (b, a, x, si, dim)
```

Apply a 1-D digital filter to the data $x$.

`filter` returns the solution to the following linear, time-invariant difference equation:

$$\sum_{k=0}^{N} a_{k+1} y_{n-k} = \sum_{k=0}^{M} b_{k+1} x_{n-k}, \qquad 1 \le n \le P$$

where $a \in \Re^{N-1}$, $b \in \Re^{M-1}$, and $x \in \Re^{P}$. The result is calculated over the first non-singleton dimension of $x$ or over *dim* if supplied.

An equivalent form of the equation is:

$$y_n = -\sum_{k=1}^{N} c_{k+1} y_{n-k} + \sum_{k=0}^{M} d_{k+1} x_{n-k}, \qquad 1 \le n \le P$$

where $c = a/a_1$ and $d = b/a_1$.

If the fourth argument *si* is provided, it is taken as the initial state of the system and the final state is returned as *sf*. The state vector is a column vector whose length is equal to the length of the longest coefficient vector minus one. If *si* is not supplied, the initial state vector is set to all zeros.

In terms of the Z Transform, $y$ is the result of passing the discrete-time signal $x$ through a system characterized by the following rational system function:

$$H(z) = \frac{\sum_{k=0}^{M} d_{k+1} z^{-k}}{1 + \sum_{k+1}^{N} c_{k+1} z^{-k}}$$

**See also:** [filter2], page 929, [fftfilt], page 928, [freqz], page 929.

`y = filter2 (b, x)`
`y = filter2 (b, x, shape)`

Apply the 2-D FIR filter $b$ to $x$.

If the argument *shape* is specified, return an array of the desired shape. Possible values are:

`"full"` pad $x$ with zeros on all sides before filtering.

`"same"` unpadded $x$ (default)

`"valid"` trim $x$ after filtering so edge effects are no included.

Note this is just a variation on convolution, with the parameters reversed and $b$ rotated 180 degrees.

**See also:** [conv2], page 880.

`[h, w] = freqz (b, a, n, "whole")`
`[h, w] = freqz (b)`
`[h, w] = freqz (b, a)`
`[h, w] = freqz (b, a, n)`
`h = freqz (b, a, w)`
`[h, w] = freqz (..., Fs)`
`freqz (...)`

Return the complex frequency response $h$ of the rational IIR filter whose numerator and denominator coefficients are $b$ and $a$, respectively.

The response is evaluated at $n$ angular frequencies between 0 and $2\pi$.

The output value $w$ is a vector of the frequencies.

If *a* is omitted, the denominator is assumed to be 1 (this corresponds to a simple FIR filter).

If *n* is omitted, a value of 512 is assumed. For fastest computation, *n* should factor into a small number of small primes.

If the fourth argument, `"whole"`, is omitted the response is evaluated at frequencies between 0 and $\pi$.

`freqz (b, a, w)`

Evaluate the response at the specific frequencies in the vector *w*. The values for *w* are measured in radians.

`[...] = freqz (..., Fs)`

Return frequencies in Hz instead of radians assuming a sampling rate *Fs*. If you are evaluating the response at specific frequencies *w*, those frequencies should be requested in Hz rather than radians.

`freqz (...)`

Plot the magnitude and phase response of *h* rather than returning them.

**See also:** [freqz_plot], page 930.

`freqz_plot (w, h)`
`freqz_plot (w, h, freq_norm)`

Plot the magnitude and phase response of *h*.

If the optional *freq_norm* argument is true, the frequency vector *w* is in units of normalized radians. If *freq_norm* is false, or not given, then *w* is measured in Hertz.

**See also:** [freqz], page 929.

`y = sinc (x)`

Compute the sinc function.

Return $\sin(\pi x)/(\pi x)$.

`b = unwrap (x)`
`b = unwrap (x, tol)`
`b = unwrap (x, tol, dim)`

Unwrap radian phases by adding or subtracting multiples of 2*pi as appropriate to remove jumps greater than *tol*.

*tol* defaults to pi.

`unwrap` will work along the dimension *dim*. If *dim* is unspecified it defaults to the first non-singleton dimension.

`unwrap` ignores all non-finite input values (Inf, NaN, NA).

`[a, b] = arch_fit (y, x, p, iter, gamma, a0, b0)`

Fit an ARCH regression model to the time series *y* using the scoring algorithm in Engle's original ARCH paper.

The model is

```
y(t) = b(1) * x(t,1) + ... + b(k) * x(t,k) + e(t),
h(t) = a(1) + a(2) * e(t-1)^2 + ... + a(p+1) * e(t-p)^2
```

in which $e(t)$ is $N(0, h(t))$, given a time-series vector $y$ up to time $t-1$ and a matrix of (ordinary) regressors $x$ up to $t$. The order of the regression of the residual variance is specified by $p$.

If invoked as `arch_fit (y, k, p)` with a positive integer $k$, fit an ARCH($k$, $p$) process, i.e., do the above with the $t$-th row of $x$ given by

```
[1, y(t-1), ..., y(t-k)]
```

Optionally, one can specify the number of iterations *iter*, the updating factor *gamma*, and initial values $a0$ and $b0$ for the scoring algorithm.

`y = arch_rnd (a, b, t)`

Simulate an ARCH sequence of length $t$ with AR coefficients $b$ and CH coefficients $a$.

The result $y(t)$ follows the model

```
y(t) = b(1) + b(2) * y(t-1) + ... + b(lb) * y(t-lb+1) + e(t),
```

where $e(t)$, given $y$ up to time $t-1$, is $N(0, h(t))$, with

```
h(t) = a(1) + a(2) * e(t-1)^2 + ... + a(la) * e(t-la+1)^2
```

`[pval, lm] = arch_test (y, x, p)`

For a linear regression model

```
y = x * b + e
```

perform a Lagrange Multiplier (LM) test of the null hypothesis of no conditional heteroscedascity against the alternative of CH($p$).

I.e., the model is

```
y(t) = b(1) * x(t,1) + ... + b(k) * x(t,k) + e(t),
```

given $y$ up to $t-1$ and $x$ up to $t$, $e$(t) is $N(0, h(t))$ with

```
h(t) = v + a(1) * e(t-1)^2 + ... + a(p) * e(t-p)^2,
```

and the null is $a(1) == \ldots == a(p) == 0$.

If the second argument is a scalar integer, $k$, perform the same test in a linear autoregression model of order $k$, i.e., with

```
[1, y(t-1), ..., y(t-k)]
```

as the $t$-th row of x.

Under the null, LM approximately has a chisquare distribution with $p$ degrees of freedom and *pval* is the $p$-value (1 minus the CDF of this distribution at LM) of the test.

If no output argument is given, the $p$-value is displayed.

`x = arma_rnd (a, b, v, t, n)`

Return a simulation of the ARMA model.

The ARMA model is defined by

```
x(n) = a(1) * x(n-1) + ... + a(k) * x(n-k)
     + e(n) + b(1) * e(n-1) + ... + b(l) * e(n-l)
```

in which $k$ is the length of vector $a$, $l$ is the length of vector $b$ and $e$ is Gaussian white noise with variance $v$. The function returns a vector of length $t$.

The optional parameter *n* gives the number of dummy *x*(*i*) used for initialization, i.e., a sequence of length *t*+*n* is generated and *x*(*n*+1:*t*+*n*) is returned. If *n* is omitted, *n* = 100 is used.

`x = autoreg_matrix (y, k)`

Given a time series (vector) *y*, return a matrix with ones in the first column and the first *k* lagged values of *y* in the other columns.

In other words, for *t* > *k*, `[1, y(t-1), ..., y(t-k)]` is the t-th row of the result.

The resulting matrix may be used as a regressor matrix in autoregressions.

`c = bartlett (m)`

Return the filter coefficients of a Bartlett (triangular) window of length *m*.

For a definition of the Bartlett window see, e.g., A.V. Oppenheim & R. W. Schafer, *Discrete-Time Signal Processing*.

`c = blackman (m)`
`c = blackman (m, "`*periodic*`")`
`c = blackman (m, "`*symmetric*`")`

Return the filter coefficients of a Blackman window of length *m*.

If the optional argument `"periodic"` is given, the periodic form of the window is returned. This is equivalent to the window of length *m*+1 with the last coefficient removed. The optional argument `"symmetric"` is equivalent to not specifying a second argument.

For a definition of the Blackman window, see, e.g., A.V. Oppenheim & R. W. Schafer, *Discrete-Time Signal Processing*.

`y = detrend (x, p)`

If *x* is a vector, `detrend (x, p)` removes the best fit of a polynomial of order *p* from the data *x*.

If *x* is a matrix, `detrend (x, p)` does the same for each column in *x*.

The second argument *p* is optional. If it is not specified, a value of 1 is assumed. This corresponds to removing a linear trend.

The order of the polynomial can also be given as a string, in which case *p* must be either `"constant"` (corresponds to `p=0`) or `"linear"` (corresponds to `p=1`).

**See also:** [polyfit], page 883.

`[d, dd] = diffpara (x, a, b)`

Return the estimator *d* for the differencing parameter of an integrated time series.

The frequencies from $[2 * pi * a/t, 2 * pi * b/T]$ are used for the estimation. If *b* is omitted, the interval $[2 * pi/T, 2 * pi * a/T]$ is used. If both *b* and *a* are omitted then $a = 0.5 * sqrt(T)$ and $b = 1.5 * sqrt(T)$ is used, where *T* is the sample size. If *x* is a matrix, the differencing parameter of each column is estimated.

The estimators for all frequencies in the intervals described above is returned in *dd*.

The value of *d* is simply the mean of *dd*.

Reference: P.J. Brockwell & R.A. Davis, *Time Series: Theory and Methods*, Springer, 1987.

`[`*newphi*`, `*newv*`] = durbinlevinson (`*c*`, `*oldphi*`, `*oldv*`)`

Perform one step of the Durbin-Levinson algorithm.

The vector *c* specifies the autocovariances `[gamma_0, ..., gamma_t]` from lag 0 to *t*, *oldphi* specifies the coefficients based on *c*(*t*-1) and *oldv* specifies the corresponding error.

If *oldphi* and *oldv* are omitted, all steps from 1 to *t* of the algorithm are performed.

*y* = `fftshift (`*x*`)`
*y* = `fftshift (`*x*`, `*dim*`)`

Perform a shift of the vector *x*, for use with the `fft` and `ifft` functions, in order to move the frequency 0 to the center of the vector or matrix.

If *x* is a vector of *N* elements corresponding to *N* time samples spaced by *dt*, then `fftshift (fft (x))` corresponds to frequencies

```
f = [ -(ceil((N-1)/2):-1:1), 0, (1:floor((N-1)/2)) ] * df
```

where $df = 1/(N * dt)$.

If *x* is a matrix, the same holds for rows and columns. If *x* is an array, then the same holds along each dimension.

The optional *dim* argument can be used to limit the dimension along which the permutation occurs.

**See also:** [ifftshift], page 933.

*x* = `ifftshift (`*y*`)`
*x* = `ifftshift (`*y*`, `*dim*`)`

Undo the action of the `fftshift` function.

For even length *x*, `fftshift` is its own inverse, but odd lengths differ slightly.

**See also:** [fftshift], page 933.

*fd* = `fractdiff (`*x*`, `*d*`)`

Compute the fractional differences $(1 - L)^d x$ where $L$ denotes the lag-operator and $d$ is greater than -1.

*c* = `hamming (`*m*`)`
*c* = `hamming (`*m*`, "`*periodic*`")`
*c* = `hamming (`*m*`, "`*symmetric*`")`

Return the filter coefficients of a Hamming window of length *m*.

If the optional argument `"periodic"` is given, the periodic form of the window is returned. This is equivalent to the window of length *m*+1 with the last coefficient removed. The optional argument `"symmetric"` is equivalent to not specifying a second argument.

For a definition of the Hamming window see, e.g., A.V. Oppenheim & R. W. Schafer, *Discrete-Time Signal Processing*.

*c* = `hanning (`*m*`)`
*c* = `hanning (`*m*`, "`*periodic*`")`
*c* = `hanning (`*m*`, "`*symmetric*`")`

Return the filter coefficients of a Hanning window of length *m*.

If the optional argument `"periodic"` is given, the periodic form of the window is returned. This is equivalent to the window of length $m$+1 with the last coefficient removed. The optional argument `"symmetric"` is equivalent to not specifying a second argument.

For a definition of the Hanning window see, e.g., A.V. Oppenheim & R. W. Schafer, *Discrete-Time Signal Processing*.

`H = hurst (x)`

Estimate the Hurst parameter of sample $x$ via the rescaled range statistic.

If $x$ is a matrix, the parameter is estimated for every column.

`pp = pchip (x, y)`
`yi = pchip (x, y, xi)`

Return the Piecewise Cubic Hermite Interpolating Polynomial (pchip) of points $x$ and $y$.

If called with two arguments, return the piecewise polynomial $pp$ that may be used with `ppval` to evaluate the polynomial at specific points.

When called with a third input argument, `pchip` evaluates the pchip polynomial at the points $xi$. The third calling form is equivalent to `ppval (pchip (x, y), xi)`.

The variable $x$ must be a strictly monotonic vector (either increasing or decreasing) of length $n$.

$y$ can be either a vector or array. If $y$ is a vector then it must be the same length $n$ as $x$. If $y$ is an array then the size of $y$ must have the form

$$[s_1, s_2, \cdots, s_k, n]$$

The array is reshaped internally to a matrix where the leading dimension is given by

$$s_1 s_2 \cdots s_k$$

and each row of this matrix is then treated separately. Note that this is exactly opposite to `interp1` but is done for MATLAB compatibility.

**See also:** [spline], page 898, [ppval], page 892, [mkpp], page 891, [unmkpp], page 891.

`[Pxx, w] = periodogram (x)`
`[Pxx, w] = periodogram (x, win)`
`[Pxx, w] = periodogram (x, win, nfft)`
`[Pxx, f] = periodogram (x, win, nfft, Fs)`
`[Pxx, f] = periodogram (..., "range")`
`periodogram (...)`

Return the periodogram (Power Spectral Density) of $x$.

The possible inputs are:

$x$

data vector. If $x$ is real-valued a one-sided spectrum is estimated. If $x$ is complex-valued, or `"range"` specifies `"twosided"`, the full spectrum is estimated.

*win* window weight data. If window is empty or unspecified a default rectangular window is used. Otherwise, the window is applied to the signal (`x .* win`) before computing the periodogram. The window data must be a vector of the same length as *x*.

*nfft* number of frequency bins. The default is 256 or the next higher power of 2 greater than the length of *x* (`max (256, 2.^nextpow2 (length (x)))`). If *nfft* is greater than the length of the input then *x* will be zero-padded to the length of *nfft*.

*Fs* sampling rate. The default is 1.

*range* range of spectrum. `"onesided"` computes spectrum from [0:nfft/2+1]. `"twosided"` computes spectrum from [0:nfft-1].

The optional second output *w* are the normalized angular frequencies. For a one-sided calculation *w* is in the range [0, pi] if *nfft* is even and [0, pi) if *nfft* is odd. Similarly, for a two-sided calculation *w* is in the range [0, 2*pi] or [0, 2*pi) depending on *nfft*.

If a sampling frequency is specified, *Fs*, then the output frequencies *f* will be in the range [0, *Fs*/2] or [0, *Fs*/2) for one-sided calculations. For two-sided calculations the range will be [0, *Fs*).

When called with no outputs the periodogram is immediately plotted in the current figure window.

**See also:** [fft], page 925.

`y = sinetone (freq, rate, sec, ampl)`

Return a sinetone of frequency *freq* with a length of *sec* seconds at sampling rate *rate* and with amplitude *ampl*.

The arguments *freq* and *ampl* may be vectors of common size.

The defaults are *rate* = 8000, *sec* = 1, and *ampl* = 64.

**See also:** [sinewave], page 935.

`y = sinewave (m, n, d)`

Return an *m*-element vector with *i*-th element given by `sin (2 * pi * (i+d-1) / n)`.

The default value for *d* is 0 and the default value for *n* is *m*.

**See also:** [sinetone], page 935.

`sde = spectral_adf (c)`
`sde = spectral_adf (c, win)`
`sde = spectral_adf (c, win, b)`

Return the spectral density estimator given a vector of autocovariances *c*, window name *win*, and bandwidth, *b*.

The window name, e.g., `"triangle"` or `"rectangle"` is used to search for a function called `win_lw`.

If *win* is omitted, the triangle window is used.

If *b* is omitted, `1 / sqrt (length (x))` is used.

**See also:** [spectral_xdf], page 936.

*sde* = spectral_xdf (*x*)
*sde* = spectral_xdf (*x*, *win*)
*sde* = spectral_xdf (*x*, *win*, *b*)

Return the spectral density estimator given a data vector *x*, window name *win*, and bandwidth, *b*.

The window name, e.g., "triangle" or "rectangle" is used to search for a function called *win*_sw.

If *win* is omitted, the triangle window is used.

If *b* is omitted, 1 / sqrt (length (*x*)) is used.

**See also:** [spectral_adf], page 935.

*savg* = spencer (*x*)

Return Spencer's 15-point moving average of each column of *x*.

*y* = stft (*x*)
*y* = stft (*x*, *win_size*)
*y* = stft (*x*, *win_size*, *inc*)
*y* = stft (*x*, *win_size*, *inc*, *num_coef*)
*y* = stft (*x*, *win_size*, *inc*, *num_coef*, *win_type*)
[*y*, *c*] = stft (...)

Compute the short-time Fourier transform of the vector *x* with *num_coef* coefficients by applying a window of *win_size* data points and an increment of *inc* points.

Before computing the Fourier transform, one of the following windows is applied:

"hanning"
: win_type = 1

"hamming"
: win_type = 2

"rectangle"
: win_type = 3

The window names can be passed as strings or by the *win_type* number.

The following defaults are used for unspecified arguments: *win_size* = 80, *inc* = 24, *num_coef* = 64, and *win_type* = 1.

*y* = stft (*x*, ...) returns the absolute values of the Fourier coefficients according to the *num_coef* positive frequencies.

[*y*, *c*] = stft (*x*, ...) returns the entire STFT-matrix *y* and a 3-element vector *c* containing the window size, increment, and window type, which is needed by the synthesis function.

**See also:** [synthesis], page 936.

*x* = synthesis (*y*, *c*)

Compute a signal from its short-time Fourier transform *y* and a 3-element vector *c* specifying window size, increment, and window type.

The values *y* and *c* can be derived by

```
[y, c] = stft (x , ...)
```

**See also:** [stft], page 936.

`[a, v] = yulewalker (c)`

Fit an AR (p)-model with Yule-Walker estimates given a vector $c$ of autocovariances `[gamma_0, ..., gamma_p]`.

Returns the AR coefficients, $a$, and the variance of white noise, $v$.

# 32 Image Processing

Since an image is basically a matrix, Octave is a very powerful environment for processing and analyzing images. To illustrate how easy it is to do image processing in Octave, the following example will load an image, smooth it by a 5-by-5 averaging filter, and compute the gradient of the smoothed image.

```
I = imread ("myimage.jpg");
S = conv2 (I, ones (5, 5) / 25, "same");
[Dx, Dy] = gradient (S);
```

In this example `S` contains the smoothed image, and `Dx` and `Dy` contains the partial spatial derivatives of the image.

## 32.1 Loading and Saving Images

The first step in most image processing tasks is to load an image into Octave which is done with the `imread` function. The `imwrite` function is the corresponding function for writing images to the disk.

In summary, most image processing code will follow the structure of this code

```
I = imread ("my_input_image.img");
J = process_my_image (I);
imwrite (J, "my_output_image.img");
```

`[`*img*`,` *map*`,` *alpha*`] = imread (`*filename*`)`
`[...] = imread (`*url*`)`
`[...] = imread (...,` *ext*`)`
`[...] = imread (...,` *idx*`)`
`[...] = imread (...,` *param1*`,` *value1*`, ...)`

Read images from various file formats.

Read an image as a matrix from the file *filename* or from the online resource *url*. If neither is given, but *ext* was specified, look for a file with the extension *ext*.

The size and class of the output depends on the format of the image. A color image is returned as an MxNx3 matrix. Grayscale and black-and-white images are of size MxN. Multipage images will have an additional 4th dimension.

The bit depth of the image determines the class of the output: `"uint8"`, `"uint16"`, or `"single"` for grayscale and color, and `"logical"` for black-and-white. Note that indexed images always return the indexes for a colormap, independent of whether *map* is a requested output. To obtain the actual RGB image, use `ind2rgb`. When more than one indexed image is being read, *map* is obtained from the first. In some rare cases this may be incorrect and `imfinfo` can be used to obtain the colormap of each image.

See the Octave manual for more information in representing images. (see Section 32.3 [Representing Images], page 947)

Some file formats, such as TIFF and GIF, are able to store multiple images in a single file. *idx* can be a scalar or vector specifying the index of the images to read. By default, Octave will read only the first page.

Depending on the file format, it is possible to configure the reading of images with *parameter*, *value* pairs. The following options are supported:

`"Frames"` or `"Index"`
: This is an alternative method to specify *idx*. When specifying it in this way, its value can also be the string `"all"`.

`"Info"`
: This option exists for MATLAB compatibility, but has no effect. For maximum performance when reading multiple images from a single file, use the `"Index"` option.

`"PixelRegion"`
: Controls the image region that is read. The value must be a cell array with two arrays of 3 elements `{[rows], [cols]}`. The elements in the array are the start, increment, and end pixel to be read. If the increment value is omitted it defaults to 1. For example, the following are all equivalent:

```
imread (filename, "PixelRegion", {[200 600], [300 700]})
imread (filename, "PixelRegion", {[200 1 600], [300 1 700]})
imread (filename)(200:600, 300:700)
```

**See also:** [imwrite], page 940, [imfinfo], page 942, [imformats], page 944.

`imwrite (`*`img`*`, `*`filename`*`)`
`imwrite (`*`img`*`, `*`filename`*`, `*`ext`*`)`
`imwrite (`*`img`*`, `*`map`*`, `*`filename`*`)`
`imwrite (...,` *`param1`*`,` *`val1`*`, ...)`
: Write images in various file formats.

The image *img* can be a binary, grayscale, RGB, or multi-dimensional image. The size and class of *img* should be the same as what should be expected when reading it with `imread`: the 3rd and 4th dimensions reserved for color space, and multiple pages respectively. If it's an indexed image, the colormap *map* must also be specified.

If *ext* is not supplied, the file extension of *filename* is used to determine the format. The actual supported formats are dependent on options made during the build of Octave. Use `imformats` to check the support of the different image formats.

Depending on the file format, it is possible to configure the writing of images with *param*, *val* pairs. The following options are supported:

'`Alpha`'
: Alpha (transparency) channel for the image. This must be a matrix with same class, and number of rows and columns of *img*. In case of a multipage image, the size of the 4th dimension must also match and the third dimension must be a singleton. By default, image will be completely opaque.

'`Compression`'
: Compression to use one the image. Can be one of the following: `"none"` (default), `"bzip"`, `"fax3"`, `"fax4"`, `"jpeg"`, `"lzw"`, `"rle"`, or `"deflate"`. Note that not all compression types are available for all image formats in which it defaults to your Magick library.

'DelayTime'
: For formats that accept animations (such as GIF), controls for how long a frame is displayed until it moves to the next one. The value must be scalar (which will applied to all frames in *img*), or a vector of length equal to the number of frames in *im*. The value is in seconds, must be between 0 and 655.35, and defaults to 0.5.

'DisposalMethod'
: For formats that accept animations (such as GIF), controls what happens to a frame before drawing the next one. Its value can be one of the following strings: "doNotSpecify" (default); "leaveInPlace"; "restoreBG"; and "restorePrevious", or a cell array of those string with length equal to the number of frames in *img*.

'LoopCount'
: For formats that accept animations (such as GIF), controls how many times the sequence is repeated. A value of Inf means an infinite loop (default), a value of 0 or 1 that the sequence is played only once (loops zero times), while a value of 2 or above loops that number of times (looping twice means it plays the complete sequence 3 times). This option is ignored when there is only a single image at the end of writing the file.

'Quality'
: Set the quality of the compression. The value should be an integer between 0 and 100, with larger values indicating higher visual quality and lower compression. Defaults to 75.

'WriteMode'
: Some file formats, such as TIFF and GIF, are able to store multiple images in a single file. This option specifies if *img* should be appended to the file (if it exists) or if a new file should be created for it (possibly overwriting an existing file). The value should be the string "Overwrite" (default), or "Append".

  Despite this option, the most efficient method of writing a multipage image is to pass a 4 dimensional *img* to `imwrite`, the same matrix that could be expected when using `imread` with the option "Index" set to "all".

**See also:** [imread], page 939, [imfinfo], page 942, [imformats], page 944.

*val* = IMAGE_PATH ()
*old_val* = IMAGE_PATH (*new_val*)
*old_val* = IMAGE_PATH (*new_val*, "*local*")
: Query or set the internal variable that specifies a colon separated list of directories in which to search for image files.

  When called from inside a function with the "local" option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

  **See also:** [EXEC_PATH], page 1062, [OCTAVE_HOME], page 1074, [OCTAVE_EXEC_HOME], page 1074.

It is possible to get information about an image file on disk, without actually reading it into Octave. This is done using the `imfinfo` function which provides read access to many of the parameters stored in the header of the image file.

*info* = imfinfo (*filename*)
*info* = imfinfo (*url*)
*info* = imfinfo (..., *ext*)

Read image information from a file.

`imfinfo` returns a structure containing information about the image stored in the file *filename*. If there is no file *filename*, and *ext* was specified, it will look for a file named *filename* and extension *ext*, i.e., a file named *filename.ext*.

The output structure *info* contains the following fields:

'`Filename`'
: The full name of the image file.

'`FileModDate`'
: Date of last modification to the file.

'`FileSize`'
: Number of bytes of the image on disk

'`Format`' Image format (e.g., `"jpeg"`).

'`Height`' Image height in pixels.

'`Width`' Image Width in pixels.

'`BitDepth`'
: Number of bits per channel per pixel.

'`ColorType`'
: Image type. Value is `"grayscale"`, `"indexed"`, `"truecolor"`, `"CMYK"`, or `"undefined"`.

'`XResolution`'
: X resolution of the image.

'`YResolution`'
: Y resolution of the image.

'`ResolutionUnit`'
: Units of image resolution. Value is `"Inch"`, `"Centimeter"`, or `"undefined"`.

'`DelayTime`'
: Time in 1/100ths of a second (0 to 65535) which must expire before displaying the next image in an animated sequence.

'`LoopCount`'
: Number of iterations to loop an animation.

'`ByteOrder`'
: Endian option for formats that support it. Value is `"little-endian"`, `"big-endian"`, or `"undefined"`.

'`Gamma`' Gamma level of the image. The same color image displayed on two different workstations may look different due to differences in the display monitor.

'`Quality`' JPEG/MIFF/PNG compression level. Value is an integer in the range [0 100].

'`DisposalMethod`'
Only valid for GIF images, control how successive frames are rendered (how the preceding frame is disposed of) when creating a GIF animation. Values can be `"doNotSpecify"`, `"leaveInPlace"`, `"restoreBG"`, or `"restorePrevious"`. For non-GIF files, value is an empty string.

'`Chromaticities`'
Value is a 1x8 Matrix with the x,y chromaticity values for white, red, green, and blue points, in that order.

'`Comment`' Image comment.

'`Compression`'
Compression type. Value can be `"none"`, `"bzip"`, `"fax3"`, `"fax4"`, `"jpeg"`, `"lzw"`, `"rle"`, `"deflate"`, `"lzma"`, `"jpeg2000"`, `"jbig2"`, `"jbig2"`, or `"undefined"`.

'`Colormap`'
Colormap for each image.

'`Orientation`'
The orientation of the image with respect to the rows and columns. Value is an integer between 1 and 8 as defined in the TIFF 6 specifications, and for MATLAB compatibility.

'`Software`'
Name and version of the software or firmware of the camera or image input device used to generate the image.

'`Make`' The manufacturer of the recording equipment. This is the manufacture of the DSC, scanner, video digitizer or other equipment that generated the image.

'`Model`' The model name or model number of the recording equipment as mentioned on the field `"Make"`.

'`DateTime`'
The date and time of image creation as defined by the Exif standard, i.e., it is the date and time the file was changed.

'`ImageDescription`'
The title of the image as defined by the Exif standard.

'`Artist`' Name of the camera owner, photographer or image creator.

'`Copyright`'
Copyright notice of the person or organization claiming rights to the image.

'DigitalCamera'
: A struct with information retrieved from the Exif tag.

'GPSInfo'
: A struct with geotagging information retrieved from the Exif tag.

**See also:** [imread], page 939, [imwrite], page 940, [imshow], page 945, [imformats], page 944.

By default, Octave's image IO functions (`imread`, `imwrite`, and `imfinfo`) use the `GraphicsMagick` library for their operations. This means a vast number of image formats is supported but considering the large amount of image formats in science and its commonly closed nature, it is impossible to have a library capable of reading them all. Because of this, the function `imformats` keeps a configurable list of available formats, their extensions, and what functions should the image IO functions use. This allows one to expand Octave's image IO capabilities by creating functions aimed at acting on specific file formats.

While it would be possible to call the extra functions directly, properly configuring Octave with `imformats` allows one to keep a consistent code that is abstracted from file formats.

It is important to note that a file format is not actually defined by its file extension and that `GraphicsMagick` is capable to read and write more file formats than the ones listed by `imformats`. What this means is that even with an incorrect or missing extension the image may still be read correctly, and that even unlisted formats are not necessarily unsupported.

`imformats ()`
*`formats`* `= imformats (`*`ext`*`)`
*`formats`* `= imformats (`*`format`*`)`
*`formats`* `= imformats ("`*`add`*`", `*`format`*`)`
*`formats`* `= imformats ("`*`remove`*`", `*`ext`*`)`
*`formats`* `= imformats ("`*`update`*`", `*`ext`*`, `*`format`*`)`
*`formats`* `= imformats ("`*`factory`*`")`

Manage supported image formats.

*formats* is a structure with information about each supported file format, or from a specific format *ext*, the value displayed on the field *ext*. It contains the following fields:

ext
: The name of the file format. This may match the file extension but Octave will automatically detect the file format.

description
: A long description of the file format.

isa
: A function handle to confirm if a file is of the specified format.

write
: A function handle to write if a file is of the specified format.

read
: A function handle to open files the specified format.

info
: A function handle to obtain image information of the specified format.

alpha
: Logical value if format supports alpha channel (transparency or matte).

multipage
: Logical value if format supports multipage (multiple images per file).

It is possible to change the way Octave manages file formats with the options `"add"`, `"remove"`, and `"update"`, and supplying a structure *format* with the required fields. The option `"factory"` resets the configuration to the default.

This can be used by Octave packages to extend the image reading capabilities Octave, through use of the PKG_ADD and PKG_DEL commands.

**See also:** [imfinfo], page 942, [imread], page 939, [imwrite], page 940.

## 32.2 Displaying Images

A natural part of image processing is visualization of an image. The most basic function for this is the `imshow` function that shows the image given in the first input argument.

`imshow (`*`im`*`)`
`imshow (`*`im`*`, `*`limits`*`)`
`imshow (`*`im`*`, `*`map`*`)`
`imshow (`*`rgb`*`, ...)`
`imshow (`*`filename`*`)`
`imshow (..., `*`string_param1`*`, `*`value1`*`, ...)`
*`h`* `= imshow (...)`

Display the image *im*, where *im* can be a 2-dimensional (grayscale image) or a 3-dimensional (RGB image) matrix.

If *limits* is a 2-element vector `[`*`low, high`*`]`, the image is shown using a display range between *low* and *high*. If an empty matrix is passed for *limits*, the display range is computed as the range between the minimal and the maximal value in the image.

If *map* is a valid colormap, the image will be shown as an indexed image using the supplied colormap.

If a filename is given instead of an image, the file will be read and shown.

If given, the parameter *string_param1* has value *value1*. *string_param1* can be any of the following:

`"displayrange"`
: *value1* is the display range as described above.

`"colormap"`
: *value1* is the colormap to use when displaying an indexed image.

`"xdata"`
: If *value1* is a 2-element vector, it must contain horizontal image limits in the form [xfirst, xlast], where xfirst and xlast are the abscissa of the centers of the corner pixels. Otherwise *value1* must be a vector and only the first and last elements will be used for xfirst and xlast respectively.

`"ydata"`
: If *value1* is a 2-element vector, it must contain vertical image limits in the form [yfirst, ylast], where yfirst and ylast are the ordinates of the center of the corner pixels. Otherwise *value1* must be a vector and only the first and last elements will be used for yfirst and ylast respectively.

The optional return value *h* is a graphics handle to the image.

**See also:** [image], page 946, [imagesc], page 946, [colormap], page 951, [gray2ind], page 948, [rgb2ind], page 948.

`image (`*`img`*`)`
`image (`*`x`*`, `*`y`*`, `*`img`*`)`
`image ("`*`CData`*`", `*`img`*`)`
`image ("`*`XData`*`", `*`x`*`, "`*`YData`*`", `*`y`*`, "`*`CData`*`", `*`img`*`)`
`image (...,` *`prop`*`, `*`val`*`)`
`image (`*`hax`*`, ...)`
*`h`* `= image (...)`

Display a matrix as an image.

*img* may be a 2-D matrix where each element is an index into the current colormap. For floating point data, the value 1 chooses the first color in the colormap. For integer data, the value 0 chooses the first color in the colormap.

Or *img* may be a 3-D matrix where the third dimension is an RGB triplet specifying the color. If the image data is floating point then the data must be in the range [0, 1]. If the image data is of integer type (uint8 or uint16) then the data must be in the range [0, INTMAX].

*x* and *y* are optional 1-element (`[min]`) or 2-element vectors (`[min, max]`) which specify the coordinate(s) of the *center* of a corner pixel. If unspecified, the default minimum value is 1 and the maximum is the length of *img* along the specific dimension. If a range is specified as `[max, min]` then the image will be reversed along that axis. For convenience, *x* and *y* may be specified as vectors, however, only the first and last elements will be used to determine the axis limits.

Multiple property/value pairs may be specified for the image object, but they must appear in pairs.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

The optional return value *h* is a graphics handle to the image.

Implementation Note: The origin (0, 0) for images is located in the upper left. For ordinary plots, the origin is located in the lower left. Octave handles this inversion by plotting the data normally, and then reversing the direction of the y-axis by setting the `ydir` property to `"reverse"`. This has implications whenever an image and an ordinary plot need to be overlaid. The recommended solution is to display the image and then plot the reversed ydata using, for example, `flipud (ydata)`.

Calling Forms: The `image` function can be called in two forms: High-Level and Low-Level. When invoked with normal options, the High-Level form is used which first calls `newplot` to prepare the graphic figure and axes. When the only inputs to `image` are property/value pairs the Low-Level form is used which creates a new instance of an image object and inserts it in the current axes (as if `hold on` was in effect).

Graphic Properties: The full list of properties is documented at Section 15.3.3.7 [Image Properties], page 492.

**See also:** [imshow], page 945, [imagesc], page 946, [colormap], page 951.

`imagesc (`*`img`*`)`
`imagesc (`*`x`*`, `*`y`*`, `*`img`*`)`
`imagesc (...,` *`climits`*`)`
`imagesc (..., "`*`prop`*`", `*`val`*`, ...)`

`imagesc ("prop1", val1, ...)`
`imagesc (hax, ...)`
`h = imagesc (...)`

Display a scaled version of the matrix *img* as a color image.

The colormap is scaled so that the entries of the matrix occupy the entire colormap. If `climits = [lo, hi]` is given, then that range is set to the `"clim"` of the current axes.

*x* and *y* are optional 2-element vectors, `[min, max]`, which specify the coordinates of the centers of the corner pixels. If a range is specified as `[max, min]` then the image will be reversed along that axis. For convenience, *x* and *y* may be specified as N-element vectors matching the length of the data in *img*. However, only the first and last elements will be used to determine the image limits.

The optional return value *h* is a graphics handle to the image.

Calling Forms: The `imagesc` function can be called in two forms: High-Level and Low-Level. When invoked with normal options, the High-Level form is used which first calls `newplot` to prepare the graphic figure and axes. When the only inputs to `image` are property/value pairs the Low-Level form is used which creates a new instance of an image object and inserts it in the current axes. The full list of properties is documented at Section 15.3.3.7 [Image Properties], page 492.

**See also:** [image], page 946, [imshow], page 945, [clim], page 368.

## 32.3 Representing Images

In general Octave supports four different kinds of images, grayscale images, RGB images, binary images, and indexed images. A grayscale image is represented with an M-by-N matrix in which each element corresponds to the intensity of a pixel. An RGB image is represented with an M-by-N-by-3 array where each 3-vector corresponds to the red, green, and blue intensities of each pixel.

The actual meaning of the value of a pixel in a grayscale or RGB image depends on the class of the matrix. If the matrix is of class `double` pixel intensities are between 0 and 1, if it is of class `uint8` intensities are between 0 and 255, and if it is of class `uint16` intensities are between 0 and 65535.

A binary image is an M-by-N matrix of class `logical`. A pixel in a binary image is black if it is `false` and white if it is `true`.

An indexed image consists of an M-by-N matrix of integers and a C-by-3 colormap. Each integer corresponds to an index in the colormap, and each row in the colormap corresponds to an RGB color. The colormap must be of class `double` with values between 0 and 1.

The following convenience functions are available for conversion between image formats.

`dimg = im2double (img)`
`dimg = im2double (img, "indexed")`

Convert image to double precision.

The conversion of *img* to double precision, is dependent on the type of input image. The following input classes are supported:

'`uint8, uint16, and int16`'

The range of values from the class is scaled to the interval [0 1].

'`logical`' True and false values are assigned a value of 1 and 0 respectively.

'`single`' Values are cast to double.

'`double`' Returns the same image.

If *img* is an indexed image, then the second argument should be the string `"indexed"`. If so, then *img* must either be of floating point class, or unsigned integer class and it will simply be cast to double. If it is an integer class, an offset of `+1` is applied.

**See also:** [double], page 51.

*img* = `gray2ind` (*I*)
*img* = `gray2ind` (*I*, *n*)
*img* = `gray2ind` (*BW*)
*img* = `gray2ind` (*BW*, *n*)
[*img*, *map*] = `gray2ind` (...)

Convert a grayscale or binary intensity image to an indexed image.

The indexed image will consist of *n* different intensity values. If not given *n* defaults to 64 for grayscale images or 2 for binary black and white images.

The output *img* is of class uint8 if *n* is less than or equal to 256; Otherwise the return class is uint16.

**See also:** [ind2gray], page 948, [rgb2ind], page 948.

*I* = `ind2gray` (*x*, *map*)

Convert a color indexed image to a grayscale intensity image.

The image *x* must be an indexed image which will be converted using the colormap *map*. Pixels in *x* that are outside the colormap are **clipped** to the range of *map*.

The output *I* is of the same class as the input *x* and may be one of `uint8`, `uint16`, `single`, or `double`.

Programming Notes: There are several ways of converting colors to grayscale intensities. This functions uses the luminance value obtained from `rgb2gray` which is `I = 0.299*R + 0.587*G + 0.114*B`. Other possibilities include the value component from `rgb2hsv` or using a single color channel from `ind2rgb`.

Index values in *x* outside of the number of colors in *map* are clipped to the range of *map*. For example, negative indices are clipped to the first color while large values, including exceptional values such as `NaN` and `Inf`, are clipped to the last color.

**See also:** [gray2ind], page 948, [ind2rgb], page 949.

[*x*, *map*] = `rgb2ind` (*rgb*)
[*x*, *map*] = `rgb2ind` (*R*, *G*, *B*)

Convert an image in red-green-blue (RGB) color space to an indexed image.

The input image *rgb* can be specified as a single matrix of size MxNx3, or as three separate variables, *R*, *G*, and *B*, its three color channels, red, green, and blue.

It outputs an indexed image *x* and a colormap *map* to interpret an image exactly the same as the input. No dithering or other form of color quantization is performed. The output class of the indexed image *x* can be uint8, uint16 or double, whichever is

required to specify the number of unique colors in the image (which will be equal to the number of rows in *map*) in order.

Multi-dimensional indexed images (of size MxNx3xK) are also supported, both via a single input (*rgb*) or its three color channels as separate variables.

**See also:** [ind2rgb], page 949, [rgb2hsv], page 960, [rgb2gray], page 961.

`rgb = ind2rgb (x, map)`
`[R, G, B] = ind2rgb (x, map)`

Convert an indexed image to red, green, and blue color components.

The image *x* must be an indexed image which will be converted using the colormap *map*. Pixels in *x* that are outside the colormap are **clipped** to the range of *map*.

Multi-dimensional indexed images (of size MxNx1xK) are also supported.

The output may be a single RGB image (MxNx3 matrix where M and N are the original image *x* dimensions, one for each of the red, green and blue channels). Alternatively, the individual red, green, and blue color matrices of size MxN may be returned.

Programming Note: Index values in *x* outside of the number of colors in *map* are clipped to the range of *map*. For example, negative indices are clipped to the first color while large values, including exceptional values such as `NaN` and `Inf`, are clipped to the last color.

**See also:** [rgb2ind], page 948, [ind2gray], page 948, [hsv2rgb], page 960.

The following function allows processing an existing image.

`X = dither (RGB, map)`
`X = dither (RGB, map, Qm, Qe)`
`BW = dither (I)`

Quantize an image using dithering to increase the apparent color resolution.

`X = dither (RGB,map)` creates an indexed image approximation. It uses the color provided in the colormap, and uses dithering to increase apparent color resolution. The Floyd-Steinberg error filter is:

$$\begin{bmatrix} & x & 7 \\ 3 & 5 & 1 \end{bmatrix} /16$$

It uses a raster scan and no weight renormalization at boundaries. The default values are used: *Qm*=5 and *Qe*=8.

*RGB* is a MxNx3 array with values in `[0, 1]` (double) or `[0, 255]` (uint8).

*map* is Cx3 matrix holding RGB triplets in `[0, 1]` (double).

*Qm* is the number of quantization bits per axis for inverse colormap (default: 5).

*Qe* is the number of quantization bits for error diffusion (default: 8, max 16).

*X* is a MxN indexed image (uint8 if c≤256, else uint16) for the colormap *map* provided.

*Qm* is the number of quantization bits along each color axis for the inverse colormap. *Qm* determines the resolution of this grid along each color axis (R, G, B). *Qm* defines

the precision of the color space discretization used to map input RGB values to those colors available in the colormap. *Qe* is the number of quantization bits for the color space error calculations in the Floyd-Steinberg error diffusion algorithm. It controls the precision of the error values that are calculated and propagated during dithering. If *Qe* < *Qm*, the error diffusion process may lose precision. Therefore dithering cannot be performed, and the function returns an undithered indexed image.

`BW = dither (I)` converts the grayscale input image *I* into binary applying dithering in the process. The output image *BW* is a black and white image where dithering creates the illusion of shades of gray.

References:

- R. W. Floyd and L. Steinberg, "An Adaptive Algorithm for Spatial Gray Scale", *International Symposium Digest of Technical Papers*, Society for Information Displays, p. 36, 1975.
- R. Ulichney, *Digital Halftoning*, The MIT Press, 1987.

**See also:** [rgb2ind], page 948, [gray2ind], page 948.

Octave also provides tools to produce and work with movie frame structures. Those structures encapsulate the image data (`"cdata"` field) together with the corresponding colormap (`"colormap"` field).

`frame = getframe ()`
`frame = getframe (hax)`
`frame = getframe (hfig)`
`frame = getframe (..., rect)`

Capture a figure or axes as a movie frame structure.

Without an argument, capture the current axes excluding ticklabels, title, and x/y/zlabels. The returned structure *frame* has a field `cdata`, which contains the actual image data in the form of an NxMx3 (RGB) uint8 matrix in physical screen pixels, and a field `colormap` which is provided for MATLAB compatibility but is always empty.

If the first argument *hax* is an axes handle, then capture this axes, rather than the current axes returned by `gca`.

If the first argument *hfig* is a figure handle then the entire corresponding figure canvas is captured.

Finally, if a second argument *rect* is provided it must be a four-element vector ([left bottom width height]) defining the region inside the figure to be captured. Regardless of the figure `"units"` property, *rect* must be defined in **pixels**.

**See also:** [im2frame], page 951, [frame2im], page 951, [movie], page 950.

`movie (mov)`
`movie (mov, n)`
`movie (mov, n, fps)`
`movie (h, ...)`

Play a movie defined by an array of frame structures.

The movie *mov* must be a struct array of frames with fields `"cdata"` and `"colormap"`, as returned by the `getframe` function. By default all images are displayed once, at 12 fps, in the current axes.

The optional argument *n* is a scalar or vector of integers that controls the number of times the movie is displayed and which particular frames are shown:

First element:

*n*(1) > 0 Play the movie *n*(1) times.

*n*(1) < 0 Play the movie `abs (n(1)` times alternatively in forward and backward order.

Other elements (if any):

Indices of the frames in *mov* that will be displayed.

If the first argument is a handle to a figure or axes *h*, the movie is played in that figure or axes instead of the current axes.

**See also:** [getframe], page 950, [im2frame], page 951, [frame2im], page 951.

`[`*x*`, `*map*`] = frame2im (`*frame*`)`

Convert movie frame to indexed image.

A movie frame is simply a struct with the fields `"cdata"` and `"colormap"`.

Support for N-dimensional images or movies is given when *frame* is a struct array. In such cases, *x* will be a MxNx1xK or MxNx3xK for indexed and RGB movies respectively, with each frame concatenated along the 4th dimension.

**See also:** [im2frame], page 951, [getframe], page 950.

*frame* `= im2frame (`*rgb*`)`
*frame* `= im2frame (`*x*`, `*map*`)`

Convert image to movie frame.

A movie frame is simply a struct with the fields `"cdata"` and `"colormap"`.

Support for N-dimensional images is given when each image projection, matrix sizes of MxN and MxNx3 for RGB images, is concatenated along the fourth dimension. In such cases, the returned value is a struct array.

**See also:** [frame2im], page 951.

The `colormap` function is used to change the colormap of the current axes or figure.

*cmap* `= colormap ()`
*cmap* `= colormap (`*map*`)`
*cmap* `= colormap ("default")`
*cmap* `= colormap (`*map_name*`)`
*cmap* `= colormap (`*hax*`, ...)`
*cmap* `= colormap (`*hfig*`, ...)`
`colormap` *map_name*

Query or set the current colormap.

With no input arguments, `colormap` returns the current colormap. If there is no current figure, a new figure will be opened and the default colormap will be returned.

`colormap (map)` sets the current colormap to *map*. The colormap should be an *n* row by 3 column matrix. The columns contain red, green, and blue intensities respectively. All entries must be between 0 and 1 inclusive. The new colormap is returned.

`colormap ("default")` restores the default colormap (the `viridis` map with 256 entries). The default colormap is returned.

The map may also be specified by a string, *map_name*, which is the name of a function that returns a colormap.

If the first argument *hax* is an axes handle, then the colormap for the specified axes is queried or set. If the first argument *hfig* is a figure handle, then the colormap for the specified figure is queried or set.

For convenience, it is also possible to use this function with the command form, `colormap` *map_name*.

The list of built-in colormaps is:

| **Map** | **Description** |
|---|---|
| viridis | default |
| turbo | colormap traversing blue, cyan, green, yellow, red; modern replacement for jet. |
| jet | colormap traversing blue, cyan, green, yellow, red. |
| cubehelix | colormap traversing black, blue, green, red, white with increasing intensity. |
| hsv | cyclic colormap traversing Hue, Saturation, Value space. |
| rainbow | colormap traversing red, yellow, blue, green, violet. |
| hot | colormap traversing black, red, orange, yellow, white. |
| cool | colormap traversing cyan, purple, magenta. |
| spring | colormap traversing magenta to yellow. |
| summer | colormap traversing green to yellow. |
| autumn | colormap traversing red, orange, yellow. |
| winter | colormap traversing blue to green. |
| gray | colormap traversing black to white in shades of gray. |
| bone | colormap traversing black, gray-blue, white. |
| copper | colormap traversing black to light copper. |
| pink | colormap traversing black, gray-pink, white. |
| ocean | colormap traversing black, dark-blue, white. |
| colorcube | equally spaced colors in RGB color space. |
| flag | cyclic 4-color map of red, white, blue, black. |
| lines | cyclic colormap with colors from axes `"ColorOrder"` property. |
| prism | cyclic 6-color map of red, orange, yellow, green, blue, violet. |

white all white colormap (no colors).

**See also:** [viridis], page 957, [turbo], page 957, [jet], page 955, [cubehelix], page 954, [hsv], page 955, [rainbow], page 956, [hot], page 955, [cool], page 954, [spring], page 957, [summer], page 957, [autumn], page 953, [winter], page 957, [gray], page 955, [bone], page 954, [copper], page 954, [pink], page 956, [ocean], page 956, [colorcube], page 954, [flag], page 955, [lines], page 956, [prism], page 956, [white], page 957.

*tf* = `iscolormap` (*cmap*)

Return true if *cmap* is a colormap.

A colormap is a real matrix, of class single or double, with 3 columns. Each row represents a single color. The 3 columns contain red, green, and blue intensities respectively.

All values in a colormap should be in the range [0, 1], but this is not enforced. Each function must decide what to do for values outside this range.

**See also:** [colormap], page 951, [rgbplot], page 953.

The following functions return predefined colormaps, the same that can be requested by name using the `colormap` function.

`rgbplot` (*cmap*)
`rgbplot` (*cmap*, *style*)
*h* = `rgbplot` (`...`)

Plot the components of a colormap.

Two different *style*s are available for displaying the *cmap*:

profile (default)
: Plot the RGB line profile of the colormap for each of the channels (red, green and blue) with the plot lines colored appropriately. Each line represents the intensity of an RGB component across the colormap.

composite
: Draw the colormap across the X-axis so that the actual index colors are visible rather than the individual color components.

The optional return value *h* is a graphics handle to the created plot.

Run `demo rgbplot` to see an example of `rgbplot` and each style option.

**See also:** [colormap], page 951.

*map* = `autumn` ()
*map* = `autumn` (*n*)

Create color colormap. This colormap ranges from red through orange to yellow.

The argument *n* must be a scalar. If *n* is not specified the length of the current colormap is used. If there is no current colormap the default value of 256 is used.

**See also:** [colormap], page 951.

*map* = bone ()
*map* = bone (*n*)

Create color colormap. This colormap varies from black to white with gray-blue shades.

The argument *n* must be a scalar. If *n* is not specified the length of the current colormap is used. If there is no current colormap the default value of 256 is used.

**See also:** [colormap], page 951.

*map* = colorcube ()
*map* = colorcube (*n*)

Create color colormap. This colormap is composed of as many equally spaced colors (not grays) in the RGB color space as possible.

If there are not a perfect number *n* of regularly spaced colors then the remaining entries in the colormap are gradients of pure red, green, blue, and gray.

The argument *n* must be a scalar. If *n* is not specified the length of the current colormap is used. If there is no current colormap the default value of 256 is used.

**See also:** [colormap], page 951.

*map* = cool ()
*map* = cool (*n*)

Create color colormap. The colormap varies from cyan to magenta.

The argument *n* must be a scalar. If *n* is not specified the length of the current colormap is used. If there is no current colormap the default value of 256 is used.

**See also:** [colormap], page 951.

*map* = copper ()
*map* = copper (*n*)

Create color colormap. This colormap varies from black to a light copper tone.

The argument *n* must be a scalar. If *n* is not specified the length of the current colormap is used. If there is no current colormap the default value of 256 is used.

**See also:** [colormap], page 951.

*map* = cubehelix ()
*map* = cubehelix (*n*)
*map* = cubehelix (*n*, *start*, *rots*, *hue*, *gamma*)

Create cubehelix colormap.

This colormap varies from black to white going though blue, green, and red tones while maintaining a monotonically increasing perception of intensity. This is achieved by traversing a color cube from black to white through a helix, hence the name cubehelix, while taking into account the perceived brightness of each channel according to the NTSC specifications from 1953.

```
rgbplot (cubehelix (256))
```

The argument *n* must be a scalar. If *n* is not specified the length of the current colormap is used. If there is no current colormap the default value of 256 is used.

Reference: D.A Green, "A colour scheme for the display of astronomical intensity images", *Bulletin of the Astronomical Society of India*, 39, p. 289, 2011.

**See also:** [colormap], page 951.

*map* = `flag ()`
*map* = `flag (`*n*`)`

Create color colormap. This colormap cycles through red, white, blue, and black with each index change.

The argument *n* must be a scalar. If *n* is not specified the length of the current colormap is used. If there is no current colormap the default value of 256 is used.

**See also:** [colormap], page 951.

*map* = `gray ()`
*map* = `gray (`*n*`)`

Create gray colormap. This colormap varies from black to white with shades of gray.

The argument *n* must be a scalar. If *n* is not specified the length of the current colormap is used. If there is no current colormap the default value of 256 is used.

**See also:** [colormap], page 951.

*map* = `hot ()`
*map* = `hot (`*n*`)`

Create color colormap. This colormap ranges from black through dark red, red, orange, yellow, to white.

The argument *n* must be a scalar. If *n* is not specified the length of the current colormap is used. If there is no current colormap the default value of 256 is used.

**See also:** [colormap], page 951.

*map* = `hsv ()`
*map* = `hsv (`*n*`)`

Create color colormap. This colormap begins with red, changes through yellow, green, cyan, blue, and magenta, before returning to red.

It is useful for displaying periodic functions. The map is obtained by linearly varying the hue through all possible values while keeping constant maximum saturation and value. The equivalent code is `hsv2rgb ([(0:N-1)'/N, ones(N,2)])`.

The argument *n* must be a scalar. If *n* is not specified the length of the current colormap is used. If there is no current colormap the default value of 256 is used.

**See also:** [colormap], page 951.

*map* = `jet ()`
*map* = `jet (`*n*`)`

Create color colormap. This colormap ranges from dark blue through blue, cyan, green, yellow, red, to dark red.

The argument *n* must be a scalar. If *n* is not specified the length of the current colormap is used. If there is no current colormap the default value of 256 is used.

Programming Note: The `jet` colormap is not perceptually uniform. Try the `viridis` colormap if that is important. For a drop-in replacement for `jet` with better perceptual characteristics try the `turbo` colormap.

**See also:** [colormap], page 951, [turbo], page 957.

*map* = `lines` ()
*map* = `lines` (*n*)

Create color colormap. This colormap is composed of the list of colors in the current axes `"ColorOrder"` property. The default is blue, orange, yellow, purple, green, light blue, and dark red.

The argument *n* must be a scalar. If *n* is not specified the length of the current colormap is used. If there is no current colormap the default value of 256 is used.

**See also:** [colormap], page 951.

*map* = `ocean` ()
*map* = `ocean` (*n*)

Create color colormap. This colormap varies from black to white with shades of blue.

The argument *n* must be a scalar. If *n* is not specified the length of the current colormap is used. If there is no current colormap the default value of 256 is used.

**See also:** [colormap], page 951.

*map* = `pink` ()
*map* = `pink` (*n*)

Create color colormap. This colormap varies from black to white with shades of gray-pink.

This colormap gives a sepia tone when used on grayscale images.

The argument *n* must be a scalar. If *n* is not specified the length of the current colormap is used. If there is no current colormap the default value of 256 is used.

**See also:** [colormap], page 951.

*map* = `prism` ()
*map* = `prism` (*n*)

Create color colormap. This colormap cycles through red, orange, yellow, green, blue and violet with each index change.

The argument *n* must be a scalar. If *n* is not specified the length of the current colormap is used. If there is no current colormap the default value of 256 is used.

**See also:** [colormap], page 951.

*map* = `rainbow` ()
*map* = `rainbow` (*n*)

Create color colormap. This colormap ranges from red through orange, yellow, green, blue, to violet.

The argument *n* must be a scalar. If *n* is not specified the length of the current colormap is used. If there is no current colormap the default value of 256 is used.

**See also:** [colormap], page 951.

*map* = spring ()
*map* = spring (*n*)

Create color colormap. This colormap varies from magenta to yellow.

The argument *n* must be a scalar. If *n* is not specified the length of the current colormap is used. If there is no current colormap the default value of 256 is used.

**See also:** [colormap], page 951.

*map* = summer ()
*map* = summer (*n*)

Create color colormap. This colormap varies from green to yellow.

The argument *n* must be a scalar. If *n* is not specified the length of the current colormap is used. If there is no current colormap the default value of 256 is used.

**See also:** [colormap], page 951.

*map* = turbo ()
*map* = turbo (*n*)

Create color colormap. This colormap ranges from dark blue through green to dark red; similar to the outdated `jet` colormap but perceptually uniform.

The argument *n* must be a scalar. If *n* is not specified the length of the current colormap is used. If there is no current colormap the default value of 256 is used.

**See also:** [colormap], page 951.

*map* = viridis ()
*map* = viridis (*n*)

Create color colormap. This colormap ranges from dark purplish-blue through blue, green, to yellow.

The argument *n* must be a scalar. If *n* is not specified the length of the current colormap is used. If there is no current colormap the default value of 256 is used.

**See also:** [colormap], page 951.

*map* = white ()
*map* = white (*n*)

Create color colormap. This colormap is completely white.

The argument *n* must be a scalar. If *n* is not specified the length of the current colormap is used. If there is no current colormap the default value of 256 is used.

**See also:** [colormap], page 951.

*map* = winter ()
*map* = winter (*n*)

Create color colormap. This colormap varies from blue to green.

The argument *n* must be a scalar. If *n* is not specified the length of the current colormap is used. If there is no current colormap the default value of 256 is used.

**See also:** [colormap], page 951.

*cmap* = contrast (*x*)
*cmap* = contrast (*x*, *n*)

Return a gray colormap that maximizes the contrast in an image.

The returned colormap will have *n* rows. If *n* is not defined then the size of the current colormap is used.

**See also:** [colormap], page 951, [brighten], page 958.

The following three functions modify the existing colormap rather than replace it.

*map_out* = brighten (*beta*)
*map_out* = brighten (*map*, *beta*)
*map_out* = brighten (*h*, *beta*)
brighten (...)

Brighten or darken a colormap.

The argument *beta* must be a scalar between -1 and 1, where a negative value darkens and a positive value brightens the colormap.

If the *map* argument is omitted, the function is applied to the current colormap.

The first argument can also be a valid graphics handle *h*, in which case `brighten` is applied to the colormap associated with this handle.

If no output is specified then the result is written to the current colormap.

**See also:** [colormap], page 951, [contrast], page 958.

spinmap ()
spinmap (*t*)
spinmap (*t*, *inc*)
spinmap ("*inf*")

Cycle the colormap for *t* seconds with a color increment of *inc*.

Both parameters are optional. The default cycle time is 5 seconds and the default increment is 2. If the option `"inf"` is given then cycle continuously until *Control-C* is pressed.

When rotating, the original color 1 becomes color 2, color 2 becomes color 3, etc. A positive or negative increment is allowed and a higher value of *inc* will cause faster cycling through the colormap.

**See also:** [colormap], page 951.

whitebg ()
whitebg (*color*)
whitebg ("*none*")
whitebg (*hfig*)
whitebg (*hfig*, *color*)
whitebg (*hfig*, "*none*")

Invert the colors in the current color scheme.

The root properties are also inverted such that all subsequent plots will use the new color scheme.

If the optional argument *color* is present then the background color is set to *color* rather than inverted. *color* may be a string representing one of the eight known colors or an RGB triplet. The special string argument `"none"` restores the plot to the factory default colors.

If the first argument *hfig* is a figure handle or list of figure handles, then operate on these figures rather than the current figure returned by `gcf`. The root properties will not be changed unless 0 is in the list of figures.

Programming Note: `whitebg` operates by changing the color properties of the children of the specified figures. Only objects with a single color are affected. For example, a patch with a single `"FaceColor"` will be changed, but a patch with shading (`"interp"`) will not be modified. For inversion, the new color is simply the inversion in RGB space: *cnew* `= [1-`*R* `1-`*G* `1-`*B*`]`. When a color is specified, the axes and figure are set to the new color, and the color of child objects are then adjusted to have some contrast (visibility) against the new background.

**See also:** [reset], page 544, [get], page 457, [set], page 457.

The following functions can be used to manipulate colormaps.

`[`*Y*`, `*newmap*`] = cmunique (`*X*`, `*map*`)`
`[`*Y*`, `*newmap*`] = cmunique (`*RGB*`)`
`[`*Y*`, `*newmap*`] = cmunique (`*I*`)`

Convert an input image $X$ to an output indexed image $Y$ which uses the smallest colormap possible *newmap*.

When the input is an indexed image ($X$ with colormap *map*) the output is a colormap *newmap* from which any repeated rows have been eliminated. The output image, $Y$, is the original input image with the indices adjusted to match the new, possibly smaller, colormap.

When the input is an RGB image (an MxNx3 array), the output colormap will contain one entry for every unique color in the original image. In the worst case the new map could have as many rows as the number of pixels in the original image.

When the input is a grayscale image $I$, the output colormap will contain one entry for every unique intensity value in the original image. In the worst case the new map could have as many rows as the number of pixels in the original image.

Implementation Details:

*newmap* is always an Mx3 matrix, even if the input image is an intensity grayscale image $I$ (all three RGB planes are assigned the same value).

The output image is of class uint8 if the size of the new colormap is less than or equal to 256. Otherwise, the output image is of class double.

**See also:** [rgb2ind], page 948, [gray2ind], page 948.

`[`*Y*`, `*newmap*`] = cmpermute (`*X*`, `*map*`)`
`[`*Y*`, `*newmap*`] = cmpermute (`*X*`, `*map*`, `*index*`)`

Reorder colors in a colormap.

When called with only two arguments, `cmpermute` randomly rearranges the colormap *map* and returns a new colormap *newmap*. It also returns the indexed image $Y$ which is the equivalent of the original input image $X$ when displayed using *newmap*.

When called with an optional third argument the order of colors in the new colormap is defined by *index*.

**Caution:** *index* should not have repeated elements or the function will fail.

## 32.4 Plotting on top of Images

If gnuplot is being used to display images it is possible to plot on top of images. Since an image is a matrix it is indexed by row and column values. The plotting system is, however, based on the traditional $(x, y)$ system. To minimize the difference between the two systems Octave places the origin of the coordinate system in the point corresponding to the pixel at $(1, 1)$. So, to plot points given by row and column values on top of an image, one should simply call `plot` with the column values as the first argument and the row values as the second. As an example the following code generates an image with random intensities between 0 and 1, and shows the image with red circles over pixels with an intensity above 0.99.

```
I = rand (100, 100);
[row, col] = find (I > 0.99);
hold ("on");
imshow (I);
plot (col, row, "ro");
hold ("off");
```

## 32.5 Color Conversion

Octave supports conversion from the RGB color system to the HSV color system and vice versa. It is also possible to convert from a color RGB image to a grayscale image.

*hsv_map* = rgb2hsv (*rgb_map*)
*hsv_img* = rgb2hsv (*rgb_img*)

Transform a colormap or image from RGB to HSV color space.

A color in the RGB space consists of red, green, and blue intensities.

A color in HSV space is represented by hue, saturation and value (brightness) levels in a cylindrical coordinate system. Hue is the azimuth and describes the dominant color. Saturation is the radial distance and gives the amount of hue mixed into the color. Value is the height and is the amount of light in the color.

Output class and size will be the same as input.

**See also:** [hsv2rgb], page 960, [rgb2ind], page 948, [rgb2gray], page 961.

*rgb_map* = hsv2rgb (*hsv_map*)
*rgb_img* = hsv2rgb (*hsv_img*)

Transform a colormap or image from HSV to RGB color space.

A color in HSV space is represented by hue, saturation and value (brightness) levels in a cylindrical coordinate system. Hue is the azimuth and describes the dominant color. Saturation is the radial distance and gives the amount of hue mixed into the color. Value is the height and is the amount of light in the color.

The input can be both a colormap or RGB image. In the case of floating point input, values are expected to be on the [0 1] range. In the case of hue (azimuth), since the value corresponds to an angle, `mod (h, 1)` is used.

```
>> hsv2rgb ([0.5 1 1])
⇒ ans = 0 1 1

>> hsv2rgb ([2.5 1 1])
⇒ ans = 0 1 1

>> hsv2rgb ([3.5 1 1])
⇒ ans = 0 1 1
```

Output class and size will be the same as input.

**See also:** [rgb2hsv], page 960, [ind2rgb], page 949.

*I* = rgb2gray (*rgb_img*)
*gray_map* = rgb2gray (*rgb_map*)

Transform an image or colormap from red-green-blue (RGB) color space to a grayscale intensity image.

The input may be of class uint8, int8, uint16, int16, single, or double. The output is of the same class as the input.

Implementation Note: The grayscale intensity is calculated as

```
I = 0.298936*R + 0.587043*G + 0.114021*B
```

which corresponds to the luminance channel when RGB is translated to YIQ as documented in https://en.wikipedia.org/wiki/YIQ.

**See also:** [rgb2hsv], page 960, [rgb2ind], page 948.

# 33 Audio Processing

## 33.1 Audio File Utilities

The following functions allow you to read, write and retrieve information about audio files. Various formats are supported including wav, flac and ogg vorbis.

*info* = audioinfo (*filename*)

Return information about an audio file specified by *filename*.

The output *info* is a structure containing the following fields:

'Filename'
: Name of the audio file.

'CompressionMethod'
: Audio compression method. Unused, only present for compatibility with MATLAB.

'NumChannels'
: Number of audio channels.

'SampleRate'
: Sample rate of the audio, in Hertz.

'TotalSamples'
: Number of samples in the file.

'Duration'
: Duration of the audio, in seconds.

'BitsPerSample'
: Number of bits per sample.

'BitRate'
: Audio bit rate. Unused, only present for compatibility with MATLAB.

'Title'
: "Title" audio metadata value as a string, or empty if not present.

'Artist'
: "Artist" audio metadata value as a string, or empty if not present.

'Comment'
: "Comment" audio metadata value as a string, or empty if not present.

**See also:** [audioread], page 963, [audiowrite], page 964.

[*y*, *fs*] = audioread (*filename*)
[*y*, *fs*] = audioread (*filename*, *samples*)
[*y*, *fs*] = audioread (*filename*, *datatype*)
[*y*, *fs*] = audioread (*filename*, *samples*, *datatype*)

Read the audio file *filename* and return the audio data *y* and sampling rate *fs*.

The audio data is stored as matrix with rows corresponding to audio frames and columns corresponding to channels.

The optional two-element vector argument *samples* specifies starting and ending frames.

The optional argument *datatype* specifies the datatype to return. If it is "native", then the type of data depends on how the data is stored in the audio file.

**See also:** [audiowrite], page 964, [audioformats], page 964, [audioinfo], page 963.

`audiowrite (`*`filename`*`, `*`y`*`, `*`fs`*`)`
`audiowrite (`*`filename`*`, `*`y`*`, `*`fs`*`, `*`name`*`, `*`value`*`, ...)`

Write audio data from the matrix *y* to *filename* at sampling rate *fs* with the file format determined by the file extension.

Additional name/value argument pairs may be used to specify the following options:

'`BitsPerSample`'
: Number of bits per sample. Valid values are 8, 16, 24, and 32. Default is 16.

'`BitRate`'
: Valid argument name, but ignored. Left for compatibility with MATLAB.

'`Quality`'
: Quality setting for the Ogg Vorbis compressor. Values can range between 0 and 100 with 100 being the highest quality setting. Default is 75.

'`Title`'
: Title for the audio file.

'`Artist`'
: Artist name.

'`Comment`'
: Comment.

**See also:** [audioread], page 963, [audioformats], page 964, [audioinfo], page 963.

`audioformats ()`
`audioformats (`*`format`*`)`

Display information about all supported audio formats.

If the optional argument *format* is given, then display only formats with names that start with *format*.

**See also:** [audioread], page 963, [audiowrite], page 964.

## 33.2 Audio Device Information

*`devinfo`* `= audiodevinfo ()`
*`devs`* `= audiodevinfo (`*`io`*`)`
*`name`* `= audiodevinfo (`*`io`*`, `*`id`*`)`
*`id`* `= audiodevinfo (`*`io`*`, `*`name`*`)`
*`driverversion`* `= audiodevinfo (`*`io`*`, `*`id`*`, "`*`DriverVersion`*`")`
*`id`* `= audiodevinfo (`*`io`*`, `*`rate`*`, `*`bits`*`, `*`chans`*`)`
*`supports`* `= audiodevinfo (`*`io`*`, `*`id`*`, `*`rate`*`, `*`bits`*`, `*`chans`*`)`

Return a structure describing the available audio input and output devices.

The *devinfo* structure has two fields `"input"` and `"output"`. The value of each field is a structure array with fields `"Name"`, `"DriverVersion"` and `"ID"` describing an audio device.

If the optional argument *io* is 1, return information about input devices only. If it is 0, return information about output devices only. If *io* is the only argument supplied, return the number of input or output devices available.

If the optional argument *id* is provided, return information about the corresponding device.

If the optional argument *name* is provided, return the ID of the named device.

If the optional argument `"DriverVersion"` is given, return the name of the driver for the specified device.

Given a sampling rate, bits per sample, and number of channels for an input or output device, return the ID of the first device that supports playback or recording using the specified parameters.

If also given a device ID, return true if the device supports playback or recording using those parameters.

## 33.3 Audio Player

The following methods are used to create and use audioplayer objects. These objects can be used to play back audio data stored in Octave matrices and arrays. The audioplayer object supports playback from various devices available to the system, blocking and non-blocking playback, convenient pausing and resuming and much more.

*player* = audioplayer (*y*, *fs*)
*player* = audioplayer (*y*, *fs*, *nbits*)
*player* = audioplayer (*y*, *fs*, *nbits*, *id*)
*player* = audioplayer (*recorder*)
*player* = audioplayer (*recorder*, *id*)

Create an audioplayer object that will play back data *y* at sample rate *fs*.

The signal *y* can be a vector (mono audio) or a two-dimensional array (multi-channel audio).

The optional arguments *nbits* and *id* specify the number of bits per sample and player device ID, respectively. Device IDs may be found using the `audiodevinfo` function.

Given an audiorecorder object *recorder*, use the data from the object to initialize the player.

The list of actions for an audioplayer object are shown below. All methods require an audioplayer object as the first argument.

| **Method** | **Description** |
|---|---|
| get | Read audioplayer property values |
| isplaying | Return true if audioplayer is playing |
| pause | Pause audioplayer playback |
| play | Play audio stored in audioplayer object w/o blocking |
| playblocking | Play audio stored in audioplayer object |
| resume | Resume playback after pause |
| set | Write audioplayer property values |
| stop | Stop playback |

Example

Create an audioplayer object that will play back one second of white noise at 44100 sample rate using 8 bits per sample.

```
y = 0.25 * randn (2, 44100);
player = audioplayer (y, 44100, 8);
play (player);
```

**See also:** [@audioplayer/get], page 967, [@audioplayer/isplaying], page 967, [@audioplayer/pause], page 966, [@audioplayer/play], page 966, [@audioplayer/playblocking], page 966, [@audioplayer/resume], page 966, [@audioplayer/set], page 967, [@audioplayer/stop], page 966, [audiodevinfo], page 964, [@audiorecorder/audiorecorder], page 967, [sound], page 971, [soundsc], page 971.

### 33.3.1 Playback

The following methods are used to control player playback.

`play (`*`player`*`)`
`play (`*`player`*`, `*`start`*`)`
`play (`*`player`*`, [`*`start`*`, `*`end`*`])`

Play audio stored in the audioplayer object *player* without blocking.

If the optional argument *start* is provided, begin playing *start* samples in to the recording.

If the optional argument *end* is provided, stop playing at *end* samples into the recording.

**See also:** [@audioplayer/playblocking], page 966, [@audioplayer/pause], page 966, [@audioplayer/stop], page 966, [@audioplayer/audioplayer], page 965.

`playblocking (`*`player`*`)`
`playblocking (`*`player`*`, `*`start`*`)`
`playblocking (`*`player`*`, [`*`start`*`, `*`end`*`])`

Play audio stored in the audioplayer object *player* with blocking (synchronous I/O).

If the optional argument *start* is provided, begin playing *start* samples into the recording.

If the optional argument *end* is provided, stop playing at *end* samples into the recording.

**See also:** [@audioplayer/play], page 966, [@audioplayer/pause], page 966, [@audioplayer/stop], page 966, [@audioplayer/audioplayer], page 965.

`pause (`*`player`*`)`

Pause playback of audioplayer *player*.

**See also:** [@audioplayer/resume], page 966, [@audioplayer/stop], page 966, [@audioplayer/audioplayer], page 965.

`resume (`*`player`*`)`

Resume playback for the paused audioplayer object *player*.

**See also:** [@audioplayer/pause], page 966, [@audioplayer/stop], page 966, [@audioplayer/audioplayer], page 965.

`stop (`*`player`*`)`

Stop playback of the audioplayer *player* and reset relevant variables to their initial values.

**See also:** [@audioplayer/pause], page 966, [@audioplayer/resume], page 966, [@audioplayer/audioplayer], page 965.

*tf* = isplaying (*player*)
: Return true if the audioplayer object *player* is currently playing back audio and false otherwise.

  **See also:** [@audioplayer/pause], page 966, [@audioplayer/audioplayer], page 965.

### 33.3.2 Player Properties

The remaining couple of methods are used to get and set various properties of the audioplayer object.

*value* = get (*player*, *name*)
*values* = get (*player*, {*name1*, *name2*, ...})
*values* = get (*player*)
: Return the *value* of the property identified by *name*.

  If *name* is a cell array return the values of the properties identified by the elements of the cell array. Given only the player object, return a scalar structure with values for all properties of *player*. The field names of the structure correspond to the property names.

  **See also:** [@audioplayer/set], page 967, [@audioplayer/audioplayer], page 965.

set (*player*, *name*, *value*)
set (*player*, *name_cell*, *value_cell*)
set (*player*, *properties_struct*)
*properties* = set (*player*)
: Set the value of property specified by *name* to a given *value*.

  If *name* and *value* are cell arrays, set each property to the corresponding value. Given a structure of properties with fields corresponding to property names, set the value of those properties to the corresponding field values. Given only an audioplayer object, return a structure of configurable properties (i.e., writeable properties).

  **See also:** [@audioplayer/get], page 967, [@audioplayer/audioplayer], page 965.

## 33.4 Audio Recorder

The following methods are used to create and use audiorecorder objects. These objects can be used to record audio data from various devices available to the system. You can use convenient methods to retrieve that data or audioplayer objects created from that data. Methods for blocking and non-blocking recording, pausing and resuming recording and much more is available.

*recorder* = audiorecorder ()
*recorder* = audiorecorder (*fs*, *nbits*, *nchannels*)
*recorder* = audiorecorder (*fs*, *nbits*, *nchannels*, *id*)
: Create an audiorecorder object recording 8-bit mono audio at 8000 Hz sample rate.

  The optional arguments *fs*, *nbits*, *nchannels*, and *id* specify the sample rate, number of bits per sample, number of channels, and recording device ID, respectively. Device IDs may be found using the `audiodevinfo` function.

  The list of actions for an audiorecorder object are shown below. All methods require an audiorecorder object as the first argument.

| Method | Description |
|---|---|
| get | Read audiorecorder property values |
| getaudiodata | Return audio data as a numeric matrix |
| getplayer | Return audioplayer loaded with data from audiorecorder |
| isrecording | Return true if audiorecorder is recording |
| pause | Pause recording |
| play | Play audio stored in audiorecorder object |
| record | Record audio in audiorecorder object w/o blocking |
| recordblocking | Record audio in audiorecorder object |
| resume | Resume recording after pause |
| set | Write audiorecorder property values |
| stop | Stop recording |

**See also:** [@audiorecorder/get], page 969, [@audiorecorder/getaudiodata], page 969, [@audiorecorder/getplayer], page 969, [@audiorecorder/isrecording], page 969, [@audiorecorder/pause], page 968, [@audiorecorder/play], page 969, [@audiorecorder/record], page 968, [@audiorecorder/recordblocking], page 968, [@audioplayer/resume], page 966, [@audiorecorder/set], page 970, [@audiorecorder/stop], page 969, [audiodevinfo], page 964, [@audioplayer/audioplayer], page 965, [record], page 971.

### 33.4.1 Recording

The following methods control the recording process.

`record (`*`recorder`*`)`
`record (`*`recorder`*`, `*`length`*`)`

Record audio without blocking using the audiorecorder object *recorder* until paused or stopped by the *pause* or *stop* methods.

Given the optional argument *length*, record for *length* seconds.

**See also:** [@audiorecorder/recordblocking], page 968, [@audiorecorder/audiorecorder], page 967.

`recordblocking (`*`recorder`*`, `*`length`*`)`

Record audio with blocking (synchronous I/O).

The length of the recording in seconds (*length*) must be specified.

**See also:** [@audiorecorder/record], page 968, [@audiorecorder/audiorecorder], page 967.

`pause (`*`recorder`*`)`

Pause recording for audiorecorder *recorder*.

**See also:** [@audiorecorder/resume], page 968, [@audiorecorder/stop], page 969, [@audiorecorder/audiorecorder], page 967.

`resume (`*`recorder`*`)`

Resume recording with the paused audiorecorder object *recorder*.

**See also:** [@audiorecorder/pause], page 968, [@audiorecorder/stop], page 969, [@audiorecorder/audiorecorder], page 967.

`stop (recorder)`
: Stop recording with audiorecorder object *recorder* and clean up any audio streams.

  **See also:** [@audiorecorder/pause], page 968, [@audiorecorder/resume], page 968, [@audiorecorder/audiorecorder], page 967.

`tf = isrecording (recorder)`
: Return true if the audiorecorder object *recorder* is currently recording audio and false otherwise.

  **See also:** [@audiorecorder/pause], page 968, [@audiorecorder/audiorecorder], page 967.

### 33.4.2 Data Retrieval

The following methods allow you to retrieve recorded audio data in various ways.

`data = getaudiodata (recorder)`
`data = getaudiodata (recorder, datatype)`
: Return audio data from audiorecorder object *recorder* as a double matrix with values between -1.0 and 1.0 and with as many columns as there are channels in *recorder*.

  Given the optional argument *datatype*, convert the recorded data to the specified type, which may be one of `"double"`, `"single"`, `"int16"`, `"int8"` or `"uint8"`.

  **See also:** [@audiorecorder/audiorecorder], page 967.

`player = getplayer (recorder)`
: Return an audioplayer object with data recorded by the audiorecorder object *recorder*.

  **See also:** [@audioplayer/audioplayer], page 965, [@audiorecorder/audiorecorder], page 967.

`player = play (recorder)`
`player = play (recorder, start)`
`player = play (recorder, [start, end])`
: Play the audio recorded in *recorder* without blocking and return a corresponding audioplayer object.

  If the optional argument *start* is provided, begin playing *start* seconds into the recording.

  If the optional argument *end* is provided, stop playing at *end* seconds into the recording.

  **See also:** [@audiorecorder/getplayer], page 969, [@audioplayer/audioplayer], page 965, [@audiorecorder/audiorecorder], page 967.

### 33.4.3 Recorder Properties

The remaining two methods allow you to read or alter the properties of audiorecorder objects.

`value = get (recorder, name)`
`values = get (recorder, {name1, name2, ...})`

*values* = get (*recorder*)

Return the *value* of the property identified by *name*.

If *name* is a cell array return the values of the properties identified by the elements of the cell array. Given only the recorder object, return a scalar structure with values for all properties of *recorder*. The field names of the structure correspond to the property names.

**See also:** [@audiorecorder/set], page 970, [@audiorecorder/audiorecorder], page 967.

set (*recorder*, *name*, *value*)
set (*recorder*, *name_cell*, *value_cell*)
set (*recorder*, *properties_struct*)
*properties* = set (*recorder*)

Set the value of property specified by *name* to a given *value*.

If *name* and *value* are cell arrays, set each property to a corresponding value. Given a structure of properties with fields corresponding to property names, set the value of those properties to the corresponding field values. Given only a recorder object, return a structure of configurable properties (i.e., writeable properties).

**See also:** [@audiorecorder/get], page 969, [@audiorecorder/audiorecorder], page 967.

## 33.5 Audio Data Processing

Octave provides a few functions for dealing with audio data. An audio 'sample' is a single output value from an A/D converter, i.e., a small integer number (usually 8 or 16 bits), and audio data is just a series of such samples. It can be characterized by three parameters: the sampling rate (measured in samples per second or Hz, e.g., 8000 or 44100), the number of bits per sample (e.g., 8 or 16), and the number of channels (1 for mono, 2 for stereo, etc.).

There are many different formats for representing such data. Currently, only the two most popular, *linear encoding* and *mu-law encoding*, are supported by Octave. There is an excellent FAQ on audio formats by Guido van Rossum `guido@cwi.nl` which can be found at any FAQ ftp site, in particular in the directory `/pub/usenet/news.answers/audio-fmts` of the archive site `rtfm.mit.edu`.

Octave simply treats audio data as vectors of samples (non-mono data are not supported yet). It is assumed that audio files using linear encoding have one of the extensions `lin` or `raw`, and that files holding data in mu-law encoding end in `au`, `mu`, or `snd`.

*y* = lin2mu (*x*)
*y* = lin2mu (*x*, *n*)

Convert audio data from linear to mu-law.

Linear values use floating point values in the range -1 $\leq x \leq$ 1 if $n$ is 0 (default), or $n$-bit signed integers if $n$ is 8 or 16. Mu-law values are 8-bit unsigned integers in the range $0 \leq y \leq 255$.

**See also:** [mu2lin], page 970.

*y* = mu2lin (*x*)
*y* = mu2lin (*x*, *n*)

Convert audio data from mu-law to linear.

Mu-law values are 8-bit unsigned integers in the range $0 \leq y \leq 255$. Linear values use floating point values in the range -*linmax* $\leq x$ *linmax* (where `linmax = 32124/32768 =~ 0.98`) when *n* is zero (default). If *n* is 8 or 16 then *n*-bit signed integers are used instead.

Programming Note: `mu2lin` maps maximum mu-law inputs to values slightly below the maximum ([-0.98, +0.98]) representable with a linear scale. Because of this, `mu2lin (lin2mu (x))` might not reproduce the original input.

**See also:** [lin2mu], page 970.

`data = record (sec)`
`data = record (sec, fs)`

Record *sec* seconds of audio from the system's default audio input at a sampling rate of 8000 samples per second.

If the optional argument *fs* is given, it specifies the sampling rate for recording.

For more control over audio recording, use the `audiorecorder` class.

**See also:** [@audiorecorder/audiorecorder], page 967, [sound], page 971, [soundsc], page 971.

`sound (y)`
`sound (y, fs)`
`sound (y, fs, nbits)`

Play audio data *y* at sample rate *fs* to the default audio device.

The audio signal *y* can be a vector or a two-column array representing mono or stereo audio, respectively.

If *fs* is not given, a default sample rate of 8000 samples per second is used.

The optional argument *nbits* specifies the bit depth to play to the audio device and defaults to 8 bits.

For more control over audio playback, use the `audioplayer` class.

**See also:** [soundsc], page 971, [@audioplayer/audioplayer], page 965, [record], page 971.

`soundsc (y)`
`soundsc (y, fs)`
`soundsc (y, fs, nbits)`
`soundsc (..., [ymin, ymax])`

Scale the audio data *y* and play it at sample rate *fs* to the default audio device.

The audio signal *y* can be a vector or a two-column array representing mono or stereo audio, respectively.

If *fs* is not given, a default sample rate of 8000 samples per second is used.

The optional argument *nbits* specifies the bit depth to play to the audio device and defaults to 8 bits.

By default, *y* is automatically normalized to the range [-1, 1]. If the range [*ymin*, *ymax*] is given, then elements of *y* that fall within the range *ymin* $\leq y \leq$ *ymax* are scaled to the range [-1, 1] instead.

For more control over audio playback, use the `audioplayer` class.

**See also:** [sound], page 971, [@audioplayer/audioplayer], page 965, [record], page 971.

# 34 Object Oriented Programming

Octave has the ability to create user-defined classes—including the capabilities of operator and function overloading. Classes can protect internal properties so that they may not be altered accidentally which facilitates data encapsulation. In addition, rules can be created to address the issue of class precedence in mixed class operations.

This chapter discusses the means of constructing a user class, how to query and set the properties of a class, and how to overload operators and functions. Throughout this chapter real code examples are given using a class designed for polynomials.

## 34.1 Creating a Class

This chapter illustrates user-defined classes and object oriented programming through a custom class designed for polynomials. This class was chosen for its simplicity which does not distract unnecessarily from the discussion of the programming features of Octave. Even so, a bit of background on the goals of the polynomial class is necessary before the syntax and techniques of Octave object oriented programming are introduced.

The polynomial class is used to represent polynomials of the form

$$a_0 + a_1 x + a_2 x^2 + \ldots a_n x^n$$

where $a_0$, $a_1$, etc. are elements of $\Re$. Thus the polynomial can be represented by a vector

```
a = [a0, a1, a2, ..., an];
```

This is a sufficient specification to begin writing the constructor for the polynomial class. All object oriented classes in Octave must be located in a directory that is the name of the class prepended with the '`@`' symbol. For example, the polynomial class will have all of its methods defined in the `@polynomial` directory.

The constructor for the class must be the name of the class itself; in this example the constructor resides in the file `@polynomial/polynomial.m`. Ideally, even when the constructor is called with no arguments it should return a valid object. A constructor for the polynomial class might look like

```
## -*- texinfo -*-
## @deftypefn  {} {} polynomial ()
## @deftypefnx {} {} polynomial (@var{a})
## Create a polynomial object representing the polynomial
##
## @example
## a0 + a1 * x + a2 * x^2 + @dots{} + an * x^n
## @end example
##
## @noindent
## from a vector of coefficients [a0 a1 a2 @dots{} an].
## @end deftypefn

function p = polynomial (a)
```

```
    if (nargin == 0)
      p.poly = 0;
      p = class (p, "polynomial");
    else
      if (isa (a, "polynomial"))
        p = a;
      elseif (isreal (a) && isvector (a))
        p.poly = a(:).';  # force row vector
        p = class (p, "polynomial");
      else
        error ("@polynomial: A must be a real vector");
      endif
    endif

  endfunction
```

Note that the return value of the constructor must be the output of the `class` function. The first argument to the `class` function is a structure and the second is the name of the class itself. An example of calling the class constructor to create an instance is

```
p = polynomial ([1, 0, 1]);
```

Methods are defined by m-files in the class directory and can have embedded documentation the same as any other m-file. The help for the constructor can be obtained by using the constructor name alone, that is, for the polynomial constructor `help polynomial` will return the help string. Help can be restricted to a particular class by using the class directory name followed by the method. For example, `help @polynomial/polynomial` is another way of displaying the help string for the polynomial constructor. This second means is the only way to obtain help for the overloaded methods and functions of a class.

The same specification mechanism can be used wherever Octave expects a function name. For example `type @polynomial/disp` will print the code of the `disp` method of the polynomial class to the screen, and `dbstop @polynomial/disp` will set a breakpoint at the first executable line of the `disp` method of the polynomial class.

To check whether a variable belongs to a user class, the `isobject` and `isa` functions can be used. For example:

```
p = polynomial ([1, 0, 1]);
isobject (p)
  ⇒ 1
isa (p, "polynomial")
  ⇒ 1
```

*tf* = `isobject` (*x*)

Return true if *x* is a class object.

**See also:** [class], page 41, [typeinfo], page 41, [isa], page 41, [ismethod], page 975, [isprop], page 449.

The available methods of a class can be displayed with the `methods` function.

`methods` (*obj*)
`methods` ("*classname*")

`methods (...,` "*-full*")
*mtds* = `methods (...)`

List the names of the public methods for the object *obj* or the named class *classname*.

*obj* may be an Octave class object or a Java object. *classname* may be the name of an Octave class or a Java class.

If the optional argument `"-full"` is given then Octave returns full method signatures which include output type, name of method, and the number and type of inputs.

When called with no output arguments, `methods` prints the list of method names to the screen. Otherwise, the output argument *mtds* contains the list in a cell array of strings.

**See also:** [ismethod], page 975, [properties], page 994, [fieldnames], page 125.

To inquire whether a particular method exists for a user class, the `ismethod` function can be used.

*tf* = `ismethod (`*obj*, *method*`)`
*tf* = `ismethod (`*class_name*, *method*`)`

Return true if the string *method* is a valid method of the object *obj* or of the class *clsname*.

**See also:** [isprop], page 449, [isobject], page 974, [isjava], page 1143, [methods], page 974.

For a polynomial class it makes sense to have a method to compute its roots.

```
function r = roots (p)
  r = roots (fliplr (p.poly));
endfunction
```

We can check for the existence of the `roots`-method by calling:

```
p = polynomial ([1, 0, 1]);
ismethod (p, "roots")
  ⇒ 1
```

## 34.2 Class Methods

There are a number of basic class methods that can (and should) be defined to allow the contents of the classes to be queried and set. The most basic of these is the `disp` method. The `disp` method is used by Octave whenever a class should be displayed on the screen. Usually this is the result of an Octave expression that doesn't end with a semicolon. If this method is not defined, then Octave won't print anything when displaying the contents of a class which can be confusing.

An example of a `disp` method for the polynomial class might be

```
function disp (p)

  a = p.poly;
  first = true;
  for i = 1 : length (a);
    if (a(i) != 0)
```

```
      if (first)
        first = false;
      elseif (a(i) > 0 || isnan (a(i)))
        printf (" +");
      endif
      if (a(i) < 0)
        printf (" -");
      endif
      if (i == 1)
        printf (" %.5g", abs (a(i)));
      elseif (abs (a(i)) != 1)
        printf (" %.5g *", abs (a(i)));
      endif
      if (i > 1)
        printf (" X");
      endif
      if (i > 2)
        printf (" ^ %d", i - 1);
      endif
    endif
  endfor

  if (first)
    printf (" 0");
  endif
  printf ("\n");

endfunction
```

To be consistent with the Octave graphic handle classes, a class should also define the `get` and `set` methods. The `get` method accepts one or two arguments. The first argument is an object of the appropriate class. If no second argument is given then the method should return a structure with all the properties of the class. If the optional second argument is given it should be a property name and the specified property should be retrieved.

```
function val = get (p, prop)

  if (nargin < 1)
    print_usage ();
  endif

  if (nargin == 1)
    val.poly = p.poly;
  else
    if (! ischar (prop))
      error ("@polynomial/get: PROPERTY must be a string");
    endif
```

```
    switch (prop)
      case "poly"
        val = p.poly;
      otherwise
        error ('@polynomial/get: invalid PROPERTY "%s"', prop);
    endswitch
  endif

endfunction
```

Similarly, the first argument to the `set` method should be an object and any additional arguments should be property/value pairs.

```
function pout = set (p, varargin)

  if (numel (varargin) < 2 || rem (numel (varargin), 2) != 0)
    error ("@polynomial/set: expecting PROPERTY/VALUE pairs");
  endif

  pout = p;
  while (numel (varargin) > 1)
    prop = varargin{1};
    val  = varargin{2};
    varargin(1:2) = [];
    if (! ischar (prop) || ! strcmp (prop, "poly"))
      error ("@polynomial/set: invalid PROPERTY for polynomial class");
    elseif (! (isreal (val) && isvector (val)))
      error ("@polynomial/set: VALUE must be a real vector");
    endif

    pout.poly = val(:).';  # force row vector
  endwhile

endfunction
```

Note that Octave does not implement pass by reference; Therefore, to modify an object requires an assignment statement using the return value from the `set` method.

```
p = set (p, "poly", [1, 0, 0, 0, 1]);
```

The `set` method makes use of the `subsasgn` method of the class, and therefore this method must also be defined. The `subsasgn` method is discussed more thoroughly in the next section (see Section 34.3 [Indexing Objects], page 978).

Finally, user classes can be considered to be a special type of a structure, and they can be saved to a file in the same manner as a structure. For example:

```
p = polynomial ([1, 0, 1]);
save userclass.mat p
clear p
load userclass.mat
```

All of the file formats supported by `save` and `load` are supported. In certain circumstances a user class might contain a field that it doesn't make sense to save, or a field that needs to be initialized before it is saved. This can be done with the `saveobj` method of the class.

`b = saveobj (a)`

Method of a class to manipulate an object prior to saving it to a file.

The function `saveobj` is called when the object *a* is saved using the `save` function. An example of the use of `saveobj` might be to remove fields of the object that don't make sense to be saved or it might be used to ensure that certain fields of the object are initialized before the object is saved. For example:

```
function b = saveobj (a)
  b = a;
  if (isempty (b.field))
     b.field = initfield (b);
  endif
endfunction
```

**See also:** [loadobj], page 978, [class], page 41.

`saveobj` is called just prior to saving the class to a file. Similarly, the `loadobj` method is called just after a class is loaded from a file, and can be used to ensure that any removed fields are reinserted into the user object.

`b = loadobj (a)`

Method of a class to manipulate an object after loading it from a file.

The function `loadobj` is called when the object *a* is loaded using the `load` function. An example of the use of `saveobj` might be to add fields to an object that don't make sense to be saved. For example:

```
function b = loadobj (a)
  b = a;
  b.addmissingfield = addfield (b);
endfunction
```

**See also:** [saveobj], page 978, [class], page 41.

## 34.3 Indexing Objects

### 34.3.1 Defining Indexing And Indexed Assignment

Objects can be indexed with parentheses or braces, either like `obj(idx)` or like `obj{idx}`, or even like `obj(idx).field`. However, it is up to the programmer to decide what this indexing actually means. In the case of the polynomial class `p(n)` might mean either the coefficient of the *n*-th power of the polynomial, or it might be the evaluation of the polynomial at *n*. The meaning of this subscripted referencing is determined by the `subsref` method.

`newval = subsref (val, idx)`

Perform the subscripted element selection operation on *val* according to the subscript specified by *idx*.

The subscript *idx* must be a structure array with fields 'type' and 'subs'. Valid values for 'type' are "()", "{}", and ".". The 'subs' field may be either ":" or a cell array of index values.

The following example shows how to extract the first two columns of a matrix

```
val = magic (3)
    ⇒ val = [ 8   1   6
              3   5   7
              4   9   2 ]
idx.type = "()";
idx.subs = {":", 1:2};
subsref (val, idx)
     ⇒ [ 8   1
         3   5
         4   9 ]
```

Note that this is the same as writing `val(:, 1:2)`.

If *idx* is an empty structure array with fields 'type' and 'subs', return *val*.

The keyword `end` cannot be used within `subsref` for indexing assignments.

**See also:** [subsasgn], page 980, [substruct], page 128.

For example, this class uses the convention that indexing with "()" evaluates the polynomial and indexing with "{}" returns the $n$-th coefficient (of the $n$-th power). The code for the `subsref` method looks like

```
function r = subsref (p, s)

  if (isempty (s))
    error ("@polynomial/subsref: missing index");
  endif

  switch (s(1).type)

    case "()"
      idx = s(1).subs;
      if (numel (idx) != 1)
        error ("@polynomial/subsref: need exactly one index");
      endif
      r = polyval (fliplr (p.poly), idx{1});

    case "{}"
      idx = s(1).subs;
      if (numel (idx) != 1)
        error ("@polynomial/subsref: need exactly one index");
      endif

      if (isnumeric (idx{1}))
        r = p.poly(idx{1}+1);
```

```
      else
        r = p.poly(idx{1});
      endif

    case "."
      fld = s.subs;
      if (! strcmp (fld, "poly"))
        error ('@polynomial/subsref: invalid property "%s"', fld);
      endif
      r = p.poly;

    otherwise
      error ("@polynomial/subsref: invalid subscript type");

  endswitch

  if (numel (s) > 1)
    r = subsref (r, s(2:end));
  endif

endfunction
```

The equivalent functionality for subscripted assignments uses the `subsasgn` method.

*newval* = **subsasgn** (*val*, *idx*, *rhs*)

Perform the subscripted assignment operation according to the subscript specified by *idx*.

The subscript *idx* must be a structure array with fields '`type`' and '`subs`'. Valid values for '`type`' are `"()"`, `"{}"`, and `"."`. The '`subs`' field may be either `":"` or a cell array of index values.

The following example shows how to set the two first columns of a 3-by-3 matrix to zero.

```
val = magic (3);
idx.type = "()";
idx.subs = {":", 1:2};
val = subsasgn (val, idx, 0)
     ⇒  [ 0   0   6
          0   0   7
          0   0   2 ]
```

Note that this is the same as writing `val(:, 1:2) = 0`.

If *idx* is an empty structure array with fields '`type`' and '`subs`', return *rhs*.

The keyword `end` cannot be used within `subsasgn` for indexing assignments.

**See also:** [subsref], page 978, [substruct], page 128, [optimize_subsasgn_calls], page 980.

*val* = **optimize_subsasgn_calls** ()
*old_val* = **optimize_subsasgn_calls** (*new_val*)

*old_val* = optimize_subsasgn_calls (*new_val*, "*local*")

Query or set the internal flag for `subsasgn` method call optimizations.

If true, Octave will attempt to eliminate the redundant copying when calling the `subsasgn` method of a user-defined class.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [subsasgn], page 980.

Note that the `subsref` and `subsasgn` methods always receive the whole index chain, while they usually handle only the first element. It is the responsibility of these methods to handle the rest of the chain (if needed), usually by forwarding it again to `subsref` or `subsasgn`.

*n* = numArgumentsFromSubscript (*obj*, *idx*, *unused*)

Override *nargout* for overloaded `subsref` method.

*obj* is the object for which the overloaded `subsref` method is called.

*idx* is a structure array with fields '`type`' and '`subs`'. See [subsref], page 978 for a description of that structure.

The third input argument *unused* is currently unused. It is always the empty matrix `[]`.

The function must return a scalar integer which will be passed as *nargout* to the overloaded `subsref` method

**See also:** [subsref], page 978, [substruct], page 128.

If you wish to use the `end` keyword in subscripted expressions of an object, then there must be an `end` method defined. For example, the `end` method for the polynomial class might look like

```
function r = end (obj, index_pos, num_indices)

  if (num_indices != 1)
    error ("polynomial object may only have one index");
  endif

  r = length (obj.poly) - 1;

endfunction
```

which is a fairly generic `end` method that has a behavior similar to the `end` keyword for Octave Array classes. An example using the polynomial class is then

```
p = polynomial ([1,2,3,4]);
p{end-1}
  ⇒ 3
```

Objects can also be used themselves as the index in a subscripted expression and this is controlled by the `subsindex` function.

*idx* = subsindex (*obj*)

Convert an object to an index vector.

When *obj* is a class object defined with a class constructor, then `subsindex` is the overloading method that allows the conversion of this class object to a valid indexing vector. It is important to note that `subsindex` must return a zero-based real integer vector of the class `"double"`. For example, if the class constructor were

```
function obj = myclass (a)
  obj = class (struct ("a", a), "myclass");
endfunction
```

then the `subsindex` function

```
function idx = subsindex (obj)
  idx = double (obj.a) - 1.0;
endfunction
```

could be used as follows

```
a = myclass (1:4);
b = 1:10;
b(a)
⇒ 1  2  3  4
```

**See also:** [class], page 41, [subsref], page 978, [subsasgn], page 980.

Finally, objects can be used like ranges by providing a `colon` method.

*r* = colon (*base*, *limit*)
*r* = colon (*base*, *increment*, *limit*)

Return the result of the colon expression corresponding to *base*, *limit*, and optionally, *increment*.

This function is equivalent to the operator syntax *base* : *limit* or *base* : *increment* : *limit*.

**See also:** [linspace], page 584.

### 34.3.2 Indexed Assignment Optimization

Octave's ubiquitous lazily-copied pass-by-value semantics implies a problem for performance of user-defined `subsasgn` methods. Imagine the following call to `subsasgn`

```
ss = substruct ("()", {1});
x = subsasgn (x, ss, 1);
```

where the corresponding method looking like this:

```
function x = subsasgn (x, ss, val)
  ...
  x.myfield (ss.subs{1}) = val;
endfunction
```

The problem is that on entry to the `subsasgn` method, `x` is still referenced from the caller's scope, which means that the method will first need to unshare (copy) `x` and `x.myfield` before performing the assignment. Upon completing the call, unless an error occurs, the result is immediately assigned to `x` in the caller's scope, so that the previous

value of `x.myfield` is forgotten. Hence, the Octave language implies a copy of N elements (N being the size of `x.myfield`), where modifying just a single element would actually suffice. In other words, a constant-time operation is degraded to linear-time one. This may be a real problem for user classes that intrinsically store large arrays.

To partially solve the problem Octave uses a special optimization for user-defined `subsasgn` methods coded as m-files. When the method gets called as a result of the built-in assignment syntax (not a direct `subsasgn` call as shown above), i.e., `x(1) = 1`, **AND** if the `subsasgn` method is declared with identical input and output arguments, as in the example above, then Octave will ignore the copy of `x` inside the caller's scope; therefore, any changes made to `x` during the method execution will directly affect the caller's copy as well. This allows, for instance, defining a polynomial class where modifying a single element takes constant time.

It is important to understand the implications that this optimization brings. Since no extra copy of `x` will exist in the caller's scope, it is *solely* the callee's responsibility to not leave `x` in an invalid state if an error occurs during the execution. Also, if the method partially changes `x` and then errors out, the changes *will* affect `x` in the caller's scope. Deleting or completely replacing `x` inside subsasgn will not do anything, however, only indexed assignments matter.

Since this optimization may change the way code works (especially if badly written), a function `optimize_subsasgn_calls` is provided to control it. This feature is enabled by default. Another way to avoid the optimization is to declare subsasgn methods with different output and input arguments like this:

```
function y = subsasgn (x, ss, val)
  ...
endfunction
```

## 34.4 Overloading Objects

### 34.4.1 Function Overloading

Any Octave function can be overloaded, and this allows an object-specific version of a function to be called as needed. A pertinent example for the polynomial class might be to overload the `polyval` function.

```
function [y, dy] = polyval (p, varargin)

  if (nargout > 1)
    [y, dy] = polyval (fliplr (p.poly), varargin{:});
  else
    y = polyval (fliplr (p.poly), varargin{:});
  endif

endfunction
```

This function just hands off the work to the normal Octave `polyval` function. Another interesting example of an overloaded function for the polynomial class is the `plot` function.

```
function h = plot (p, varargin)

  n = 128;
  rmax = max (abs (roots (p.poly)));
  x = [0 : (n - 1)] / (n - 1) * 2.2 * rmax - 1.1 * rmax;
  if (nargout > 0)
    h = plot (x, polyval (p, x), varargin{:});
  else
    plot (x, polyval (p, x), varargin{:});
  endif

endfunction
```

which allows polynomials to be plotted in the domain near the region of the roots of the polynomial.

Functions that are of particular interest for overloading are the class conversion functions such as `double`. Overloading these functions allows the `cast` function to work with a user class. It can also aid in the use of a class object with methods and functions from other classes since the object can be transformed to the requisite input form for the new function. An example `double` function for the polynomial class might look like

```
function a = double (p)
  a = p.poly;
endfunction
```

### 34.4.2 Operator Overloading

The following table shows, for each built-in numerical operation, the corresponding function name to use when providing an overloaded method for a user class.

| Operation | Method | Description |
|---|---|---|
| `a + b` | `plus (a, b)` | Binary addition |
| `a - b` | `minus (a, b)` | Binary subtraction |
| `+a` | `uplus (a)` | Unary addition |
| `-a` | `uminus (a)` | Unary subtraction |
| `a .* b` | `times (a, b)` | Element-wise multiplication |
| `a * b` | `mtimes (a, b)` | Matrix multiplication |
| `a ./ b` | `rdivide (a, b)` | Element-wise right division |
| `a / b` | `mrdivide (a, b)` | Matrix right division |
| `a .\ b` | `ldivide (a, b)` | Element-wise left division |
| `a \ b` | `mldivide (a, b)` | Matrix left division |
| `a .^ b` | `power (a, b)` | Element-wise power |
| `a ^ b` | `mpower (a, b)` | Matrix power |
| `a < b` | `lt (a, b)` | Less than |
| `a <= b` | `le (a, b)` | Less than or equal to |
| `a > b` | `gt (a, b)` | Greater than |
| `a >= b` | `ge (a, b)` | Greater than or equal to |
| `a == b` | `eq (a, b)` | Equal to |
| `a != b` | `ne (a, b)` | Not equal to |
| `a & b` | `and (a, b)` | Logical and |
| `a \| b` | `or (a, b)` | Logical or |
| `!a` | `not (a)` | Logical not |
| `a'` | `ctranspose (a)` | Complex conjugate transpose |
| `a.'` | `transpose (a)` | Transpose |
| `a:b` | `colon (a, b)` | Two element range |
| `a:b:c` | `colon (a, b, c)` | Three element range |
| `[a, b]` | `horzcat (a, b)` | Horizontal concatenation |
| `[a; b]` | `vertcat (a, b)` | Vertical concatenation |
| `a(`$s_1$`,...,`$s_n$`)` | `subsref (a, s)` | Subscripted reference |
| `a(`$s_1$`,...,`$s_n$`) = b` | `subsasgn (a, s, b)` | Subscripted assignment |
| `b(a)` | `subsindex (a)` | Convert object to index |
| `disp` | `disp (a)` | Object display |

Table 34.1: Available overloaded operators and their corresponding class method

An example `mtimes` method for the polynomial class might look like

```
function p = mtimes (a, b)
  p = polynomial (conv (double (a), double (b)));
endfunction
```

### 34.4.3 Precedence of Objects

Many functions and operators take two or more arguments and the situation can easily arise where these functions are called with objects of different classes. It is therefore necessary to determine the precedence of which method from which class to call when there are mixed objects given to a function or operator. To do this the `superiorto` and `inferiorto` functions can be used

`superiorto (`*`class_name`*`, ...)`

When called from a class constructor, mark the object currently constructed as having a higher precedence than *class_name*.

More that one such class can be specified in a single call. This function may *only* be called from a class constructor.

**See also:** [inferiorto], page 986.

`inferiorto (`*`class_name`*`, ...)`

When called from a class constructor, mark the object currently constructed as having a lower precedence than *class_name*.

More that one such class can be specified in a single call. This function may *only* be called from a class constructor.

**See also:** [superiorto], page 985.

With the polynomial class, consider the case

```
2 * polynomial ([1, 0, 1]);
```

that mixes an object of the class `"double"` with an object of the class `"polynomial"`. In this case the return type should be `"polynomial"` and so the `superiorto` function is used in the class constructor. In particular the polynomial class constructor would be modified to

```
## -*- texinfo -*-
## @deftypefn  {} {} polynomial ()
## @deftypefnx {} {} polynomial (@var{a})
## Create a polynomial object representing the polynomial
##
## @example
## a0 + a1 * x + a2 * x^2 + @dots{} + an * x^n
## @end example
##
## @noindent
## from a vector of coefficients [a0 a1 a2 @dots{} an].
## @end deftypefn

function p = polynomial (a)

  if (nargin == 0)
    p.poly = [0];
    p = class (p, "polynomial");
  else
    if (strcmp (class (a), "polynomial"))
      p = a;
    elseif (isreal (a) && isvector (a))
      p.poly = a(:).';  # force row vector
      p = class (p, "polynomial");
    else
      error ("@polynomial: A must be a real vector");
```

```
    endif
  endif

  superiorto ("double");

endfunction
```

Note that user classes *always* have higher precedence than built-in Octave types. Thus, marking the polynomial class higher than the `"double"` class is not actually necessary.

When confronted with two objects of equal precedence, Octave will use the method of the object that appears first in the list of arguments.

## 34.5 Inheritance and Aggregation

Using classes to build new classes is supported by Octave through the use of both inheritance and aggregation.

Class inheritance is provided by Octave using the `class` function in the class constructor. As in the case of the polynomial class, the Octave programmer will create a structure that contains the data fields required by the class, and then call the `class` function to indicate that an object is to be created from the structure. Creating a child of an existing object is done by creating an object of the parent class and providing that object as the third argument of the class function.

This is most easily demonstrated by example. Suppose the programmer needs a FIR filter, i.e., a filter with a numerator polynomial but a denominator of 1. In traditional Octave programming this would be performed as follows.

```
>> x = [some data vector];
>> n = [some coefficient vector];
>> y = filter (n, 1, x);
```

The equivalent behavior can be implemented as a class `@FIRfilter`. The constructor for this class is the file `FIRfilter.m` in the class directory `@FIRfilter`.

```
## -*- texinfo -*-
## @deftypefn  {} {} FIRfilter ()
## @deftypefnx {} {} FIRfilter (@var{p})
## Create a FIR filter with polynomial @var{p} as coefficient vector.
## @end deftypefn

function f = FIRfilter (p)

  if (nargin == 0)
    p = @polynomial ([1]);
  elseif (! isa (p, "polynomial"))
    error ("@FIRfilter: P must be a polynomial object");
  endif

  f.polynomial = [];
  f = class (f, "FIRfilter", p);
```

```
endfunction
```

As before, the leading comments provide documentation for the class constructor. This constructor is very similar to the polynomial class constructor, except that a polynomial object is passed as the third argument to the `class` function, telling Octave that the `FIRfilter` class will be derived from the polynomial class. The FIR filter class itself does not have any data fields, but it must provide a struct to the `class` function. Given that the `@polynomial` constructor will add an element named *polynomial* to the object struct, the `@FIRfilter` just initializes a struct with a dummy field *polynomial* which will later be overwritten.

Note that the sample code always provides for the case in which no arguments are supplied. This is important because Octave will call a constructor with no arguments when loading objects from saved files in order to determine the inheritance structure.

A class may be a child of more than one class (see [class], page 41), and inheritance may be nested. There is no limitation to the number of parents or the level of nesting other than memory or other physical issues.

For the `FIRfilter` class, more control about the object display is desired. Therefore, the `display` method rather than the `disp` method is overloaded (see Section 34.2 [Class Methods], page 975). A simple example might be

```
function display (f)
  printf ("%s.polynomial", inputname (1));
  disp (f.polynomial);
endfunction
```

Note that the `FIRfilter`'s display method relies on the `disp` method from the `polynomial` class to actually display the filter coefficients. Furthermore, note that in the `display` method it makes sense to start the method with the line `printf ("%s =", inputname (1))` to be consistent with the rest of Octave which prints the variable name to be displayed followed by the value. In general it is not recommended to overload the `display` function.

`display (`*obj*`)`

Display the contents of the object *obj* prepended by its name.

The Octave interpreter calls the `display` function whenever it needs to present a class on-screen. Typically, this would be a statement which does not end in a semicolon to suppress output. For example:

```
myclass (...)
```

Or:

```
myobj = myclass (...)
```

In general, user-defined classes should overload the `disp` method to avoid the default output:

```
myobj = myclass (...)
  ⇒ myobj =

    <class myclass>
```

When overloading the `display` method instead, one has to take care of properly displaying the object's name. This can be done by using the `inputname` function.

**See also:** [disp], page 287, [class], page 41, [subsref], page 978, [subsasgn], page 980.

Once a constructor and display method exist, it is possible to create an instance of the class. It is also possible to check the class type and examine the underlying structure.

```
octave:1> f = FIRfilter (polynomial ([1 1 1]/3))
f.polynomial = 0.33333 + 0.33333 * X + 0.33333 * X ^ 2
octave:2> class (f)
ans = FIRfilter
octave:3> isa (f, "FIRfilter")
ans =  1
octave:4> isa (f, "polynomial")
ans =  1
octave:5> struct (f)
ans =

  scalar structure containing the fields:

polynomial = 0.33333 + 0.33333 * X + 0.33333 * X ^ 2
```

The only thing remaining to make this class usable is a method for processing data. But before that, it is usually desirable to also have a way of changing the data stored in a class. Since the fields in the underlying struct are private by default, it is necessary to provide a mechanism to access the fields. The `subsref` method may be used for both tasks.

```
function r = subsref (f, x)

  switch (x.type)

    case "()"
      n = f.polynomial;
      r = filter (n.poly, 1, x.subs{1});

    case "."
      fld = x.subs;
      if (! strcmp (fld, "polynomial"))
        error ('@FIRfilter/subsref: invalid property "%s"', fld);
      endif
      r = f.polynomial;

    otherwise
      error ("@FIRfilter/subsref: invalid subscript type for FIR filter");

  endswitch

endfunction
```

The `"()"` case allows us to filter data using the polynomial provided to the constructor.

```
octave:2> f = FIRfilter (polynomial ([1 1 1]/3));
octave:3> x = ones (5,1);
octave:4> y = f(x)
y =

   0.33333
   0.66667
   1.00000
   1.00000
   1.00000
```

The "." case allows us to view the contents of the polynomial field.

```
octave:1> f = FIRfilter (polynomial ([1 1 1]/3));
octave:2> f.polynomial
ans = 0.33333 + 0.33333 * X + 0.33333 * X ^ 2
```

In order to change the contents of the object a `subsasgn` method is needed. For example, the following code makes the polynomial field publicly writable

```
function fout = subsasgn (f, index, val)

  switch (index.type)
    case "."
      fld = index.subs;
      if (! strcmp (fld, "polynomial"))
        error ('@FIRfilter/subsasgn: invalid property "%s"', fld);
      endif
      fout = f;
      fout.polynomial = val;

    otherwise
      error ("@FIRfilter/subsasgn: Invalid index type")
  endswitch

endfunction
```

so that

```
octave:1> f = FIRfilter ();
octave:2> f.polynomial = polynomial ([1 2 3])
f.polynomial = 1 + 2 * X + 3 * X ^ 2
```

Defining the FIRfilter class as a child of the polynomial class implies that a FIRfilter object may be used any place that a polynomial object may be used. This is not a normal use of a filter. It may be a more sensible design approach to use aggregation rather than inheritance. In this case, the polynomial is simply a field in the class structure. A class constructor for the aggregation case might be

```
## -*- texinfo -*-
## @deftypefn  {} {} FIRfilter ()
## @deftypefnx {} {} FIRfilter (@var{p})
## Create a FIR filter with polynomial @var{p} as coefficient vector.
```

```
## @end deftypefn

function f = FIRfilter (p)

  if (nargin == 0)
    f.polynomial = @polynomial ([1]);
  else
    if (! isa (p, "polynomial"))
      error ("@FIRfilter: P must be a polynomial object");
    endif

    f.polynomial = p;
  endif

  f = class (f, "FIRfilter");

endfunction
```

For this example only the constructor needs changing, and all other class methods stay the same.

## 34.6 `classdef` Classes

Since version 4.0, Octave has limited support for `classdef` classes. In contrast to the aforementioned classes, called *old style classes* in this section, `classdef` classes can be defined within a single m-file. Other innovations of `classdef` classes are:

- **access rights** for properties and methods,
- **static methods**, i.e., methods that are independent of an object, and
- the distinction between **value and handle classes**.

Several features have to be added in future versions of Octave to be fully compatible to MATLAB. An overview of what is missing can be found at https://wiki.octave.org/Classdef.

### 34.6.1 Creating a `classdef` Class

A very basic `classdef` value class (see Section 34.6.5 [Value Classes vs. Handle Classes], page 996) is defined by:

```
classdef some_class
  properties
  endproperties

  methods
  endmethods
endclassdef
```

In contrast to old style classes, the `properties-endproperties` block as well as the `methods-endmethods` block can be used to define properties and methods of the class. Because both blocks are empty, they can be omitted in this particular case.

For simplicity, a more advanced implementation of a `classdef` class is shown using the `polynomial` example again (see Section 34.1 [Creating a Class], page 973):

```
classdef polynomial2
  properties
    poly = 0;
  endproperties

  methods
    function p = polynomial2 (a)
      if (nargin == 1)
        if (isa (a, "polynomial2"))
          p.poly = a.poly;
        elseif (isreal (a) && isvector (a))
          p.poly = a(:).';  # force row vector
        else
          error ("polynomial2: A must be a real vector");
        endif
      endif
    endfunction

    function disp (p)
      a = p.poly;
      first = true;
      for i = 1 : length (a);
        if (a(i) != 0)
          if (first)
            first = false;
          elseif (a(i) > 0 || isnan (a(i)))
            printf (" +");
          endif
          if (a(i) < 0)
            printf (" -");
          endif
          if (i == 1)
            printf (" %.5g", abs (a(i)));
          elseif (abs (a(i)) != 1)
            printf (" %.5g *", abs (a(i)));
          endif
          if (i > 1)
            printf (" X");
          endif
          if (i > 2)
            printf (" ^ %d", i - 1);
          endif
        endif
      endfor
```

```
        if (first)
          printf (" 0");
        endif
        printf ("\n");
      endfunction
    endmethods
  endclassdef
```

An object of class `polynomial2` is created by calling the class constructor:

```
>> p = polynomial2 ([1, 0, 1])
⇒ p =

 1 + X ^ 2
```

### 34.6.2 Properties

All class properties must be defined within `properties` blocks. The definition of a default value for a property is optional and can be omitted. The default initial value for each class property is `[]`.

A `properties` block can have additional attributes to specify access rights or to define constants:

```
classdef some_class
  properties (Access = mode)
    prop1
  endproperties

  properties (SetAccess = mode, GetAccess = mode)
    prop2
  endproperties

  properties (Constant = true)
    prop3 = pi ()
  endproperties

  properties
    prop4 = 1337
  endproperties
endclassdef
```

where *mode* can be one of:

`public` The properties can be accessed from everywhere.

`private` The properties can only be accessed from class methods. Subclasses of that class cannot access them.

`protected`
The properties can only be accessed from class methods and from subclasses of that class.

When creating an object of `some_class`, *prop1* has the default value `[]` and reading from and writing to *prop1* is defined by a single *mode*. For *prop2* the read and write access can be set differently. Finally, *prop3* is a constant property which can only be initialized once within the `properties` block.

By default, in the example *prop4*, properties are not constant and have public read and write access.

`properties (`*obj*`)`
`properties (`*class_name*`)`
*proplist* `= properties (...)`

Display or return the public properties for the classdef object *obj* or the named class *class_name*.

If an output value is requested, return the list of property names in a cell array.

Programming Note: Property names are returned if the `GetAccess` attribute is public and if the `Hidden` attribute is false.

**See also:** [methods], page 974.

### 34.6.3 Methods

All class methods must be defined within `methods` blocks. An exception to this rule is described at the end of this subsection. Those `methods` blocks can have additional attributes specifying the access rights or whether the methods are static, i.e., methods that can be called without creating an object of that class.

```
classdef some_class
  methods
    function obj = some_class ()
      disp ("New instance created.");
    endfunction

    function disp (obj)
      disp ("Here is some_class.");
    endfunction
  endmethods

  methods (Access = mode)
    function r = func (obj, r)
      r = 2 * r;
    endfunction
  endmethods

  methods (Static = true)
    function c = circumference (radius)
      c = 2 * pi () .* radius;
    endfunction
  endmethods
endclassdef
```

The constructor of the class is declared in the `methods` block and must have the same name as the class and exactly one output argument which is an object of its class.

It is also possible to overload built-in or inherited methods, like the `disp` function in the example above to tell Octave how objects of `some_class` should be displayed (see Section 34.2 [Class Methods], page 975).

In general, the first argument in a method definition is always the object that it is called from. Class methods can either be called by passing the object as the first argument to that method or by calling the object followed by a dot (`"."`) and the method's name with subsequent arguments:

```
>> obj = some_class ();
New instance created.
>> disp (obj);   # both are
>> obj.disp ();  # equal
```

In `some_class`, the method `func` is defined within a `methods` block setting the `Access` attribute to *mode*, which is one of:

`public` The methods can be accessed from everywhere.

`private` The methods can only be accessed from other class methods. Subclasses of that class cannot access them.

`protected`
The methods can only be accessed from other class methods and from subclasses of that class.

The default access for methods is `public`.

Finally, the method `circumference` is defined in a static `methods` block and can be used without creating an object of `some_class`. This is useful for methods, that do not depend on any class properties. The class name and the name of the static method, separated by a dot (`"."`), call this static method. In contrast to non-static methods, the object is not passed as first argument even if called using an object of `some_class`.

```
>> some_class.circumference (3)
⇒ ans =  18.850
>> obj = some_class ();
New instance created.
>> obj.circumference (3)
⇒ ans =  18.850
```

Additionally, class methods can be defined as functions in a folder of the same name as the class prepended with the '`@`' symbol (see Section 34.1 [Creating a Class], page 973). The main `classdef` file has to be stored in this class folder as well.

### 34.6.4 Inheritance

Classes can inherit from other classes. In this case all properties and methods of the superclass are inherited to the subclass, considering their access rights. Use this syntax to inherit from `superclass`:

```
classdef subclass < superclass
  ...
endclassdef
```

### 34.6.5 Value Classes vs. Handle Classes

There are two intrinsically different types of `classdef` classes, whose major difference is the behavior regarding variable assignment. The first type are **value classes**:

```
classdef value_class
  properties
    prop1
  endproperties

  methods
    function obj = set_prop1 (obj, val)
      obj.prop1 = val;
    endfunction
  endmethods
endclassdef
```

Assigning an object of that class to another variable essentially creates a new object:

```
>> a = value_class ();
>> a.prop1 = 1;
>> b = a;
>> b.prop1 = 2;
>> b.prop1
⇒ ans =  2
>> a.prop1
⇒ ans =  1
```

But that also means that you might have to assign the output of a method that changes properties back to the object manually:

```
>> a = value_class ();
>> a.prop1 = 1;
>> a.set_prop1 (3);
⇒ ans =

<object value_class>

>> ans.prop1
⇒ ans =  3
>> a.prop1
⇒ ans =  1
```

The second type are **handle classes**. Those classes have to be derived from the abstract `handle` class:

```
classdef handle_class < handle
  properties
    prop1
  endproperties

  methods
    function set_prop1 (obj, val)
      obj.prop1 = val;
    endfunction
  endmethods
endclassdef
```

In the following example, the variables `a` and `b` refer to the very same object of class `handle_class`:

```
>> a = handle_class ();
>> a.prop1 = 1;
>> b = a;
>> b.prop1 = 2;
>> b.prop1
⇒ ans =  2
>> a.prop1
⇒ ans =  2
```

Object properties that are modified by a method of an handle class are changed persistently:

```
>> a.set_prop1 (3);
>> a.prop1
⇒ ans =  3
```

# 35 GUI Development

Octave is principally a batch or command-line language. However, it does offer some features for constructing graphical interfaces that interact with users.

The GUI elements available are I/O dialogs, a progress bar, and UI elements for plot windows. For example, rather than hardcoding a filename for output results a script can open a dialog box and allow the user to choose a file. Similarly, if a calculation is expected to take a long time a script can display a progress bar. The various UI elements can be used to fully customize the plot window with menubars, toolbars, context menus, pushbuttons, sliders, etc.

Several utility functions make it possible to store private data for use with a GUI which will not pollute the user's variable space.

Finally, a program written in Octave might want to have long term storage of preferences or state variables. This can be done with user-defined preferences.

## 35.1 I/O Dialogs

Simple dialog menus are available for choosing directories or files. They return a string variable which can then be used with any command requiring a filename.

*dirname* = uigetdir ()
*dirname* = uigetdir (*init_path*)
*dirname* = uigetdir (*init_path*, *dialog_name*)

Open a GUI dialog for selecting a directory.

If *init_path* is not given the current working directory is used.

*dialog_name* may be used to customize the dialog title.

The output *dirname* is a character string with the name of the selected directory. However, if the 'Cancel' button is clicked the output is of type double with the value 0.

**See also:** [uigetfile], page 999, [uiputfile], page 1000.

[*fname*, *fpath*, *fltidx*] = uigetfile ()
[...] = uigetfile (*flt*)
[...] = uigetfile (*flt*, *dialog_name*)
[...] = uigetfile (*flt*, *dialog_name*, *default_file*)
[...] = uigetfile (..., "*MultiSelect*", *mode*)

Open a GUI dialog for selecting a file and return the filename *fname*, the path to this file *fpath*, and the filter index *fltidx*.

*flt* contains a (list of) file filter string(s) in one of the following formats:

"/path/to/filename.ext"
: If a filename is given then the file extension is extracted and used as filter. In addition, the path is selected as current path in the dialog and the filename is selected as default file. Example: `uigetfile ("myfcn.m")`

A single file extension "*.ext"
: Example: `uigetfile ("*.ext")`

A 2-column cell array
: containing a file extension in the first column and a brief description in the second column. Example: `uigetfile ({"*.ext", "My Description";"*.xyz", "XYZ-Format"})`

  The filter string can also contain a semicolon separated list of filter extensions. Example: `uigetfile ({"*.gif;*.png;*.jpg", "Supported Picture Formats"})`

A directory name or path name
: If the folder name of path name contains a trailing file separator, the contents of that folder will be displayed. If no trailing file separator is present the parent directory is listed. The substring to the right of the rightmost file separator (if any) will be interpreted as a file or directory name and if that file or directory exists it will be highlighted. If the path name or directory name is entirely or partly nonexistent, the current working directory will be displayed. No filter will be active.

*dialog_name* can be used to customize the dialog title.

If *default_file* is given then it will be selected in the GUI dialog. If, in addition, a path is given it is also used as current path.

Two or more files can be selected when setting the `"MultiSelect"` key to `"on"`. In that case, *fname* is a cell array containing the files.

The outputs *fname* and *fpath* are strings returning the chosen name and path, respectively. However, if the '`Cancel`' button is clicked the outputs are of type double with a value of `0`. *fltidx* is the index in the list of filter extensions *flt* that was selected.

**See also:** [uiputfile], page 1000, [uigetdir], page 999.

`[`*fname*`, `*fpath*`, `*fltidx*`] = uiputfile ()`
`[`*fname*`, `*fpath*`, `*fltidx*`] = uiputfile (`*flt*`)`
`[`*fname*`, `*fpath*`, `*fltidx*`] = uiputfile (`*flt*`, `*dialog_name*`)`
`[`*fname*`, `*fpath*`, `*fltidx*`] = uiputfile (`*flt*`, `*dialog_name*`, `*default_file*`)`

Open a GUI dialog for selecting a file.

*flt* contains a (list of) file filter string(s) in one of the following formats:

`"/path/to/filename.ext"`
: If a filename is given the file extension is extracted and used as filter. In addition the path is selected as current path in the dialog and the filename is selected as default file. Example: `uiputfile ("myfcn.m")`

`"*.ext"` A single file extension. Example: `uiputfile ("*.ext")`

`{"*.ext", "My Description"}`
: A 2-column cell array containing the file extension in the 1st column and a brief description in the 2nd column. Example: `uiputfile ({"*.ext","My Description";"*.xyz", "XYZ-Format"})`

The filter string can also contain a semicolon separated list of filter extensions. Example: `uiputfile ({"*.gif;*.png;*.jpg", "Supported Picture Formats"})`

*dialog_name* can be used to customize the dialog title. If *default_file* is given it is preselected in the GUI dialog. If, in addition, a path is given it is also used as current path.

*fname* and *fpath* return the chosen name and path, respectively. *fltidx* is the index in the list of filter extensions *flt* that was selected.

**See also:** [uigetfile], page 999, [uigetdir], page 999.

Additionally, there are dialog boxes for displaying help messages, warnings, or errors, and for getting text input from the user.

`errordlg ()`
`errordlg (`*`msg`*`)`
`errordlg (`*`msg`*`, `*`title`*`)`
`errordlg (`*`msg`*`, `*`title`*`, `*`opt`*`)`
*h* `= errordlg (...)`

Display an error dialog box with error message *msg* and caption *title*.

The default error message is `"This is the default error string."` and the default caption is `"Error Dialog"`.

The error message may have multiple lines separated by newline characters (`"\n"`), or it may be a cellstr array with one element for each line.

The third optional argument *opt* controls the behavior of the dialog. For details, see [msgbox], page 1003.

The return value *h* is a handle to the figure object used for building the dialog.

Examples:

```
errordlg ("Some fancy error occurred.");
errordlg ("Some fancy error\nwith two lines.");
errordlg ({"Some fancy error", "with two lines."});
errordlg ("Some fancy error occurred.", "Fancy caption");
```

**See also:** [helpdlg], page 1001, [warndlg], page 1005, [msgbox], page 1003, [inputdlg], page 1002, [listdlg], page 1002, [questdlg], page 1004.

`helpdlg ()`
`helpdlg (`*`msg`*`)`
`helpdlg (`*`msg`*`, `*`title`*`)`
*h* `= helpdlg (...)`

Display a help dialog box with help message *msg* and caption *title*.

The default help message is `"This is the default help string."` and the default caption is `"Help Dialog"`.

The help message may have multiple lines separated by newline characters (`"\n"`), or it may be a cellstr array with one element for each line.

The return value *h* is a handle to the figure object used for building the dialog.

Examples:

```
helpdlg ("Some helpful text for the user.");
helpdlg ("Some helpful text\nwith two lines.");
helpdlg ({"Some helpful text", "with two lines."});
helpdlg ("Some helpful text for the user.", "Fancy caption");
```

**See also:** [errordlg], page 1001, [warndlg], page 1005, [msgbox], page 1003, [inputdlg], page 1002, [listdlg], page 1002, [questdlg], page 1004.

*cstr* = inputdlg (*prompt*)
*cstr* = inputdlg (*prompt*, *title*)
*cstr* = inputdlg (*prompt*, *title*, *rowscols*)
*cstr* = inputdlg (*prompt*, *title*, *rowscols*, *defaults*)
*cstr* = inputdlg (*prompt*, *title*, *rowscols*, *defaults*, *options*)

Return user input from a multi-textfield dialog box.

Inputs:

*prompt* A cell array with strings labeling each text field. This input is required.

*title* String to use for the caption of the dialog. The default is `"Input Dialog"`.

*rowscols* Specifies the size of the text fields and can take four forms:

1. a scalar value which defines the number of rows used for each text field.
2. a row or column vector with the number of elements matching the number of prompts. Each element defines the number of rows used for the corresponding text field.
3. a 1x2 row vector which defines the number of rows and the number of columns used for all text fields.
4. an Mx2 matrix which defines the individual number of rows and columns used for each text field. In the matrix, each row describes a single text field. The first column specifies the number of input rows to use and the second column specifies the text field width.

*defaults* A list of default values to place in each text field. It must be a cell array of strings with the same size as *prompt*.

*options* Not supported, only for MATLAB compatibility.

Output: The output is a cell array of strings with each element being the text that was entered by the user. An empty cell array is returned if the dialog was closed by the Cancel button.

Example:

```
prompt = {"Width", "Height", "Depth"};
defaults = {"1.10", "2.20", "3.30"};
rowscols = [1,10; 2,20; 3,30];
dims = inputdlg (prompt, "Enter Box Dimensions", ...
                 rowscols, defaults);
```

**See also:** [errordlg], page 1001, [helpdlg], page 1001, [listdlg], page 1002, [msgbox], page 1003, [questdlg], page 1004, [warndlg], page 1005.

[*sel*, *ok*] = listdlg (*key*, *value*, ...)

Return user inputs from a list dialog box in a vector of selection indices (*sel*) and a flag indicating how the user closed the dialog box (*ok*).

The indices in *sel* are 1-based.

The value of *ok* is 1 if the user closed the box with the OK button, otherwise it is 0 and *sel* is empty.

Input arguments are specified in form of *key*, *value* pairs. The `"ListString"` argument pair **must** be specified.

Valid *key* and *value* pairs are:

`"ListString"`
: a cell array of strings specifying the items to list in the dialog.

`"SelectionMode"`
: can be either `"Single"` (only one item may be selected at a time) or `"Multiple"` (default).

`"ListSize"`
: a two-element vector `[`*`width`*`, `*`height`*`]` specifying the size of the list field in pixels. The default is [160, 300].

`"InitialValue"`
: a vector containing 1-based indices of elements which will be pre-selected when the list dialog is first displayed. The default is 1 (first item).

`"Name"`
: a string to be used as the dialog caption. Default is `""`.

`"PromptString"`
: a cell array of strings to be displayed above the list of items. Default is {}.

`"OKString"`
: a string used to label the OK button. Default is `"OK"`.

`"CancelString"`
: a string used to label the Cancel button. Default is `"Cancel"`.

Example:

```
my_options = {"An item", "another", "yet another"};
[sel, ok] = listdlg ("ListString", my_options,
                     "SelectionMode", "Multiple");
if (ok == 1)
  disp ("You selected:");
  for i = 1:numel (sel)
    disp (sprintf ("\t%s", my_options{sel(i)}));
  endfor
else
  disp ("You cancelled.");
endif
```

**See also:** [menu], page 293, [errordlg], page 1001, [helpdlg], page 1001, [inputdlg], page 1002, [msgbox], page 1003, [questdlg], page 1004, [warndlg], page 1005.

`h = msgbox (`*`msg`*`)`
`h = msgbox (`*`msg`*`, `*`title`*`)`
`h = msgbox (`*`msg`*`, `*`title`*`, `*`icon`*`)`

*h* = msgbox (*msg*, *title*, "*custom*", *cdata*)
*h* = msgbox (*msg*, *title*, "*custom*", *cdata*, *colormap*)
*h* = msgbox (*...*, *opt*)

Display *msg* using a message dialog box.

The message may have multiple lines separated by newline characters ("\n"), or it may be a cellstr array with one element for each line.

The optional input *title* (character string) can be used to decorate the dialog caption.

The optional argument *icon* selects a dialog icon. It can be one of "none" (default), "error", "help", "warn", or "custom". The latter must be followed by an image array *cdata*, and for indexed images the associated colormap.

The final optional argument *opt* controls the behavior of the dialog. If *opt* is a string, it may be one of

"non-modal" (default)
: The dialog is normal.

"modal"
: If any dialogs already exist with the same title, the most recent is reused and all others are closed. The dialog is displayed "modal" which means it prevents users from interacting with any other GUI element until the dialog has been closed.

"replace"
: If any dialogs already exist with the same title, the most recent is reused and all others are closed. The resulting dialog is set "non-modal".

If *opt* is a structure, it must contain fields "WindowStyle" and "Interpreter":

"WindowStyle"
: The value must be "non-modal", "modal", or "replace". See above.

"Interpreter"
: Controls the "interpreter" property of the text object used for displaying the message. The value must be "tex" (default), "none", or "latex".

The return value *h* is a handle to the figure object used for building the dialog.

Examples:

```
msgbox ("Some message for the user.");
msgbox ("Some message\nwith two lines.");
msgbox ({"Some message", "with two lines."});
msgbox ("Some message for the user.", "Fancy caption");

## A message dialog box with error icon
msgbox ("Some message for the user.", "Fancy caption", "error");
```

**See also:** [errordlg], page 1001, [helpdlg], page 1001, [inputdlg], page 1002, [listdlg], page 1002, [questdlg], page 1004, [warndlg], page 1005.

*btn* = questdlg (*msg*)
*btn* = questdlg (*msg*, *title*)
*btn* = questdlg (*msg*, *title*, *default*)

`btn = questdlg (msg, title, btn1, btn2, default)`
`btn = questdlg (msg, title, btn1, btn2, btn3, default)`

Display *msg* using a question dialog box and return the caption of the activated button.

The message may have multiple lines separated by newline characters (`"\n"`), or it may be a cellstr array with one element for each line.

The optional *title* (character string) can be used to specify the dialog caption. It defaults to `"Question Dialog"`.

The dialog may contain two or three buttons which will all close the dialog.

The string *default* identifies the default button, which is activated by pressing the `ENTER` key. It must match one of the strings given in *btn1*, *btn2*, or *btn3*.

If only *msg* and *title* are specified, three buttons with the default captions `"Yes"`, `"No"`, and `"Cancel"` are used.

If only two button captions, *btn1* and *btn2*, are specified the dialog will have only these two buttons.

Examples:

```
btn = questdlg ("Close Octave?", "Some fancy title", ...
                "Yes", "No", "No");
if (strcmp (btn, "Yes"))
  exit ();
endif
```

**See also:** [errordlg], page 1001, [helpdlg], page 1001, [inputdlg], page 1002, [listdlg], page 1002, [msgbox], page 1003, [warndlg], page 1005.

`warndlg ()`
`warndlg (msg)`
`warndlg (msg, title)`
`warndlg (msg, title, opt)`
`h = warndlg (...)`

Display a warning dialog box with warning message *msg* and caption *title*.

The default warning message is `"This is the default warning string."` and the default caption is `"Warning Dialog"`.

The warning message may have multiple lines separated by newline characters (`"\n"`), or it may be a cellstr array with one element for each line.

The third optional argument *opt* controls the behavior of the dialog. For details, see [msgbox], page 1003.

The return value *h* is a handle to the figure object used for building the dialog.

Examples:

```
warndlg ("Some warning text for the user.");
warndlg ("Some warning text\nwith two lines.");
warndlg ({"Some warning text", "with two lines."});
warndlg ("Some warning text for the user.", "Fancy caption");
```

**See also:** [errordlg], page 1001, [helpdlg], page 1001, [msgbox], page 1003, [inputdlg], page 1002, [listdlg], page 1002, [questdlg], page 1004.

```
  uisetfont ()
  uisetfont (h)
  uisetfont (fontstruct)
  uisetfont (..., title)
fontstruct = uisetfont (...)
```

Open a font selection dialog.

If the first argument is a handle to a text, axes, or uicontrol object, pressing the OK button will change the font properties of the object.

The first argument may also be a structure with fields `FontName`, `FontWeight`, `FontAngle`, `FontUnits`, and `FontSize`, indicating the initially selected font.

The title of the dialog window can be specified by using the last argument *title*.

If an output argument *fontstruct* is requested, the selected font structure is returned. Otherwise, the font information is displayed onscreen.

Programming Note: On systems that don't use FontConfig natively (all but Linux), the font cache is built when Octave is installed. You will need to run `system ("fc-cache -fv")` manually after installing new fonts.

**See also:** [listfonts], page 1018, [text], page 417, [axes], page 450, [uicontrol], page 1009.

For creating new dialog types, there is a dialog function.

```
h = dialog ()
h = dialog ("property", value, ...)
```

Create an empty modal dialog window to which other uicontrols can be added.

The dialog box is a figure object with properties as recommended for a dialog box.

The default properties differing from a figure are:

buttondownfcn
: `if isempty (allchild(gcbf)), close (gcbf), endif`

colormap
: []

color
: defaultuicontrolbackgroundcolor

dockcontrols
: off

handlevisibility
: callback

integerhandle
: off

inverthardcopy
: off

menubar
: none

numbertitle
: off

paperpositionmode
: auto

resize
: off

windowstyle
: modal

Multiple property-value pairs may be specified for the dialog object, but they must appear in pairs. The full list of properties is documented at Section 15.3.3.2 [Figure Properties], page 462.

The return value *h* is a graphics handle to the created figure.

Example:

```
## create an empty dialog window titled "Dialog Example"
h = dialog ("name", "Dialog Example");

## create a button (default style)
b = uicontrol (h, "string", "OK",
                  "position", [10 10 150 40],
                  "callback", "delete (gcf)");

## wait for dialog to resume or close
uiwait (h);
```

**See also:** [errordlg], page 1001, [msgbox], page 1003, [questdlg], page 1004, [warndlg], page 1005, [figure], page 425, [uiwait], page 1020.

## 35.2 Progress Bar

*h* = waitbar (*frac*)
*h* = waitbar (*frac*, *msg*)
*h* = waitbar (. . ., "*createcancelbtn*", *fcn*, . . .)
*h* = waitbar (. . ., *prop*, *val*, . . .)
waitbar (*frac*)
waitbar (*frac*, *h*)
waitbar (*frac*, *h*, *msg*)

Return a handle *h* to a new progress indicator ("waitbar") object.

The waitbar is filled to fraction *frac* which must be in the range [0, 1].

The optional message *msg* is centered and displayed above the waitbar.

A cancel button can be added to the bottom of the waitbar using the "createcancelbtn" property of waitbar figures. The action to be executed when the user presses the button is specified using a string or function handle *fcn*.

The appearance of the waitbar figure window can be configured by passing *prop*/*val* pairs to the function. The full list of properties is documented at Section 15.3.3.2 [Figure Properties], page 462.

When called with a single input the current waitbar, if it exists, is updated to the new value *frac*. If there are multiple outstanding waitbars they can be updated individually by passing the handle *h* of the specific waitbar to modify.

**See also:** [delete], page 430.

## 35.3 UI Elements

The ui* series of functions work best with the `qt` graphics toolkit, although some functionality is available with the `fltk` toolkit. There is no support for the `gnuplot` toolkit.

*h* = uifigure ()
*h* = uifigure ("*property*", *value*, ...)

Create a new figure window for applications.

Multiple property-value pairs may be specified for the figure object, but they must occur in pairs.

The return value *h* is a graphics handle to the created figure object.

Programming Note: The full list of properties is documented at Section 15.3.3.2 [Figure Properties], page 462. This function differs from `figure` in that the created figure is optimized for application development, rather than plotting. This means features such as menubars and toolbars are turned off. This is not perfect replacement for the MATLAB uifigure object as the properties `"AutoResizeChildren"`, `"Icon"`, and `"Scrollable"` are not implemented.

**See also:** [uipanel], page 1008, [uibuttongroup], page 1009.

*hui* = uipanel ()
*hui* = uipanel (*property*, *value*, ...)
*hui* = uipanel (*parent*)
*hui* = uipanel (*parent*, *property*, *value*, ...)

Create a uipanel object.

uipanels are used as containers to group other uicontrol objects.

If *parent* is omitted then a uipanel for the current figure is created. If no figure is available, a new figure is created first.

If *parent* is given then a uipanel relative to *parent* is created.

Any provided property value pairs will override the default values of the created uipanel object.

The full list of properties is documented at Section 15.3.3.15 [Uipanel Properties], page 521.

The optional return value *hui* is a graphics handle to the created uipanel object.

Examples:

```
## create figure and panel on it
f = figure;
p = uipanel ("title", "Panel Title", "position", [.25 .25 .5 .5]);

## add two buttons to the panel
b1 = uicontrol ("parent", p, "string", "A Button", ...
                "position", [18 10 150 36]);
b2 = uicontrol ("parent", p, "string", "Another Button", ...
                "position",[18 60 150 36]);
```

**See also:** [figure], page 425, [uicontrol], page 1009.

*hui* = uibuttongroup ()
*hui* = uibuttongroup (*property*, *value*, ...)
*hui* = uibuttongroup (*parent*)
*hui* = uibuttongroup (*parent*, *property*, *value*, ...)
uibuttongroup (*h*)

Create a uibuttongroup object and return a handle to it.

A uibuttongroup is used to group uicontrol objects.

If *parent* is omitted then a uibuttongroup for the current figure is created. If no figure is available, a new figure is created first.

If *parent* is given then a uibuttongroup relative to *parent* is created.

Any provided property value pairs will override the default values of the created uibuttongroup object.

The full list of properties is documented at Section 15.3.3.13 [Uibuttongroup Properties], page 516.

Examples:

```
## Create figure and panel on it
f = figure;
## Create a button group
gp = uibuttongroup (f, "Position", [ 0 0.5 1 1])
## Create a buttons in the group
b1 = uicontrol (gp, "style", "radiobutton", ...
                "string", "Choice 1", ...
                "Position", [ 10 150 100 50 ]);
b2 = uicontrol (gp, "style", "radiobutton", ...
                "string", "Choice 2", ...
                "Position", [ 10 50 100 30 ]);
## Create a button not in the group
b3 = uicontrol (f, "style", "radiobutton", ...
                "string", "Not in the group", ...
                "Position", [ 10 50 100 50 ]);
```

When called with a single argument *h* which is a handle to an existing uibuttongroup object, switch the focus to the specified uibuttongroup. This functionality is not currently implemented.

**See also:** [figure], page 425, [uipanel], page 1008.

*hui* = uicontrol ()
*hui* = uicontrol (*property*, *value*, ...)
*hui* = uicontrol (*parent*)
*hui* = uicontrol (*parent*, *property*, *value*, ...)
uicontrol (*h*)

Create a uicontrol object and return a handle to it.

A uicontrol object is used to create simple interactive controls such as push buttons, checkboxes, edit and list controls.

If *parent* is omitted then a uicontrol for the current figure is created. If no figure is available, a new figure is created first.

If *parent* is given then a uicontrol relative to *parent* is created.

Any provided property value pairs will override the default values of the created uicontrol object.

The full list of properties is documented at Section 15.3.3.16 [Uicontrol Properties], page 525.

The type of uicontrol created is specified by the *style* property. If no style property is provided, a push button will be created.

Valid styles for uicontrol are:

`"checkbox"`
: Create a checkbox control that allows user on/off selection.

`"edit"`
: Create an edit control that allows user input of single or multiple lines of text.

`"listbox"`
: Create a listbox control that displays a list of items and allows user selection of single or multiple items.

`"popupmenu"`
: Create a popupmenu control that displays a list of options that can be selected when the user clicks on the control.

`"pushbutton"`
: Create a push button control that allows user to press to cause an action.

`"radiobutton"`
: Create a radio button control intended to be used for mutually exclusive input in a group of radiobutton controls.

`"slider"`
: Create a slider control that allows user selection from a range of values by sliding knob on the control.

`"text"`
: Create a static text control to display single or multiple lines of text.

`"togglebutton"`
: Create a toggle button control that appears like a push button but allows the user to select between two states.

Note: For the `"edit"` and `"listbox"` styles, the single or multiple line/selection behavior is determined by the `"Min"` and `"Max"` properties, permitting multiple lines/selections when the values are set such that `Max - Min > 1`.

Examples:

```
## Create figure and panel on it
f = figure;
## Create a button (default style)
b1 = uicontrol (f, "string", "A Button", ...
                   "position", [10 10 150 40]);
## Create an edit control
e1 = uicontrol (f, "style", "edit", "string", "editable text", ...
                   "position", [10 60 300 40]);
## Create a checkbox
c1 = uicontrol (f, "style", "checkbox", "string", "a checkbox", ...
                   "position", [10 120 150 40]);
```

When called with a single argument *h* which is a handle to an existing uicontrol object, switch the keyboard focus to the specified uicontrol. As a result, the uicontrol object will receive keyboard events that can be processed using the `"keypressfcn"` callback.

**See also:** [figure], page 425, [uipanel], page 1008.

*hui* = uitable (*property*, *value*, ...)
*hui* = uitable (*parent*, *property*, *value*, ...)

Create a uitable object and return a handle to it.

A uitable object is used to show tables of data in a figure window.

If *parent* is omitted then a uitable for the current figure is created. If no figure is available, a new figure is created first.

If *parent* is given then a uitable relative to *parent* is created.

Any provided property value pairs will override the default values of the created uitable object.

The full list of properties is documented at Section 15.3.3.17 [Uitable Properties], page 529.

Examples:

```
## Create figure and place a table on it
f = figure ();
m = magic (8);
t = uitable (f, "Data", m, "ColumnWidth", { 40 });

## Create a table with labeled rows and columns
f = figure ();
d = reshape (1:9, [3, 3]);
row_names = { "Row1", "Row2", "Row3" };
col_names = { "Col1", "Col2", "Col3" };
t = uitable (f, "Data", d, ...
             "RowName", row_names, "ColumnName", col_names);

p = get (t, "Position");
e = get (t, "Extent");
p(3:4) = e(3:4);
set (t, "Position", p);

## Long demo with callbacks
function uitable_demo ()
```

```
  f = figure ("Name", "uitable Demo", "Menu", "none", ...
              "Position", [10 10 1000 680]);

  ## A basic example
  d = { "char"   , "A string";
        "double" , 12.3456789;
        "complex", 1+2i;
        "bool"   , true;
        "single" , single(12.3456789);
        "int8"   , int8(-128);
        "uint8"  , uint8(128);
        "int16"  , int16(-32768);
        "uint16" , uint16(32768);
        "int32"  , int32(-2147483648);
        "uint32" , uint32(2147483648);
        "int64"  , int64(-2147483649);
        "uint64" , uint64(2147843649)};

  popup_options = {"A", "B", "C", "D", "E"};

  columnformat_options = { "[]", "char", "pop-up", "numeric", ...
                           "short", "short e", "short eng", ...
                           "short g", "long", "long e", ...
                           "long eng", "long g", "bank", "+", ...
                           "rat", "logical"};
  columnformat_values = columnformat_options;
  columnformat_values{1} = "";
  columnformat_values{3} = popup_options;

  default_data = repmat (d(:,2), 1, columns (columnformat_options));
  b_add = uicontrol (f, "Position", [285 630 600 50], ...
            "UserData", [rows(d), 1], ...
            "Style", "pushbutton", ...
            "String", "Set data at selected point to selected datatype");

  l_type_table = uicontrol (f, "Position", [ 0 603 120 25 ], ...
      "String", "Datatype Table:", ...
      "Style", "text");
  t_type_table = uitable (f, "Position", [ 0 530 1000 70 ], ...
      "Data", transpose (d(:, 2)), ...
      "ColumnName", transpose (d(:, 1)), ...
      "RowName", "Value", ...
      "CellSelectionCallback", ...
@(x, y) set (b_add, "UserData", y.Indices ));

  l_point_table = uicontrol (f, "Position", [ 0 640 60 25 ], ...
      "String", "Point:", ...
      "Style", "text");
  t_point_table = uitable (f, "Position", [ 80 630 160 42 ], ...
      "RowName", [], ...
      "ColumnName", {"x", "y"}, ...
      "Data", [ 1, 1 ], ...
      "ColumnEditable", true);

  l_editable_table = uicontrol (f, "Position", [ 0 502 200 25 ], ...
      "Style", "text", ...
      "String", "Set Data Columns Editable:");
  t_editable_table = ...
```

```
    uitable (f, "Position", [ 0 434 1000 65 ], ...
                "Data", repmat (false, 1, columns (default_data)), ...
                "ColumnEditable", true);

  l_format_table = uicontrol (f, "Position", [ 0 406 200 25 ], ...
      "Style", "text", ...
      "String", "Set Data Column Format:");
  t_format_table = ...
    uitable (f, "Position", [ 0 338 1000 65 ], ...
                "Data", columnformat_options, ...
                "ColumnEditable", true, ...
                "ColumnFormat", arrayfun (@(x) {columnformat_options}, ...
                                          1:columns (columnformat_options)));

  l_data_table = uicontrol (f, "Style", "text", ...
                               "String", "Data:", ...
                               "Position", [ 0 310 60 25 ]);
  t_data_table = uitable (f, "Position", [ 0 15 1000 290 ], ...
      "Data", default_data, ...
      "ColumnFormat", columnformat_values);

  set (t_format_table, ...
       "CellEditCallback", ...
@(x, y) update_column_format (y.NewData, y.Indices, ...
                                      t_data_table, popup_options));
  set (t_point_table, "CellEditCallback", ...
@(x, y) validate_point_table (x, y, t_data_table));
  set (t_editable_table, "CellEditCallback", ...
@(x,y) set (t_data_table, ...
                    "ColumnEditable", get (t_editable_table, "Data")));
  set (b_add, ...
       "Callback", @(x, y) update_data (b_add, t_point_table, ...
                                        t_type_table, t_data_table));
  set (t_data_table, "CellSelectionCallback", ...
@(x, y) update_point_table (y.Indices, t_point_table));
endfunction

function validate_point_table (h, dat, t_data_table)
  if (! (dat.NewData > 0 && ...
    dat.NewData < size (get (t_data_table, "Data"), dat.Indices(1, 1)) + 1))

    d = get (h, "Data");
    d(dat.Indices) = 1;
    set (h, "Data", d);
  endif
endfunction

function update_column_format (format, indices, t_data_table, ...
                               popup_options)
  cf = get (t_data_table, "ColumnFormat");
  if (strcmp (format, "[]"))
    format = "";
  elseif (strcmp (format, "pop-up"))
    format = popup_options;
  endif
  cf{indices(1,2)} = format;
  set (t_data_table, "ColumnFormat", cf);
endfunction
```

```
function update_point_table (indices, t_point_table)
  if (isempty (indices))
    indices = [1, 1];
  endif
  set (t_point_table, "Data", indices(1,:));
endfunction

function update_data (b_add, t_point_table, t_type_table, ...
                      t_data_table)
  indices = get (b_add, "UserData");
  if (isempty (indices))
    indices = [1, 1];
  endif
  d = get (t_data_table, "Data");
  t_type_table_data = get (t_type_table, "Data");
  p = get (t_point_table, "Data");
  d(p(1,2), p(1,1)) = t_type_table_data(indices(1,2));
  set (t_data_table, "Data", d);
endfunction
```

**See also:** [figure], page 425, [uicontrol], page 1009.

*hui* = uimenu (*property*, *value*, ...)
*hui* = uimenu (*h*, *property*, *value*, ...)

Create a uimenu object and return a handle to it.

If *h* is omitted then a top-level menu for the current figure is created. If *h* is given then a submenu relative to *h* is created.

uimenu objects have the following specific properties:

`"accelerator"`
: A string containing the key, together with CTRL, to execute this menu entry (e.g., `"x"` for CTRL+x).

`"checked"`
: Can be set `"on"` or `"off"`. Sets a mark at this menu entry.

`"enable"`
: Can be set `"on"` or `"off"`. If disabled then the menu entry cannot be selected and is grayed out.

`"foregroundcolor"`
: A color value for the text of the menu entry.

`"menuselectedfcn"`
: The function called when this menu entry is executed. It can be either a function string (e.g., `"myfcn"`), a function handle (e.g., @myfcn) or a cell array containing the function handle and arguments for the callback function (e.g., {@myfcn, arg1, arg2}).

`"position"`
: A scalar value containing the relative menu position. The first position has value 1 and will be either the left or top depending on the orientation of the uimenu.

`"separator"`
: Can be set `"on"` or `"off"`. If enabled, a separator line is drawn above the current position. This property is ignored for top-level entries.

`"text"`
: A string containing the text for this menu entry. A `"&"`-symbol can be used to mark the `"accelerator"` character (e.g., `"E&xit"`).

The full list of properties is documented at Section 15.3.3.12 [Uimenu Properties], page 513.

Examples:

```
f = uimenu ("text", "&File", "accelerator", "f");
e = uimenu ("text", "&Edit", "accelerator", "e");
uimenu (f, "text", "Close", "accelerator", "q", ...
           "menuselectedfcn", "close (gcf)");
uimenu (e, "text", "Toggle &Grid", "accelerator", "g", ...
           "menuselectedfcn", "grid (gca)");
```

**See also:** [figure], page 425.

*hui* = uicontextmenu (*property*, *value*, ...)
*hui* = uicontextmenu (*h*, *property*, *value*, ...)

Create a uicontextmenu object and return a handle to it.

If *h* is omitted then a uicontextmenu for the current figure is created. If no figure is available, a new figure is created first.

If *h* is given then a uicontextmenu relative to *h* is created.

Any provided property value pairs will override the default values of the created uicontextmenu object.

The full list of properties is documented at Section 15.3.3.14 [Uicontextmenu Properties], page 519.

Examples:

```
## create figure and uicontextmenu
f = figure ();
c = uicontextmenu (f);

## create menus in the context menu
m1 = uimenu ("parent", c, "label", "Menu item 1", ...
             "callback", "disp('menu item 1')");
m2 = uimenu ("parent", c, "label", "Menu item 2", ...
             "callback", "disp('menu item 2')");

## set the context menu for the figure
set (f, "uicontextmenu", c);
```

**See also:** [figure], page 425, [uimenu], page 1014.

*hui* = uitoolbar ()
*hui* = uitoolbar (*property*, *value*, ...)
*hui* = uitoolbar (*parent*)

*hui* = uitoolbar (*parent*, *property*, *value*, ...)

Create a uitoolbar object. A uitoolbar displays uitoggletool and uipushtool buttons.

If *parent* is omitted then a uitoolbar for the current figure is created. If no figure is available, a new figure is created first.

If *parent* is given then a uitoolbar relative to *parent* is created.

Any provided property value pairs will override the default values of the created uitoolbar object.

The full list of properties is documented at Section 15.3.3.18 [Uitoolbar Properties], page 534.

The optional return value *hui* is a graphics handle to the created uitoolbar object.

Examples:

```
% create figure without a default toolbar
f = figure ("toolbar", "none");
% create empty toolbar
t = uitoolbar (f);
```

**See also:** [figure], page 425, [uitoggletool], page 1017, [uipushtool], page 1016.

*hui* = uipushtool ()
*hui* = uipushtool (*property*, *value*, ...)
*hui* = uipushtool (*parent*)
*hui* = uipushtool (*parent*, *property*, *value*, ...)

Create a uipushtool object.

uipushtools are buttons that appear on a figure toolbar. The button is created with a border that is shown when the user hovers over the button. An image can be set using the cdata property.

If *parent* is omitted then a uipushtool for the current figure is created. If no figure is available, a new figure is created first. If a figure is available, but does not contain a uitoolbar, a uitoolbar will be created.

If *parent* is given then a uipushtool is created on the *parent* uitoolbar.

Any provided property value pairs will override the default values of the created uipushtool object.

The full list of properties is documented at Section 15.3.3.19 [Uipushtool Properties], page 536.

The optional return value *hui* is a graphics handle to the created uipushtool object.

Examples:

```
% create figure without a default toolbar
f = figure ("toolbar", "none");
% create empty toolbar
t = uitoolbar (f);
% create a 19x19x3 black square
img=zeros(19,19,3);
% add pushtool button to toolbar
b = uipushtool (t, "cdata", img);
```

**See also:** [figure], page 425, [uitoolbar], page 1015, [uitoggletool], page 1017.

*hui* = uitoggletool ()
*hui* = uitoggletool (*property*, *value*, ...)
*hui* = uitoggletool (*parent*)
*hui* = uitoggletool (*parent*, *property*, *value*, ...)

Create a uitoggletool object.

uitoggletool are togglebuttons that appear on a figure toolbar. The button is created with a border that is shown when the user hovers over the button. An image can be set using the cdata property.

If *parent* is omitted then a uitoggletool for the current figure is created. If no figure is available, a new figure is created first. If a figure is available, but does not contain a uitoolbar, a uitoolbar will be created.

If *parent* is given then a uitoggletool is created on the *parent* uitoolbar.

Any provided property value pairs will override the default values of the created uitoggletool object.

The full list of properties is documented at Section 15.3.3.20 [Uitoggletool Properties], page 539.

The optional return value *hui* is a graphics handle to the created uitoggletool object.

Examples:

```
% create figure without a default toolbar
f = figure ("toolbar", "none");
% create empty toolbar
t = uitoolbar (f);
% create a 19x19x3 black square
img=zeros(19,19,3);
% add uitoggletool button to toolbar
b = uitoggletool (t, "cdata", img);
```

**See also:** [figure], page 425, [uitoolbar], page 1015, [uipushtool], page 1016.

## 35.4 GUI Utility Functions

These functions do not implement a GUI element but are useful when developing programs that do. The functions `uiwait`, `uiresume`, and `waitfor` are only available with the `qt` or `fltk` toolkits.

*data* = guidata (*h*)
guidata (*h*, *data*)

Query or set user-custom GUI data.

The GUI data is stored in the figure handle *h*. If *h* is not a figure handle then it's parent figure will be used for storage.

*data* must be a single object which means it is usually preferable for it to be a data container such as a cell array or struct so that additional data items can be added easily.

**See also:** [getappdata], page 548, [setappdata], page 548, [get], page 457, [set], page 457, [getpref], page 1021, [setpref], page 1021.

*hdata* = guihandles (*h*)
*hdata* = guihandles

Return a structure of object handles for the figure associated with handle *h*.

If no handle is specified the current figure returned by `gcf` is used.

The fieldname for each entry of *hdata* is taken from the `"tag"` property of the graphic object. If the tag is empty then the handle is not returned. If there are multiple graphic objects with the same tag then the entry in *hdata* will be a vector of handles. `guihandles` includes all possible handles, including those for which `"HandleVisibility"` is `"off"`.

**See also:** [guidata], page 1017, [findobj], page 542, [findall], page 543, [allchild], page 458.

*tf* = have_window_system ()

Return true if a window system is available (X11, Windows, or Apple OS X) and false otherwise.

**See also:** [isguirunning], page 1018.

*tf* = isguirunning ()

Return true if Octave is running in GUI mode and false otherwise.

**See also:** [have_window_system], page 1018.

*pos* = getpixelposition (*h*)
*pos* = getpixelposition (*h*, *rel_to_fig*)

Return the position of a user interface component in pixel units.

The first argument *h* must be a handle to a valid graphics object of type uibuttongroup, uicontrol, uipanel, uitable, axes, or figure. For other object types, the function returns zeros.

By default, the position is returned relative to the object's parent. If the second argument *rel_to_fig* is logically true, the position is computed relative to the enclosing figure object.

The return value *pos* is a 4-element vector with values `[lower_left_X, lower_left_Y, width, height]`.

**See also:** [get], page 457.

fonts = listfonts ()
fonts = listfonts (*h*)

List system fonts.

If a handle to a graphics object *h* is provided, also include the font from the object's `"FontName"` property in the list.

Programming Note: On systems that don't use FontConfig natively (all but Linux), the font cache is built when Octave is installed. You will need to run `system ("fc-cache -fv")` manually after installing new fonts.

**See also:** [uisetfont], page 1006, [text], page 417, [axes], page 450, [uicontrol], page 1009.

`movegui`
`movegui (`*h*`)`
`movegui (`*pos*`)`
`movegui (`*h*`, `*pos*`)`
`movegui (`*h*`, `*event*`)`
`movegui (`*h*`, `*event*`, `*pos*`)`

Move a figure specified by figure handle *h* to a position on the screen defined by *pos*.

*h* is a figure handle, or a handle to a graphics object. In the latter case, its parent figure will be used. If unspecified, *h* will be set to the handle of the relevant figure if a callback is being executed (`gcbf`), otherwise it will be set to the handle of the current figure (`gcf`).

*pos* is either a two-value numeric vector or a string. If *pos* is numeric then it must be of the form `[h, v]` specifying the horizontal and vertical offsets of the figure with respect to the screen. A positive value indicates the offset between the left (or bottom for the vertical component) of the screen, and the left (or bottom) of the figure. A negative value indicates the offset between the right (or top) of the screen and the right (or top) of the figure.

Possible values for *pos* as a string are

`north` Top center of the screen.

`south` Bottom center of the screen.

`east` Right center of the screen.

`west` Left center of the screen.

`northeast`
Top right of the screen.

`northwest`
Top left of the screen.

`southeast`
Bottom right of the screen.

`southwest`
Bottom left of the screen.

`center` Center of the screen.

`onscreen (default)`
The figure will be minimally moved to be entirely visible on the screen, with a 30 pixel extra padding from the sides of the screen. This is the default value if none is provided.

*event* contains event data that will be ignored. This construct facilitates a call to movegui from a callback.

`openvar (`*name*`)`
Open the variable *name* in the graphical Variable Editor.

`uiwait`
`uiwait (`*`h`*`)`
`uiwait (`*`h`*`, `*`timeout`*`)`

Suspend program execution until the figure with handle *h* is deleted or `uiresume` is called.

When no figure handle is specified this function uses the current figure. If the figure handle is invalid or there is no current figure, this functions returns immediately.

When specified, *timeout* defines the number of seconds to wait for the figure deletion or the `uiresume` call. The timeout value must be at least 1. If a smaller value is specified, a warning is issued and a timeout value of 1 is used instead. If a non-integer value is specified, it is truncated towards 0. If *timeout* is not specified, the program execution is suspended indefinitely.

**See also:** [uiresume], page 1020, [waitfor], page 1020.

`uiresume (`*`h`*`)`

Resume program execution suspended with `uiwait`.

The handle *h* must be the same as the on specified in `uiwait`. If the handle is invalid or there is no `uiwait` call pending for the figure with handle *h*, this function does nothing.

**See also:** [uiwait], page 1020.

`waitfor (`*`h`*`)`
`waitfor (`*`h`*`, `*`prop`*`)`
`waitfor (`*`h`*`, `*`prop`*`, `*`value`*`)`
`waitfor (`*`...`*`, "`*`timeout`*`", `*`timeout`*`)`

Suspend the execution of the current program until a condition is satisfied on the graphics handle *h*.

While the program is suspended graphics events are still processed normally, allowing callbacks to modify the state of graphics objects. This function is reentrant and can be called from a callback, while another `waitfor` call is pending at the top-level.

In the first form, program execution is suspended until the graphics object *h* is destroyed. If the graphics handle is invalid or if *h* is the root graphics handle and no property *prop* was provided, the function returns immediately.

In the second form, execution is suspended until the graphics object is destroyed or the property named *prop* is modified. If the graphics handle is invalid or the property does not exist, the function returns immediately.

In the third form, execution is suspended until the graphics object is destroyed or the property named *prop* is set to *value*. The function `isequal` is used to compare property values. If the graphics handle is invalid, the property does not exist or the property is already set to *value*, the function returns immediately.

An optional timeout can be specified using the property `"timeout"`. This timeout value is the number of seconds to wait for the condition to be true. *timeout* must be at least 1. If a smaller value is specified, a warning is issued and a value of 1 is used instead. If the timeout value is not an integer, it is truncated towards 0.

To define a condition on a property named `"timeout"`, use the string `'\timeout'` instead.

In all cases, typing CTRL-C stops program execution immediately.

**See also:** [waitforbuttonpress], page 446, [isequal], page 175.

## 35.5 User-Defined Preferences

*val* = getpref (*"group"*, *"pref"*)
*val* = getpref (*"group"*, *"pref"*, *default*)
{*val1*, *val2*, ...} = getpref (*"group"*, {*"pref1"*, *"pref2"*, ...})
*prefstruct* = getpref (*"group"*)
*prefstruct* = getpref ()

Return the preference value corresponding to the named preference *pref* in the preference group *group*.

The named preference group must be a string.

If *pref* does not exist in *group* and *default* is specified, create the preference with value *default* and return *default*.

The preference *pref* may be a string or cell array of strings. If it is a cell array of strings then a cell array of preferences is returned.

The corresponding default value *default* may be any Octave value, .e.g., double, struct, cell array, object, etc. Or, if *pref* is a cell array of strings then *default* must be a cell array of values with the same size as *pref*.

If neither *pref* nor *default* are specified, return a structure of preferences for the preference group *group*.

If no arguments are specified, return a structure containing all groups of preferences and their values.

**See also:** [addpref], page 1021, [setpref], page 1021, [ispref], page 1022, [rmpref], page 1022.

setpref (*"group"*, *"pref"*, *val*)
setpref (*"group"*, {*"pref1"*, *"pref2"*, ...}, {*val1*, *val2*, ...})

Set the preference *pref* to the given *val* in the named preference group *group*.

The named preference group must be a string.

The preference *pref* may be a string or a cell array of strings.

The corresponding value *val* may be any Octave value, .e.g., double, struct, cell array, object, etc. Or, if *pref* is a cell array of strings then *val* must be a cell array of values with the same size as *pref*.

If the named preference or group does not exist, it is added.

**See also:** [addpref], page 1021, [getpref], page 1021, [ispref], page 1022, [rmpref], page 1022.

addpref (*"group"*, *"pref"*, *val*)
addpref (*"group"*, {*"pref1"*, *"pref2"*, ...}, {*val1*, *val2*, ...})

Add the preference *pref* and associated value *val* to the named preference group *group*.

The named preference group must be a string.

The preference *pref* may be a string or a cell array of strings. An error will be issued if the preference already exists.

The corresponding value *val* may be any Octave value, .e.g., double, struct, cell array, object, etc. Or, if *pref* is a cell array of strings then *val* must be a cell array of values with the same size as *pref*.

**See also:** [setpref], page 1021, [getpref], page 1021, [ispref], page 1022, [rmpref], page 1022.

`rmpref ("group", "pref")`
`rmpref ("group", {"pref1", "pref2", ...})`
`rmpref ("group")`

Remove the named preference *pref* from the preference group *group*.

The named preference group must be a string.

The preference *pref* may be a string or cell array of strings.

If *pref* is not specified, remove the preference group *group*.

It is an error to remove a nonexistent preference or group.

**See also:** [addpref], page 1021, [ispref], page 1022, [setpref], page 1021, [getpref], page 1021.

`tf = ispref ("group", "pref")`
`tf = ispref ("group", {"pref1", "pref2", ...})`
`tf = ispref ("group")`

Return true if the named preference *pref* exists in the preference group *group*.

The named preference group must be a string.

The preference *pref* may be a string or a cell array of strings.

If *pref* is not specified, return true if the preference group *group* exists.

**See also:** [getpref], page 1021, [addpref], page 1021, [setpref], page 1021, [rmpref], page 1022.

`dir = prefdir`
`dir = prefdir (1)`

Return the directory that holds the preferences for Octave.

Examples:

Display the preferences directory

```
prefdir
```

Change to the preferences folder

```
cd (prefdir)
```

If called with an argument, the preferences directory is created if it doesn't already exist.

**See also:** [getpref], page 1021, [setpref], page 1021, [addpref], page 1021, [rmpref], page 1022, [ispref], page 1022.

`preferences`

Display the GUI preferences dialog window for Octave.

## 35.6 Octave Workspace Windows

The functions below make windows that are a part of Octave's own GUI interface visible, and move the keyboard focus to the selected window. Their utility lies in the ability to call these functions from a script and highlight a window without using a mouse for selection.

`commandhistory` ()
: Show the GUI command history window and give it the keyboard focus.

  **See also:** [commandwindow], page 1023, [filebrowser], page 1023, [workspace], page 1023.

`commandwindow` ()
: Show the GUI command window and give it the keyboard focus.

  **See also:** [commandhistory], page 1023, [filebrowser], page 1023, [workspace], page 1023.

`filebrowser` ()
: Show the GUI file browser window and give it the keyboard focus.

  **See also:** [commandwindow], page 1023, [commandhistory], page 1023, [workspace], page 1023.

`workspace` ()
: Show the GUI workspace window and give it the keyboard focus.

  **See also:** [commandwindow], page 1023, [commandhistory], page 1023, [filebrowser], page 1023.

# 36 System Utilities

This chapter describes the functions that are available to allow you to get information about what is happening outside of Octave, while it is still running, and use this information in your program. For example, you can get information about environment variables, the current time, and even start other programs from the Octave prompt.

## 36.1 Timing Utilities

Octave's core set of functions for manipulating time values are patterned after the corresponding functions from the standard C library. Several of these functions use a data structure for time that includes the following elements:

`usec` Microseconds after the second (0-999999).

`sec` Seconds after the minute (0-60). This number can be 60 to account for leap seconds.

`min` Minutes after the hour (0-59).

`hour` Hours since midnight (0-23).

`mday` Day of the month (1-31).

`mon` Months since January (0-11).

`year` Years since 1900.

`wday` Days since Sunday (0-6).

`yday` Days since January 1 (0-365).

`isdst` Daylight saving time flag.

`gmtoff` Seconds offset from UTC.

`zone` Time zone.

In the descriptions of the following functions, this structure is referred to as a *tm_struct*.

*seconds* = `time` ()

Return the current time as the number of seconds since the epoch.

The epoch is referenced to 00:00:00 UTC (Coordinated Universal Time) 1 Jan 1970. For example, on Monday February 17, 1997 at 07:15:06 UTC, the value returned by `time` was 856163706.

**See also:** [strftime], page 1027, [strptime], page 1029, [localtime], page 1026, [gmtime], page 1026, [mktime], page 1027, [now], page 1025, [date], page 1030, [clock], page 1029, [datenum], page 1032, [datestr], page 1034, [datevec], page 1036, [calendar], page 1036, [weekday], page 1037.

`t = now ()`

Return the current local date/time as a serial day number (see [`datenum`], page 1032).

The integral part, `floor (now)` corresponds to the number of days between today and Jan 1, 0000.

The fractional part, `rem (now, 1)` corresponds to the current time.

**See also:** [clock], page 1029, [date], page 1030, [datenum], page 1032.

*str* = ctime (*t*)

Convert a value returned from `time` (or any other non-negative integer), to the local time and return a string of the same form as `asctime`.

The function `ctime (time)` is equivalent to `asctime (localtime (time))`. For example:

```
ctime (time ())
   ⇒ "Mon Feb 17 01:15:06 1997\n"
```

**See also:** [asctime], page 1027, [time], page 1025, [localtime], page 1026.

*tm_struct* = gmtime (*t*)

Given a value returned from `time`, or any non-negative integer, return a time structure corresponding to UTC (Coordinated Universal Time).

For example:

```
gmtime (time ())
     ⇒ {
           usec = 0
           sec = 6
           min = 15
           hour = 7
           mday = 17
           mon = 1
           year = 97
           wday = 1
           yday = 47
           isdst = 0
           gmtoff = 0
           zone = GMT
       }
```

**See also:** [strftime], page 1027, [strptime], page 1029, [localtime], page 1026, [mktime], page 1027, [time], page 1025, [now], page 1025, [date], page 1030, [clock], page 1029, [datenum], page 1032, [datestr], page 1034, [datevec], page 1036, [calendar], page 1036, [weekday], page 1037.

*tm_struct* = localtime (*t*)

Given a value returned from `time`, or any non-negative integer, return a time structure corresponding to the local time zone.

```
localtime (time ())
     ⇒ {
           usec = 0
           sec = 6
           min = 15
           hour = 1
           mday = 17
           mon = 1
           year = 97
           wday = 1
           yday = 47
           isdst = 0
           gmtoff = -21600
           zone = CST
       }
```

**See also:** [strftime], page 1027, [strptime], page 1029, [gmtime], page 1026, [mktime], page 1027, [time], page 1025, [now], page 1025, [date], page 1030, [clock], page 1029, [datenum], page 1032, [datestr], page 1034, [datevec], page 1036, [calendar], page 1036, [weekday], page 1037.

*seconds* = mktime (*tm_struct*)

Convert a time structure corresponding to the local time to the number of seconds since the epoch.

For example:

```
mktime (localtime (time ()))
     ⇒ 856163706
```

**See also:** [strftime], page 1027, [strptime], page 1029, [localtime], page 1026, [gmtime], page 1026, [time], page 1025, [now], page 1025, [date], page 1030, [clock], page 1029, [datenum], page 1032, [datestr], page 1034, [datevec], page 1036, [calendar], page 1036, [weekday], page 1037.

*str* = asctime (*tm_struct*)

Convert a time structure to a string using the following format: `"ddd mmm mm HH:MM:SS yyyy\n"`.

For example:

```
asctime (localtime (time ()))
     ⇒ "Mon Feb 17 01:15:06 1997\n"
```

This is equivalent to `ctime (time ())`.

**See also:** [ctime], page 1026, [localtime], page 1026, [time], page 1025.

*str* = strftime (*fmt*, *tm_struct*)

Format the time structure *tm_struct* in a flexible way using the format string *fmt* that contains '%' substitutions similar to those in `printf`.

Except where noted, substituted fields have a fixed size; numeric fields are padded if necessary. Padding is with zeros by default; for fields that display a single number,

padding can be changed or inhibited by following the '%' with one of the modifiers described below. Unknown field specifiers are copied as normal characters. All other characters are copied to the output without change. For example:

```
strftime ("%r (%Z) %A %e %B %Y", localtime (time ()))
      ⇒ "01:15:06 AM (CST) Monday 17 February 1997"
```

Octave's `strftime` function supports a superset of the ANSI C field specifiers.

Literal character fields:

`%%` % character.

`%n` Newline character.

`%t` Tab character.

Numeric modifiers (a nonstandard extension):

`- (dash)` Do not pad the field.

`_ (underscore)`
Pad the field with spaces.

Time fields:

`%H` Hour (00-23).

`%I` Hour (01-12).

`%k` Hour (0-23).

`%l` Hour (1-12).

`%M` Minute (00-59).

`%p` Locale's AM or PM.

`%r` Time, 12-hour (hh:mm:ss [AP]M).

`%R` Time, 24-hour (hh:mm).

`%s` Time in seconds since 00:00:00, Jan 1, 1970 (a nonstandard extension).

`%S` Second (00-61).

`%T` Time, 24-hour (hh:mm:ss).

`%X` Locale's time representation (%H:%M:%S).

`%z` Offset from UTC (±hhmm), or nothing if no time zone is determinable.

`%Z` Time zone (EDT), or nothing if no time zone is determinable.

Date fields:

`%a` Locale's abbreviated weekday name (Sun-Sat).

`%A` Locale's full weekday name, variable length (Sunday-Saturday).

`%b` Locale's abbreviated month name (Jan-Dec).

`%B` Locale's full month name, variable length (January-December).

| | |
|---|---|
| `%c` | Locale's date and time (Sat Nov 04 12:02:33 EST 1989). |
| `%C` | Century (00-99). |
| `%d` | Day of month (01-31). |
| `%e` | Day of month ( 1-31). |
| `%D` | Date (mm/dd/yy). |
| `%h` | Same as %b. |
| `%j` | Day of year (001-366). |
| `%m` | Month (01-12). |
| `%U` | Week number of year with Sunday as first day of week (00-53). |
| `%w` | Day of week (0-6). |
| `%W` | Week number of year with Monday as first day of week (00-53). |
| `%x` | Locale's date representation (mm/dd/yy). |
| `%y` | Last two digits of year (00-99). |
| `%Y` | Year (1970-). |

**See also:** [strptime], page 1029, [localtime], page 1026, [gmtime], page 1026, [mktime], page 1027, [time], page 1025, [now], page 1025, [date], page 1030, [clock], page 1029, [datenum], page 1032, [datestr], page 1034, [datevec], page 1036, [calendar], page 1036, [weekday], page 1037.

`[`*tm_struct*`, `*nchars*`] = strptime (`*str*`, `*fmt*`)`

Convert the string *str* to the time structure *tm_struct* under the control of the format string *fmt*.

If *fmt* fails to match, *nchars* is 0; otherwise, it is set to the position of last matched character plus 1. Always check for this unless you're absolutely sure the date string will be parsed correctly.

**See also:** [strftime], page 1027, [localtime], page 1026, [gmtime], page 1026, [mktime], page 1027, [time], page 1025, [now], page 1025, [date], page 1030, [clock], page 1029, [datenum], page 1032, [datestr], page 1034, [datevec], page 1036, [calendar], page 1036, [weekday], page 1037.

Most of the remaining functions described in this section are not patterned after the standard C library. Some are available for compatibility with MATLAB and others are provided because they are useful.

*datevec* = clock ()
`[`*datevec*`, `*isdst*`] = clock ()`

Return the current local date and time as a date vector.

The date vector contains the following fields: current year, month (1-12), day (1-31), hour (0-23), minute (0-59), and second (0-61). The seconds field has a fractional part after the decimal point for extended accuracy.

The optional second output *isdst* is true if Daylight Savings Time (DST) is in effect for the system's time zone.

For example:

```
fix (clock ())
  ⇒   1993      8     20      4     56      1
```

`clock` is more accurate on systems that have the `gettimeofday` function.

**See also:** [now], page 1025, [date], page 1030, [datevec], page 1036.

*str* = `date` ()

Return the current date as a character string in the form DD-MMM-YYYY.

For example:

```
date ()
  ⇒ 20-Aug-1993
```

**See also:** [now], page 1025, [clock], page 1029, [datestr], page 1034, [localtime], page 1026.

*secs* = `etime` (*t2*, *t1*)

Return the difference in seconds between two time values returned from `clock` (*t2* − *t1*).

For example:

```
t0 = clock ();
# many computations later...
elapsed_time = etime (clock (), t0);
```

will set the variable `elapsed_time` to the number of seconds since the variable `t0` was set.

**See also:** [tic], page 1031, [toc], page 1031, [clock], page 1029, [cputime], page 1030, [addtodate], page 1036.

[*total*, *user*, *system*] = `cputime` ();

Return the CPU time used by your Octave session.

The first output is the total time spent executing your process and is equal to the sum of second and third outputs, which are the number of CPU seconds spent executing in user mode and the number of CPU seconds spent executing in system mode, respectively.

If your system does not have a way to report CPU time usage, `cputime` returns 0 for each of its output values.

Note that because Octave used some CPU time to start, it is reasonable to check to see if `cputime` works by checking to see if the total CPU time used is nonzero.

**See also:** [tic], page 1031, [toc], page 1031.

*tf* = `is_leap_year` ()
*tf* = `is_leap_year` (*year*)

Return true if *year* is a leap year and false otherwise.

If no year is specified, `is_leap_year` uses the current year.

For example:

```
is_leap_year (2000)
   ⇒ 1
```

**See also:** [weekday], page 1037, [eomday], page 1037, [calendar], page 1036.

`tic ()`
*`id`* `= tic ()`

Initialize a wall-clock timer.

Calling `tic` without an output argument resets the internal timer. Subsequent calls to `toc` return the number of seconds since the timer was set.

If called with one output argument, `tic` creates a new timer instance and returns a timer identifier *id*. The *id* is a scalar of type `uint64` that may be passed to `toc` to check elapsed time on this timer, rather than the default internal timer.

Example 1 : benchmarking code with internal timer

```
tic;
# many computations later...
elapsed_time = toc;
```

Example 2 : mixed timer id and internal timer

```
tic;
pause (1);
toc
⇒ Elapsed time is 1.0089 seconds.
id = tic;
pause (2);
toc (id)
⇒ Elapsed time is 2.01142 seconds.
toc
Elapsed time is 3.02308 seconds.
```

Calling `tic` and `toc` in this way allows nested timing calls.

If you are more interested in the CPU time that your process used, you should use the `cputime` function instead. The `tic` and `toc` functions report the actual wall clock time that elapsed between the calls. This may include time spent processing other jobs or doing nothing at all.

**See also:** [toc], page 1031, [cputime], page 1030.

`toc ()`
`toc (`*`id`*`)`
*`elapsed_time`* `= toc (...)`

Measure elapsed time on a wall-clock timer.

With no arguments, return the number of seconds elapsed on the internal timer since the last call to `tic`.

When given the identifier *id* of a specific timer, return the number of seconds elapsed since the timer *id* was initialized.

See [tic], page 1031, for examples of the use of `tic`/`toc`.

**See also:** [tic], page 1031, [cputime], page 1030.

`pause ()`
`pause (`*`n`*`)`
*`old_state`* `= pause ("`*on*`")`
*`old_state`* `= pause ("`*off*`")`
*`old_state`* `= pause ("`*query*`")`

Suspend the execution of the program or change the state of the pause function.

If invoked without an input arguments then the program is suspended until a character is typed. If argument *n* is a positive real value, it indicates the number of seconds the program shall be suspended, for example:

```
tic; pause (0.05); toc
     ⊣ Elapsed time is 0.05039 seconds.
```

The following example prints a message and then waits 5 seconds before clearing the screen.

```
disp ("wait please...");
pause (5);
clc;
```

If invoked with a string argument `"on"`, `"off"`, or `"query"`, the state of the pause function is changed or queried. When the state is `"off"`, the pause function returns immediately. The optional return value contains the previous state of the pause function. In the following example pause is disabled locally:

```
old_state = pause ("off");
tic; pause (0.05); toc
     ⊣ Elapsed time is 3.00407e-05 seconds.
pause (old_state);
```

While the program is suspended Octave still handles figures painting and graphics callbacks execution.

**See also:** [kbhit], page 294.

*`days`* `= datenum (`*`datevec`*`)`
*`days`* `= datenum (`*`year`*`, `*`month`*`, `*`day`*`)`
*`days`* `= datenum (`*`year`*`, `*`month`*`, `*`day`*`, `*`hour`*`)`
*`days`* `= datenum (`*`year`*`, `*`month`*`, `*`day`*`, `*`hour`*`, `*`minute`*`)`
*`days`* `= datenum (`*`year`*`, `*`month`*`, `*`day`*`, `*`hour`*`, `*`minute`*`, `*`second`*`)`
*`days`* `= datenum ("`*datestr*`")`
*`days`* `= datenum ("`*datestr*`", `*`f`*`)`
*`days`* `= datenum ("`*datestr*`", `*`p`*`)`
`[`*`days`*`, `*`secs`*`] = datenum (...)`

Return the date/time input as a serial day number, with Jan 1, 0000 defined as day 1.

The integer part, `floor (`*`days`*`)` counts the number of complete days in the date input.

The fractional part, `rem (`*`days`*`, 1)` corresponds to the time on the given day.

The input may be a date vector (see [datevec], page 1036), date string (see [datestr], page 1034), or directly specified as input.

When processing input datestrings, *f* is the format string used to interpret date strings (see [datestr], page 1034). If no format *f* is specified, then a relatively slow search

is performed through various formats. It is always preferable to specify the format string *f* if it is known. Formats which do not specify a particular time component will have the value set to zero. Formats which do not specify a date will default to January 1st of the current year.

When passing separate *year*, *month*, *day*, etc. arguments, each may be a scalar or nonscalar array. Nonscalar inputs must all be of the same size. Scalar inputs will be expanded to be the size of the nonscalar inputs.

*p* is the year at the start of the century to which two-digit years will be referenced. If not specified, it defaults to the current year minus 50.

The optional output *secs* holds the time on the specified day with greater precision than *days*.

Notes:

- Years can be negative and/or fractional.
- Months below 1 are considered to be January.
- Days of the month start at 1.
- Days beyond the end of the month go into subsequent months.
- Days before the beginning of the month go to the previous month.
- Days can be fractional.

Examples:

```
Convert from datestrs:
d = datenum ("1966-06-14")
⇒ d = 718232

d = datenum ({"1966-06-14", "1966-06-15", "1966-06-16"})
⇒ d =
      718232
      718233
      718234

Convert from datevec:
d = datenum ([1966 06 14])
⇒ d = 718232

d = datenum ([1966 06 14 23 59 59])
⇒ d = 718232.9999884259

Specify date components separately:
d = datenum (1966, 6, 14)
⇒ d = 718232
```

```
d = datenum (1966, magic(3), 1)
⇒ d =

      718280   718068   718219
      718127   718188   718249
      718158   718311   718099
```

**Caution:** datenums represent a specific time for the Earth as a whole. They do not take in to account time zones (shifts in time based on location), nor seasonal changes due to Daylight Savings Time (shifts in time based on local regulation). Be aware that it is possible to create datenums that, when interpreted by a function which accounts for time zone and DST shifts such as `datestr`, are nonexistent or ambiguous.

**Caution:** this function does not attempt to handle Julian calendars so dates before October 15, 1582 are wrong by as much as eleven days. Also, be aware that only Roman Catholic countries adopted the calendar in 1582. It took until 1924 for it to be adopted everywhere. See the Wikipedia entry on the Gregorian calendar for more details.

**Warning:** leap seconds are ignored. A table of leap seconds is available on the Wikipedia entry for leap seconds.

**See also:** [datestr], page 1034, [datevec], page 1036, [now], page 1025, [clock], page 1029, [date], page 1030.

*str* = `datestr` (*date*)
*str* = `datestr` (*date*, *f*)
*str* = `datestr` (*date*, *f*, *p*)

Format the given date/time according to the format *f* and return the result in *str*.

*date* is a serial date number (see [`datenum`], page 1032), a date vector (see [`datevec`], page 1036), or a string or cell array of strings. In the latter case, it is passed to `datevec` to guess the input date format.

*f* can be an integer which corresponds to one of the codes in the table below, or a date format string.

*p* is the year at the start of the century in which two-digit years are to be interpreted in. If not specified, it defaults to the current year minus 50.

For example, the date 730736.65149 (2000-09-07 15:38:09.0934) would be formatted as follows:

| **Code** | **Format** | **Example** |
|---|---|---|
| 0 | dd-mmm-yyyy HH:MM:SS | 07-Sep-2000 15:38:09 |
| 1 | dd-mmm-yyyy | 07-Sep-2000 |
| 2 | mm/dd/yy | 09/07/00 |
| 3 | mmm | Sep |
| 4 | m | S |
| 5 | mm | 09 |
| 6 | mm/dd | 09/07 |
| 7 | dd | 07 |
| 8 | ddd | Thu |

| | | |
|---|---|---|
| 9 | d | T |
| 10 | yyyy | 2000 |
| 11 | yy | 00 |
| 12 | mmmyy | Sep00 |
| 13 | HH:MM:SS | 15:38:09 |
| 14 | HH:MM:SS PM | 3:38:09 PM |
| 15 | HH:MM | 15:38 |
| 16 | HH:MM PM | 3:38 PM |
| 17 | QQ-YY | Q3-00 |
| 18 | QQ | Q3 |
| 19 | dd/mm | 07/09 |
| 20 | dd/mm/yy | 07/09/00 |
| 21 | mmm.dd,yyyy HH:MM:SS | Sep.07,2000 15:38:08 |
| 22 | mmm.dd,yyyy | Sep.07,2000 |
| 23 | mm/dd/yyyy | 09/07/2000 |
| 24 | dd/mm/yyyy | 07/09/2000 |
| 25 | yy/mm/dd | 00/09/07 |
| 26 | yyyy/mm/dd | 2000/09/07 |
| 27 | QQ-YYYY | Q3-2000 |
| 28 | mmmyyyy | Sep2000 |
| 29 | yyyy-mm-dd | 2000-09-07 |
| 30 | yyyymmddTHHMMSS | 20000907T153808 |
| 31 | yyyy-mm-dd HH:MM:SS | 2000-09-07 15:38:08 |

If *f* is a format string, the following symbols are recognized:

| **Symbol** | **Meaning** | **Example** |
|---|---|---|
| yyyy | Full year | 2005 |
| yy | Two-digit year | 05 |
| mmmm | Full month name | December |
| mmm | Abbreviated month name | Dec |
| mm | Numeric month number (padded with zeros) | 01, 08, 12 |
| m | First letter of month name (capitalized) | D |
| dddd | Full weekday name | Sunday |
| ddd | Abbreviated weekday name | Sun |
| dd | Numeric day of month (padded with zeros) | 11 |
| d | First letter of weekday name (capitalized) | S |
| HH | Hour of day, padded with zeros, | 09:00 |
| | or padded with spaces if PM is set | 9:00 AM |
| MM | Minute of hour (padded with zeros) | 10:05 |
| SS | Second of minute (padded with zeros) | 10:05:03 |
| FFF | Milliseconds of second (padded with zeros) | 10:05:03.012 |
| AM | Use 12-hour time format | 11:30 AM |
| PM | Use 12-hour time format | 11:30 PM |

If *f* is not specified or is `-1`, then use 0, 1 or 16, depending on whether the date portion or the time portion of *date* is empty.

If *p* is not specified, it defaults to the current year minus 50.

If a matrix or cell array of dates is given, a column vector of date strings is returned.

**See also:** [datenum], page 1032, [datevec], page 1036, [date], page 1030, [now], page 1025, [clock], page 1029.

`v = datevec (date)`
`v = datevec (date, f)`
`v = datevec (date, p)`
`v = datevec (date, f, p)`
`[y, m, d, h, mi, s] = datevec (...)`

Convert a serial date number (see [datenum], page 1032) or date string (see [datestr], page 1034) into a date vector.

A date vector is a row vector with six members, representing the year, month, day, hour, minute, and seconds respectively.

Date number inputs can be either a scalar or nonscalar array. Date string inputs can be either a single date string, a two-dimensional character array of dates with each row being a date string, or a cell string array of any dimension with each cell element containing a single date string.

*v* is a two-dimensional array of date vectors, one date vector per row. For array inputs, ordering of *v* is based on column major order of dates in *data*.

*f* is the format string used to interpret date strings (see [datestr], page 1034). If *date* is a string or a cell array of strings, but no format is specified, heuristics are used to guess the input format. These heuristics could lead to matches that differ from the result a user might expect. Additionally, this involves a relatively slow search through various formats. It is always preferable to specify the format string *f* if it is known. Formats which do not specify a particular time component will have the value set to zero. Formats which do not specify a particular date component will default that component to January 1st of the current year. Trailing characters are ignored for the purpose of calculating the date vector, even if the characters contain additional time/date information.

*p* is the year at the start of the century to which two-digit years will be referenced. If not specified, it defaults to the current year minus 50.

**See also:** [datenum], page 1032, [datestr], page 1034, [clock], page 1029, [now], page 1025, [date], page 1030.

`d = addtodate (d, q, f)`

Add *q* amount of time (with units *f*) to the serial datenum, *d*.

*f* must be one of `"year"`, `"month"`, `"day"`, `"hour"`, `"minute"`, `"second"`, or `"millisecond"`.

**See also:** [datenum], page 1032, [datevec], page 1036, [etime], page 1030.

`c = calendar ()`
`c = calendar (d)`
`c = calendar (y, m)`
`calendar (...)`

Return the current monthly calendar in a 6x7 matrix.

If *d* is specified, return the calendar for the month containing the date *d*, which must be a serial date number or a date string.

If *y* and *m* are specified, return the calendar for year *y* and month *m*.

If no output arguments are specified, print the calendar on the screen instead of returning a matrix.

**See also:** [datenum], page 1032, [datestr], page 1034.

`[n, s] = weekday (d)`
`[n, s] = weekday (d, format)`

Return the day of the week as a number in *n* and as a string in *s*.

The days of the week are numbered 1–7 with the first day being Sunday.

*d* is a serial date number or a date string.

If the string *format* is not present or is equal to `"short"` then *s* will contain the abbreviated name of the weekday. If *format* is `"long"` then *s* will contain the full name.

Table of return values based on *format*:

| *n* | `"short"` | `"long"` |
|---|---|---|
| 1 | Sun | Sunday |
| 2 | Mon | Monday |
| 3 | Tue | Tuesday |
| 4 | Wed | Wednesday |
| 5 | Thu | Thursday |
| 6 | Fri | Friday |
| 7 | Sat | Saturday |

**See also:** [eomday], page 1037, [is_leap_year], page 1030, [calendar], page 1036, [datenum], page 1032, [datevec], page 1036.

`e = eomday (y, m)`

Return the last day of the month *m* for the year *y*.

**See also:** [weekday], page 1037, [datenum], page 1032, [datevec], page 1036, [is_leap_year], page 1030, [calendar], page 1036.

`datetick ()`
`datetick (axis_str)`
`datetick (date_format)`
`datetick (axis_str, date_format)`
`datetick (..., "keeplimits")`
`datetick (..., "keepticks")`
`datetick (hax, ...)`

Add date-formatted tick labels to an axis.

The axis to apply the ticks to is determined by *axis_str* which can take the values `"x"`, `"y"`, or `"z"`. The default value is `"x"`.

The formatting of the labels is determined by the variable *date_format*, which can either be a string or positive integer that `datestr` accepts.

If the first argument *hax* is an axes handle, then plot into this axes, rather than the current axes returned by `gca`.

**See also:** [datenum], page 1032, [datestr], page 1034.

## 36.2 Filesystem Utilities

Octave includes many utility functions for copying, moving, renaming, and deleting files; for creating, reading, and deleting directories; for retrieving status information on files; and for manipulating file and path names.

`movefile` *f1*
`movefile` *f1 f2*
`movefile` *f1 f2 f*
`movefile (`*f1*`)`
`movefile (`*f1*`, `*f2*`)`
`movefile (`*f1*`, `*f2*`, 'f')`
`[`*status*`] = movefile (...)`
`[`*status*`, `*msg*`] = movefile (...)`
`[`*status*`, `*msg*`, `*msgid*`] = movefile (...)`

Move the source file or directory *f1* to the destination *f2*.

The name *f1* may contain globbing patterns, or may be a cell array of strings. If *f1* expands to multiple filenames, *f2* must be a directory.

If no destination *f2* is specified then the destination is the present working directory. If *f2* is a filename then *f1* is renamed to *f2*.

When the force flag `'f'` is given any existing files will be overwritten without prompting.

If successful, *status* is logical 1, and *msg*, *msgid* are empty character strings (`""`). Otherwise, *status* is logical 0, *msg* contains a system-dependent error message, and *msgid* contains a unique message identifier. Note that the status code is exactly opposite that of the `system` command.

**See also:** [rename], page 1038, [copyfile], page 1038, [unlink], page 1039, [delete], page 430, [glob], page 1044.

`rename` *old new*
`[`*status*`, `*msg*`] = rename (`*old*`, `*new*`)`

Change the name of file *old* to *new*.

If successful, *status* is 0 and *msg* is an empty string. Otherwise, *status* is -1 and *msg* contains a system-dependent error message.

**See also:** [movefile], page 1038, [copyfile], page 1038, [ls], page 1069, [dir], page 1070.

`copyfile` *f1 f2*
`copyfile` *f1 f2 f*
`copyfile (`*f1*`, `*f2*`)`
`copyfile (`*f1*`, `*f2*`, 'f')`
`[`*status*`, `*msg*`, `*msgid*`] = copyfile (...)`

Copy the source file(s) or directory *f1* to the destination *f2*.

The name *f1* may contain globbing patterns, or may be a cell array of strings. If *f1* expands to multiple filenames, *f2* must be a directory.

When the force flag `'f'` is given any existing files will be overwritten without prompting.

If successful, *status* is logical 1, and *msg*, *msgid* are empty character strings (`""`). Otherwise, *status* is logical 0, *msg* contains a system-dependent error message, and *msgid* contains a unique message identifier. Note that the status code is exactly opposite that of the `system` command.

**See also:** [movefile], page 1038, [rename], page 1038, [unlink], page 1039, [delete], page 430, [glob], page 1044.

`unlink (`*`file`*`)`
`[`*`status`*`, `*`msg`*`] = unlink (`*`file`*`)`

Delete the file named *file*.

If successful, *status* is 0 and *msg* is an empty string. Otherwise, *status* is -1 and *msg* contains a system-dependent error message.

**See also:** [delete], page 430, [rmdir], page 1040.

`link `*`old new`*
`[`*`status`*`, `*`msg`*`] = link (`*`old`*`, `*`new`*`)`

Create a new link (also known as a hard link) to an existing file.

If successful, *status* is 0 and *msg* is an empty string. Otherwise, *status* is -1 and *msg* contains a system-dependent error message.

**See also:** [symlink], page 1039, [unlink], page 1039, [readlink], page 1039, [lstat], page 1041.

`symlink `*`old new`*
`[`*`status`*`, `*`msg`*`] = symlink (`*`old`*`, `*`new`*`)`

Create a symbolic link *new* which contains the string *old*.

If successful, *status* is 0 and *msg* is an empty string. Otherwise, *status* is -1 and *msg* contains a system-dependent error message.

**See also:** [link], page 1039, [unlink], page 1039, [readlink], page 1039, [lstat], page 1041.

*`result`*` = readlink `*`symlink`*
`[`*`result`*`, `*`err`*`, `*`msg`*`] = readlink (`*`symlink`*`)`

Read the value of the symbolic link *symlink*.

If successful, *result* contains the contents of the symbolic link *symlink*, *err* is 0, and *msg* is an empty string. Otherwise, *err* is nonzero and *msg* contains a system-dependent error message.

**See also:** [lstat], page 1041, [symlink], page 1039, [link], page 1039, [unlink], page 1039, [delete], page 430.

`mkdir `*`dirname`*
`mkdir `*`parent dirname`*
`mkdir (`*`dirname`*`)`

`mkdir (`*`parent`*`, `*`dirname`*`)`
`[`*`status`*`, `*`msg`*`, `*`msgid`*`] = mkdir (...)`

Create a directory named *dirname* in the directory *parent*, creating any intermediate directories if necessary.

If *dirname* is a relative path, and no *parent* directory is specified, then the present working directory is used.

If successful, *status* is logical 1, and *msg*, *msgid* are empty character strings (""). Otherwise, *status* is logical 0, *msg* contains a system-dependent error message, and *msgid* contains a unique message identifier. Note that the status code is exactly opposite that of the `system` command.

When creating a directory permissions will be set to `0777 - UMASK`.

**See also:** [rmdir], page 1040, [pwd], page 1070, [cd], page 1069, [umask], page 1041.

`rmdir `*`dir`*
`rmdir (`*`dir`*`, "`*`s`*`")`
`[`*`status`*`, `*`msg`*`, `*`msgid`*`] = rmdir (...)`

Remove the directory named *dir*.

If the optional second parameter is supplied with value `"s"`, recursively remove all subdirectories as well.

If successful, *status* is logical 1, and *msg*, *msgid* are empty character strings (""). Otherwise, *status* is logical 0, *msg* contains a system-dependent error message, and *msgid* contains a unique message identifier.

**See also:** [mkdir], page 1039, [confirm_recursive_rmdir], page 1040, [pwd], page 1070.

*`val`* `= confirm_recursive_rmdir ()`
*`old_val`* `= confirm_recursive_rmdir (`*`new_val`*`)`
*`old_val`* `= confirm_recursive_rmdir (`*`new_val`*`, "`*`local`*`")`

Query or set the internal variable that controls whether Octave will ask for confirmation before recursively removing a directory tree.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [rmdir], page 1040.

`mkfifo (`*`name`*`, `*`mode`*`)`
`[`*`status`*`, `*`msg`*`] = mkfifo (`*`name`*`, `*`mode`*`)`

Create a FIFO special file named *name* with file mode *mode*.

*mode* is interpreted as an octal number and is subject to umask processing. The final calculated mode is *`mode - umask`*.

If successful, *status* is 0 and *msg* is an empty string. Otherwise, *status* is -1 and *msg* contains a system-dependent error message.

**See also:** [pipe], page 1063, [umask], page 1041.

*oldmask* = umask (*mask*)

Set the permission mask for file creation.

The parameter *mask* is an integer, interpreted as an octal number.

If successful, returns the previous value of the mask (as an integer to be interpreted as an octal number); otherwise an error message is printed.

The permission mask is a UNIX concept used when creating new objects on a file system such as files, directories, or named FIFOs. The object to be created has base permissions in an octal number *mode* which are modified according to the octal value of *mask*. The final permissions for the new object are *mode - mask*.

**See also:** [fopen], page 311, [mkdir], page 1039, [mkfifo], page 1040.

[*info*, *err*, *msg*] = stat (*file*)
[*info*, *err*, *msg*] = stat (*fid*)
[*info*, *err*, *msg*] = lstat (*file*)
[*info*, *err*, *msg*] = lstat (*fid*)

Return a structure *info* containing the following information about *file* or file identifier *fid*.

`dev` ID of device containing a directory entry for this file.

`ino` File number of the file.

`mode` File mode, as an integer. Use the functions `S_ISREG`, `S_ISDIR`, `S_ISCHR`, `S_ISBLK`, `S_ISFIFO`, `S_ISLNK`, or `S_ISSOCK` to extract information from this value.

`modestr` File mode, as a string of ten letters or dashes as would be returned by *ls -l*.

`nlink` Number of links.

`uid` User ID of file's owner.

`gid` Group ID of file's group.

`rdev` ID of device for block or character special files.

`size` Size in bytes.

`atime` Time of last access in the same form as time values returned from `time`. See Section 36.1 [Timing Utilities], page 1025.

`mtime` Time of last modification in the same form as time values returned from `time`. See Section 36.1 [Timing Utilities], page 1025.

`ctime` Time of last file status change in the same form as time values returned from `time`. See Section 36.1 [Timing Utilities], page 1025.

`blksize` Size of blocks in the file.

`blocks` Number of blocks allocated for file.

If the call is successful *err* is 0 and *msg* is an empty string. If the file does not exist, or some other error occurs, *info* is an empty matrix, *err* is $-1$, and *msg* contains the corresponding system error message.

If *file* is a symbolic link, `stat` will return information about the actual file that is referenced by the link. Use `lstat` if you want information about the symbolic link itself.

For example:

```
[info, err, msg] = stat ("/vmlinuz")
  ⇒ info =
     {
       atime = 855399756
       rdev = 0
       ctime = 847219094
       uid = 0
       size = 389218
       blksize = 4096
       mtime = 847219094
       gid = 6
       nlink = 1
       blocks = 768
       mode = -rw-r--r--
       modestr = -rw-r--r--
       ino = 9316
       dev = 2049
     }
  ⇒ err = 0
  ⇒ msg =
```

**See also:** [lstat], page 1041, [ls], page 1069, [dir], page 1070, [isfile], page 1044, [isfolder], page 1044.

*tf* = `S_ISBLK` (*mode*)

Return true if *mode* corresponds to a block device.

The value of *mode* is assumed to be returned from a call to `stat`.

**See also:** [stat], page 1041, [lstat], page 1041.

*tf* = `S_ISCHR` (*mode*)

Return true if *mode* corresponds to a character device.

The value of *mode* is assumed to be returned from a call to `stat`.

**See also:** [stat], page 1041, [lstat], page 1041.

*tf* = `S_ISDIR` (*mode*)

Return true if *mode* corresponds to a directory.

The value of *mode* is assumed to be returned from a call to `stat`.

**See also:** [stat], page 1041, [lstat], page 1041.

*tf* = `S_ISFIFO` (*mode*)

Return true if *mode* corresponds to a fifo.

The value of *mode* is assumed to be returned from a call to `stat`.

**See also:** [stat], page 1041, [lstat], page 1041.

`tf = S_ISLNK (mode)`
: Return true if *mode* corresponds to a symbolic link.

  The value of *mode* is assumed to be returned from a call to `stat`.

  **See also:** [stat], page 1041, [lstat], page 1041.

`tf = S_ISREG (mode)`
: Return true if *mode* corresponds to a regular file.

  The value of *mode* is assumed to be returned from a call to `stat`.

  **See also:** [stat], page 1041, [lstat], page 1041.

`tf = S_ISSOCK (mode)`
: Return true if *mode* corresponds to a socket.

  The value of *mode* is assumed to be returned from a call to `stat`.

  **See also:** [stat], page 1041, [lstat], page 1041.

`fileattrib`
`fileattrib file`
`fileattrib (file)`
`[status, attrib] = fileattrib (...)`
`[status, msg, msgid] = fileattrib (...)`
: Report attribute information about *file*.

  If no file or directory is specified, report information about the present working directory.

  If successful, the output is a structure with the following fields:

  `Name` Full name of *file*.

  `archive` True if *file* is an archive (Windows).

  `system` True if *file* is a system file (Windows).

  `hidden` True if *file* is a hidden file (Windows).

  `directory`
  : True if *file* is a directory.

  `UserRead`
  `GroupRead`
  `OtherRead`
  : True if the user (group; other users) has read permission for *file*.

  `UserWrite`
  `GroupWrite`
  `OtherWrite`
  : True if the user (group; other users) has write permission for *file*.

  `UserExecute`
  `GroupExecute`
  `OtherExecute`
  : True if the user (group; other users) has execute permission for *file*.

If an attribute does not apply (e.g., archive on a Unix system) then the field is set to NaN.

If *file* contains globbing characters, information about all matching files is returned in a structure array.

If outputs are requested, the first is *status* which takes the value 1 when the operation was successful, and 0 otherwise. The second output contains the structure described above (*attrib*) if the operation was successful; otherwise, the second output is a system-dependent error message (*msg*). The third output is an empty string (`""`) when the operation was successful, or a unique message identifier (*msgid*) in the case of failure.

**See also:** [stat], page 1041, [glob], page 1044.

`tf = isfile (f)`

Return true if *f* is a regular file and false otherwise.

If *f* is a cell array of strings, *tf* is a logical array of the same size.

**See also:** [isfolder], page 1044, [exist], page 152, [stat], page 1041, [is_absolute_filename], page 1047, [is_rooted_relative_filename], page 1047.

`tf = isdir (f)`

This function is not recommended. Use `isfolder` or `file_in_loadpath` instead.

Return true if *f* is a directory and false otherwise.

Compatibility Note: The MATLAB function of the same name will also search for *f* in the load path directories. To emulate this behavior use

```
tf = ! isempty (file_in_loadpath (f))
```

**See also:** [isfolder], page 1044, [file_in_loadpath], page 229, [exist], page 152, [stat], page 1041, [is_absolute_filename], page 1047, [is_rooted_relative_filename], page 1047.

`tf = isfolder (f)`

Return true if *f* is a directory and false otherwise.

If *f* is a cell array of strings, *tf* is a logical array of the same size.

**See also:** [isfile], page 1044, [exist], page 152, [stat], page 1041, [is_absolute_filename], page 1047, [is_rooted_relative_filename], page 1047.

`files = readdir (dir)`
`[files, err, msg] = readdir (dir)`

Return the names of files in the directory *dir* as a cell array of strings.

If an error occurs, return an empty cell array in *files*. If successful, *err* is 0 and *msg* is an empty string. Otherwise, *err* is nonzero and *msg* contains a system-dependent error message.

**See also:** [ls], page 1069, [dir], page 1070, [glob], page 1044, [what], page 156.

`cstr = glob (pattern)`

Given an array of pattern strings (as a char array or a cell array) in *pattern*, return a cell array of filenames that match any of them, or an empty cell array if no patterns match.

The pattern strings are interpreted as filename globbing patterns (as they are used by Unix shells).

Within a pattern

`*` matches any string, including the null string,

`?` matches any single character, and

`[...]` matches any of the enclosed characters.

Tilde expansion is performed on each of the patterns before looking for matching filenames. For example:

```
ls
   ⇒
      file1  file2  file3  myfile1 myfile1b
glob ("*file1")
   ⇒
      {
        [1,1] = file1
        [2,1] = myfile1
      }
glob ("myfile?")
   ⇒
      {
        [1,1] = myfile1
      }
glob ("file[12]")
   ⇒
      {
        [1,1] = file1
        [2,1] = file2
      }
```

Note: On Windows, patterns that contain non-ASCII characters are not supported.

**See also:** [ls], page 1069, [dir], page 1070, [readdir], page 1044, [what], page 156.

*fname* = file_in_path (*path*, *file*)
*fname* = file_in_path (*path*, *file*, "*all*")

Return the absolute name of *file* if it can be found in *path*.

The value of *path* should be a colon-separated list of directories in the format described for `path`. If no file is found, return an empty character string. For example:

```
file_in_path (EXEC_PATH, "sh")
     ⇒ "/bin/sh"
```

If the second argument is a cell array of strings, search each directory of the path for element of the cell array and return the first that matches.

If the third optional argument `"all"` is supplied, return a cell array containing the list of all files that have the same name in the path. If no files are found, return an empty cell array.

**See also:** [file_in_loadpath], page 229, [dir_in_loadpath], page 230, [path], page 229.

*sep* = filesep ()
filesep ("*all*")

Return the system-dependent character used to separate directory names.

If "all" is given, the function returns all valid file separators in the form of a string. The list of file separators is system-dependent. It is '/' (forward slash) under UNIX or Mac OS X, '/' and '\' (forward and backward slashes) under Windows.

**See also:** [pathsep], page 229.

[*dir*, *name*, *ext*] = fileparts (*filename*)

Return the directory, name, and extension components of *filename*.

The input *filename* is a string which is parsed. There is no attempt to check whether the filename or directory specified actually exists.

**See also:** [fullfile], page 1046, [filesep], page 1046.

*filename* = fullfile (*dir1*, *dir2*, ..., *file*)

Build complete filename from separate parts.

The function joins any number of path components intelligently. The return value is the concatenation of each component with exactly one file separator between each part of the path and at most one leading and/or trailing file separator.

The input arguments might be strings or cell strings. Any input arguments that are cell strings must contain one single string or must be equal in size. In that case, the function returns a cell string of filepaths of the same dimensions as the input cell strings, e.g.:

```
fullfile ("/home/username", "data", {"f1.csv", "f2.csv", "f3.csv"})
  ⇒
      {
        [1,1] = /home/username/data/f1.csv
        [1,2] = /home/username/data/f2.csv
        [1,3] = /home/username/data/f3.csv
      }
```

On Windows systems, while forward slash file separators do work, they are replaced by backslashes. In addition, drive letters are stripped of leading file separators to obtain a valid file path.

Note: fullfile does not perform any validation of the resulting full filename.

**See also:** [fileparts], page 1046, [filesep], page 1046.

*newstr* = tilde_expand (*string*)
*newcstr* = tilde_expand (*cellstr*)

Perform tilde expansion on *string*.

If *string* begins with a tilde character, ('~'), all of the characters preceding the first slash (or all characters, if there is no slash) are treated as a possible user name, and the tilde and the following characters up to the slash are replaced by the home directory of the named user. If the tilde is followed immediately by a slash, the tilde is replaced by the home directory of the user running Octave.

If the input is a cell array of strings *cellstr* then tilde expansion is performed on each string element.

Examples:

```
tilde_expand ("~joeuser/bin")
     ⇒ "/home/joeuser/bin"
tilde_expand ("~/bin")
     ⇒ "/home/jwe/bin"
```

`[`*`cname`*`, `*`status`*`, `*`msg`*`] = canonicalize_file_name (`*`fname`*`)`

Return the canonical name of file *fname*.

If the file does not exist the empty string (`""`) is returned. No tilde expansion of *fname* is performed.

**See also:** [make_absolute_filename], page 1047, [is_absolute_filename], page 1047, [is_rooted_relative_filename], page 1047, [is_same_file], page 1047, [tilde_expand], page 1046.

*`abs_fname`*` = make_absolute_filename (`*`file`*`)`

Return the full name of *file* beginning from the root of the file system.

No check is done for the existence of *file*. No tilde expansion of *file* is performed.

**See also:** [canonicalize_file_name], page 1047, [is_absolute_filename], page 1047, [is_rooted_relative_filename], page 1047, [isfolder], page 1044, [tilde_expand], page 1046.

*`tf`*` = is_absolute_filename (`*`file`*`)`

Return true if *file* is an absolute filename.

**See also:** [is_rooted_relative_filename], page 1047, [make_absolute_filename], page 1047, [isfolder], page 1044.

*`same`*` = is_same_file (`*`filepath1`*`, `*`filepath2`*`)`

Return true if *filepath1* and *filepath2* refer to the same file.

If either *filepath1* or *filepath2* is a cell array of strings, then an array of the same size is returned, containing the values described above for every member of the cell array. The other argument may also be a cell array of strings (of the same size) or a string.

Programming Notes: Depending on the operating system and file system, the same file or folder can be referred to with different paths. In particular, paths on the Windows platform may differ in case (`file1` vs. `FILE1`), file separator ('\' vs. '/'), and format (`A~spaces.txt` (8.3 convention) vs. `A filename with spaces.txt`). This function returns true if the paths in *filepath1* and *filepath2* actually refer to the same file or folder, and false otherwise.

Note that unlike `strcmp`, this function requires that *filepath1* and *filepath2* exist, as well as point to the same location, in order to return true.

**See also:** [canonicalize_file_name], page 1047, [strcmp], page 88.

*`tf`*` = is_rooted_relative_filename (`*`file`*`)`

Return true if *file* is a rooted-relative filename.

**See also:** [is_absolute_filename], page 1047, [make_absolute_filename], page 1047, [isfolder], page 1044.

*val* = recycle ()
*old_val* = recycle (*new_val*)

Query or set the preference for recycling deleted files.

When recycling is enabled, commands which would permanently erase files instead move them to a temporary location (such as the directory labeled Trash).

Programming Note: This function is provided for MATLAB compatibility, but recycling is not implemented in Octave. To help avoid accidental data loss an error will be raised if an attempt is made to enable file recycling.

**See also:** [delete], page 430, [rmdir], page 1040.

## 36.3 File Archiving Utilities

bunzip2 (*bzfile*)
bunzip2 (*bzfile*, *dir*)
*filelist* = bunzip2 (...)

Unpack the bzip2 archive *bzfile*.

If *dir* is specified the files are unpacked in this directory rather than the one where *bzfile* is located.

The optional output *filelist* is a list of the uncompressed files.

**See also:** [bzip2], page 1050, [unpack], page 1050, [gunzip], page 1048, [unzip], page 1049, [untar], page 1049.

*filelist* = gzip (*files*)
*filelist* = gzip (*files*, *dir*)

Compress the list of files and directories specified in *files*.

*files* is a character array or cell array of strings. Shell wildcards in the filename such as '`*`' or '`?`' are accepted and expanded. Each file is compressed separately and a new file with a `".gz"` extension is created. The original files are not modified, but existing compressed files will be silently overwritten. If a directory is specified then `gzip` recursively compresses all files in the directory.

If *dir* is defined the compressed files are placed in this directory, rather than the original directory where the uncompressed file resides. Note that this does not replicate a directory tree in *dir* which may lead to files overwriting each other if there are multiple files with the same name.

If *dir* does not exist it is created.

The optional output *filelist* is a list of the compressed files.

**See also:** [gunzip], page 1048, [unpack], page 1050, [bzip2], page 1050, [zip], page 1049, [tar], page 1049.

gunzip (*gzfile*)
gunzip (*gzfile*, *outdir*)
*filelist* = gunzip (...)

Unpack the gzip archive *gzfile*.

If *gzfile* is a directory, all gzfiles in the directory will be recursively unpacked.

If *outdir* is specified the files are unpacked in this directory rather than the one where *gzfile* is located.

The optional output *filelist* is a list of the uncompressed files.

**See also:** [gzip], page 1048, [unpack], page 1050, [bunzip2], page 1048, [unzip], page 1049, [untar], page 1049.

`filelist = tar (tarfile, files)`
`filelist = tar (tarfile, files, rootdir)`

Pack the list of files and directories specified in *files* into the TAR archive *tarfile*.

*files* is a character array or cell array of strings. Shell wildcards in the filename such as '`*`' or '`?`' are accepted and expanded. Directories are recursively traversed and all files are added to the archive.

If *rootdir* is defined then any files without absolute pathnames are located relative to *rootdir* rather than the current directory.

The optional output *filelist* is a list of the files that were included in the archive.

**See also:** [untar], page 1049, [unpack], page 1050, [bzip2], page 1050, [gzip], page 1048, [zip], page 1049.

`untar (tarfile)`
`untar (tarfile, dir)`
`filelist = untar (...)`

Unpack the TAR archive *tarfile*.

If *dir* is specified the files are unpacked in this directory rather than the current directory.

The optional output *filelist* is a list of the uncompressed files.

**See also:** [tar], page 1049, [unpack], page 1050, [bunzip2], page 1048, [gunzip], page 1048, [unzip], page 1049.

`filelist = zip (zipfile, files)`
`filelist = zip (zipfile, files, rootdir)`

Compress the list of files and directories specified in *files* into the ZIP archive *zipfile*.

*files* is a character array or cell array of strings. Shell wildcards in the filename such as '`*`' or '`?`' are accepted and expanded. Directories are recursively traversed and all files are compressed and added to the archive.

If *rootdir* is defined then any files without absolute pathnames are located relative to *rootdir* rather than the current directory.

The optional output *filelist* is a list of the files that were included in the archive.

**See also:** [unzip], page 1049, [unpack], page 1050, [bzip2], page 1050, [gzip], page 1048, [tar], page 1049.

`unzip (zipfile)`
`unzip (zipfile, dir)`
`filelist = unzip (...)`

Unpack the ZIP archive *zipfile*.

If *dir* is specified the files are unpacked in this directory rather than the current directory.

The optional output *filelist* is a list of the uncompressed files.

**See also:** [zip], page 1049, [unpack], page 1050, [bunzip2], page 1048, [gunzip], page 1048, [untar], page 1049.

*files* = unpack (*file*)
*files* = unpack (*file*, *dir*)
*files* = unpack (*file*, *dir*, *filetype*)

Unpack the archive *file* based on its extension to the directory *dir*.

If *file* is a list of strings, then each file is unpacked individually. Shell wildcards in the filename such as '`*`' or '`?`' are accepted and expanded.

If *dir* is not specified or is empty (`[]`), it defaults to the current directory. If a directory is in the file list, then *filetype* must also be specified.

The specific archive filetype is inferred from the extension of the file. The *filetype* may also be specified directly using a string which corresponds to a known extension.

Valid filetype extensions:

`bz`
`bz2` bzip archive

`gz` gzip archive

`tar` tar archive

`tarbz`
`tarbz2`
`tbz`
`tbz2` tar + bzip archive

`targz`
`tgz` tar + gzip archive

`z` compress archive

`zip` zip archive

The optional return value is a list of *files* unpacked.

**See also:** [bunzip2], page 1048, [gunzip], page 1048, [unzip], page 1049, [untar], page 1049, [bzip2], page 1050, [gzip], page 1048, [zip], page 1049, [tar], page 1049.

*filelist* = bzip2 (*files*)
*filelist* = bzip2 (*files*, *dir*)

Compress the list of files specified in *files*.

*files* is a character array or cell array of strings. Shell wildcards in the filename such as '`*`' or '`?`' are accepted and expanded. Each file is compressed separately and a new file with a `".bz2"` extension is created. The original files are not modified, but existing compressed files will be silently overwritten.

If *dir* is defined the compressed files are placed in this directory, rather than the original directory where the uncompressed file resides. Note that this does not replicate

a directory tree in *dir* which may lead to files overwriting each other if there are multiple files with the same name.

If *dir* does not exist it is created.

The optional output *filelist* is a list of the compressed files.

**See also:** [bunzip2], page 1048, [unpack], page 1050, [gzip], page 1048, [zip], page 1049, [tar], page 1049.

## 36.4 Networking Utilities

*name* = `gethostname` ()

Return the hostname of the system where Octave is running.

### 36.4.1 FTP Objects

Octave supports the FTP protocol through an object-oriented interface. Use the function `ftp` to create an FTP object which represents the connection. All FTP functions take an FTP object as the first argument.

*f* = `ftp` (*host*)
*f* = `ftp` (*host*, *username*, *password*)

Connect to the FTP server *host* with *username* and *password*.

If *username* and *password* are not specified, user `"anonymous"` with no password is used. The returned FTP object *f* represents the established FTP connection.

The list of actions for an FTP object are shown below. All functions require an FTP object as the first argument.

| **Method** | **Description** |
|---|---|
| ascii | Set transfer type to ascii |
| binary | Set transfer type to binary |
| cd | Change remote working directory |
| close | Close FTP connection |
| delete | Delete remote file |
| dir | List remote directory contents |
| mget | Download remote files |
| mkdir | Create remote directory |
| mput | Upload local files |
| rename | Rename remote file or directory |
| rmdir | Remove remote directory |

**See also:** [@ftp/ascii], page 1053, [@ftp/binary], page 1053, [@ftp/cd], page 1052, [@ftp/close], page 1051, [@ftp/delete], page 1053, [@ftp/dir], page 1052, [@ftp/mget], page 1052, [@ftp/mkdir], page 1053, [@ftp/mput], page 1052, [@ftp/rename], page 1053, [@ftp/rmdir], page 1053.

`close` (*f*)

Close the FTP connection represented by the FTP object *f*.

*f* is an FTP object returned by the `ftp` function.

**See also:** [@ftp/ftp], page 1051.

`mget (`*`f`*`, `*`file`*`)`
`mget (`*`f`*`, `*`dir`*`)`
`mget (`*`f`*`, `*`remote_name`*`, `*`target`*`)`

Download a remote file *file* or directory *dir* to the local directory on the FTP connection *f*.

*f* is an FTP object returned by the `ftp` function.

The arguments *file* and *dir* can include wildcards and any files or directories on the remote server that match will be downloaded.

If a third string argument *target* is given, then it must indicate the path to the local destination directory. *target* may be a relative or absolute path.

**See also:** [@ftp/mput], page 1052, [@ftp/ftp], page 1051.

`mput (`*`f`*`, `*`file`*`)`
*`file_list`* `= mput (`*`f`*`, `*`file`*`)`

Upload the local file *file* into the current remote directory on the FTP connection *f*.

*f* is an FTP object returned by the `ftp` function.

The argument *file* is passed through the `glob` function and any files that match the wildcards in *file* will be uploaded.

The optional output argument *file_list* contains a cell array of strings with the names of the uploaded files.

**See also:** [@ftp/mget], page 1052, [@ftp/mkdir], page 1053, [@ftp/ftp], page 1051.

*`cwd`* `= cd (`*`f`*`)`
`cd (`*`f`*`, `*`path`*`)`
*`new_cwd`* `= cd (`*`f`*`, `*`path`*`)`

Get or set the remote directory on the FTP connection *f*.

*f* is an FTP object returned by the `ftp` function.

If *path* is not specified, return the remote current working directory. Otherwise, set the remote directory to *path* and return the new remote working directory.

If the directory does not exist, an error message is printed and the working directory is not changed.

**See also:** [@ftp/dir], page 1052, [@ftp/ftp], page 1051.

`dir (`*`f`*`)`
*`lst`* `= dir (`*`f`*`)`

List the current directory in verbose form for the FTP connection *f*.

*f* is an FTP object returned by the `ftp` function.

If the optional output *lst* is requested return a struct array with one entry per file with the fields `name`, `date`, `bytes`, `isdir`, `datenum`.

**See also:** [@ftp/cd], page 1052, [@ftp/mkdir], page 1053, [@ftp/rmdir], page 1053, [@ftp/ftp], page 1051.

`ascii (f)`
: Set the FTP connection *f* to use ASCII mode for transfers.

  ASCII mode is only appropriate for text files as it will convert the remote host's newline representation to the local host's newline representation.

  *f* is an FTP object returned by the `ftp` function.

  **See also:** [@ftp/binary], page 1053, [@ftp/ftp], page 1051.

`binary (f)`
: Set the FTP connection *f* to use binary mode for transfers.

  In binary mode there is no conversion of newlines from the remote representation to the local representation.

  *f* is an FTP object returned by the `ftp` function.

  **See also:** [@ftp/ascii], page 1053, [@ftp/ftp], page 1051.

`delete (f, file)`
: Delete the remote file *file* over the FTP connection *f*.

  *f* is an FTP object returned by the `ftp` function.

  **See also:** [@ftp/rmdir], page 1053, [@ftp/rename], page 1053, [@ftp/ftp], page 1051.

`rename (f, oldname, newname)`
: Rename or move the remote file or directory *oldname* to *newname*, over the FTP connection *f*.

  *f* is an FTP object returned by the `ftp` function.

  **See also:** [@ftp/delete], page 1053, [@ftp/rmdir], page 1053, [@ftp/ftp], page 1051.

`mkdir (f, path)`
: Create the remote directory *path*, over the FTP connection *f*.

  *f* is an FTP object returned by the `ftp` function.

  **See also:** [@ftp/rmdir], page 1053, [@ftp/ftp], page 1051.

`rmdir (f, path)`
: Remove the remote directory *path*, over the FTP connection *f*.

  *f* is an FTP object returned by the `ftp` function.

  **See also:** [@ftp/delete], page 1053, [@ftp/mkdir], page 1053, [@ftp/rename], page 1053, [@ftp/ftp], page 1051.

### 36.4.2 WWW Access

Octave can communicate with websites across the Internet. The `web` function will launch an external web browser to interactively view a site. The remaining functions—`urlread`, `urlwrite`, `webread`, `webwrite`—are internal Octave functions which can import or export data to/from Octave and a website identified by a URL (Uniform Resource Locator).

*status* = web ()
*status* = web (*url*)
*status* = web (*url*, *option*)
*status* = web (*url*, *option_1*, ..., *option_N*)
[*status*, *h*, *url*] = web (...)

Open *url* in the default system web browser.

With no arguments given, the address https://www.octave.org is opened.

Additional options can be passed for MATLAB compatibility, but are ignored.

- '-browser' Open *url* in the default system browser.
- '-new' No effect on the system browser.
- '-noaddressbox' No effect on the system browser.
- '-notoolbar' No effect on the system browser.

The return value *status* has one of the values:

- '0' Found and opened system browser successfully.
- '1' Cannot find the system browser.
- '2' System browser found, but an error occurred.

The return values *handle* and *url* are currently unimplemented but given for compatibility.

**See also:** [weboptions], page 1056, [webread], page 1056, [webwrite], page 1056, [urlread], page 1054, [urlwrite], page 1055.

*s* = urlread (*url*)
*s* = urlread (*url*, *Name*, *Value*, ...)
[*s*, *success*] = urlread (*url*, ...)
[*s*, *success*, *message*] = urlread (*url*, ...)

Download a remote file specified by its *url* and return its content in string *s*.

For example:

```
s = urlread ("http://ftp.octave.org/pub/README");
```

The variable *success* is 1 if the download was successful, otherwise it is 0 in which case *message* contains an error message.

If no output argument is specified and an error occurs, then the error is signaled through Octave's error handling mechanism.

This function uses libcurl. The curl library supports, among others, the HTTP, FTP, and FILE protocols. Username and password may be specified in the URL. For example:

```
s = urlread ("http://user:password@example.com/file.txt");
```

Additional options can be specified using property/value pairs:

| | |
|---|---|
| Get | Cell array of name/value pairs for GET request parameters. |
| Post | Cell array of name/value pairs for POST request parameters. |
| Timeout | Timeout value in seconds. |

`UserAgent`
: User agent string for the request.

`Username`
: Username for authentication.

`Password`
: Password for authentication.

`Charset`
: Character encoding (stored but not currently used).

Example with property/value pairs:

```
s = urlread ("http://www.example.com/data",
             "Get", {"term", "octave"}, "Timeout", 10);
```

For backward compatibility, the old calling form is also supported:

```
s = urlread ("http://www.google.com/search",
             "get", {"query", "octave"});
```

**See also:** [urlwrite], page 1055, [weboptions], page 1056.

`urlwrite (`*`url`*`, `*`localfile`*`)`
`urlwrite (`*`url`*`, `*`localfile`*`, `*`Name`*`, `*`Value`*`, ...)`
*`f`* `= urlwrite (`*`url`*`, `*`localfile`*`, ...)`
`[`*`f`*`, `*`success`*`] = urlwrite (`*`url`*`, `*`localfile`*`, ...)`
`[`*`f`*`, `*`success`*`, `*`message`*`] = urlwrite (`*`url`*`, `*`localfile`*`, ...)`

Download a remote file specified by its *url* and save it as *localfile*.

For example:

```
urlwrite ("http://ftp.octave.org/pub/README",
          "README.txt");
```

The full path of the downloaded file is returned in *f*.

The variable *success* is 1 if the download was successful, otherwise it is 0 in which case *message* contains an error message.

If no output argument is specified and an error occurs, then the error is signaled through Octave's error handling mechanism.

This function uses libcurl. The curl library supports, among others, the HTTP, FTP, and FILE protocols. Username and password may be specified in the URL, for example:

```
urlwrite ("http://username:password@example.com/file.txt",
          "file.txt");
```

Additional options can be specified using property/value pairs:

`Get`
: Cell array of name/value pairs for GET request parameters.

`Post`
: Cell array of name/value pairs for POST request parameters.

`Timeout`
: Timeout value in seconds.

`UserAgent`
: User agent string for the request.

`Username`
: Username for authentication.

`Password`
: Password for authentication.

Example with property/value pairs:

```
urlwrite ("http://www.example.com/data.txt", "data.txt",
          "Timeout", 10, "UserAgent", "Octave");
```

For backward compatibility, the old calling form is also supported:

```
urlwrite ("http://www.google.com/search", "search.html",
          "get", {"query", "octave"});
```

**See also:** [urlread], page 1054, [weboptions], page 1056.

*response* = webread (*url*)
*response* = webread (*url*, *name1*, *value1*, ...)
*response* = webread (..., *options*)

Read content from RESTful web service.

Read content from the web service specified by *url* and return the content in *response*.

All name-value pairs given (*name1*, *value1*, ...) are appended as query parameters to *url*. To place a query in the body of the message, use `webwrite`. The web service defines the acceptable query parameters.

*options* is a `weboptions` object that may be used to add other HTTP request options. This argument can be used with either calling form. See `help weboptions` for a complete list of supported HTTP options.

**See also:** [weboptions], page 1056, [webwrite], page 1056.

*response* = webwrite (*url*, *name1*, *value1*, ...)
*response* = webwrite (*url*, *data*)
*response* = webwrite (..., *options*)

Write data to RESTful web services.

Write content to the web service specified by *url* and return the response in *response*.

All name-value pairs given (*name1*, *value1*, ...) are added as pairs of query parameters to the body of request method (`get`, `post`, `put`, etc.).

*options* is a `weboptions` object that may be used to add other HTTP request options. This argument can be used with either calling form. See `help weboptions` for a complete list of supported HTTP options.

**See also:** [weboptions], page 1056, [webread], page 1056.

*options* = weboptions ()
*options* = weboptions (*name1*, *value1*, ...)

Specify parameters for RESTful web services.

When called with no inputs return a default `weboptions` object to specify parameters for a request to a web service. A `weboptions` object is an optional input argument to the `webread` and `webwrite` functions.

Multiple name and value pair arguments may be specified in any order as *name1*, *value1*, *name2*, *value2*, etc.

The option names must match **exactly** one of those specified in the table below.

The following options are available:

- '`CharacterEncoding`' — Specify the character encoding of the data:

  '`auto`' (default), '`UTF-8`', '`US-ASCII`'. '`auto`' chooses an encoding based on the content-type of the data.
- '`UserAgent`' — Specify the User Agent for the connection.

  Default value is '`Octave/VERSION`', where '`VERSION`' is the current version of Octave as returned by `version`.
- '`Timeout`' — Specify the timeout value for the connection in seconds.

  Default is 5 seconds. The special value '`Inf`' sets the timeout to the maximum value of 2147.483647 seconds.
- '`Username`' — User identifier for a basic HTTP connection.

  Default is `""`. It must be a string or character vector.
- '`Password`' — User authentication password for HTTP connection.

  Default is `""`. It must be a string or character vector.

  Programming Note: If you display a `weboption` object with the Password property set, the value is displayed as a string containing `'*'`. However, the object stores the value of the Password property as plain text.
- '`KeyName`' — Specify the name of an additional key to be added to the HTTP request header.

  It must be a string or character vector. It should be coupled with '`KeyValue`'.
- '`KeyValue`' — Specify the value of the key '`KeyName`'.

  '`KeyName`' must already be assigned in order to specify this field.
- '`ContentType`' — Specify the content type of the data.

  The following values are available: '`auto`', '`text`', '`json`'

  Default is '`auto`'. It automatically determines the content type. All other formats like '`audio`', '`binary`', etc. available in MATLAB are not currently supported.
- '`ContentReader`' — Not yet implemented. Only for MATLAB compatibility.
- '`MediaType`' — Not yet implemented. Only for MATLAB compatibility.
- '`RequestMethod`' — Specifies the type of request to be made.

  The following methods are available: '`get`', '`put`', '`post`', '`delete`', '`patch`'

  `webread` uses the HTTP GET method. `webwrite` uses the HTTP POST method as default.
- '`ArrayFormat`' – Not yet implemented. Only for MATLAB compatibility.
- '`HeaderFields`' — Specify header fields for the connection.

  Names and values of header fields, specified as an m-by-2 cell array of strings, to add to the HTTP request header. `HeaderFields{i,1}` is the name of a field and `HeaderFields{i,2}` is its value.

  ```
  weboptions ("HeaderFields",
              {"Content-Length", "78" ;
               "Content-Type", "application/json"})
  ```

  creates a weboptions object that contains two header fields: `Content-Length` with value `78` and `Content-Type` with value `application/json`.

- '`CertificateFilename`' — Not yet implemented. Only for MATLAB compatibility.

**See also:** [webread], page 1056, [webwrite], page 1056.

### 36.4.3 Base64 and Binary Data Transmission

Some transmission channels can not accept binary data. It is customary to encode binary data in Base64 for transmission and to decode the data upon reception.

`s = base64_encode (x)`
Encode a double matrix or array *x* into the base64 format string *s*.

**See also:** [base64_decode], page 1058, [matlab.net.base64decode], page 1058, [matlab.net.base64encode], page 1058.

`x = base64_decode (s)`
`x = base64_decode (s, dims)`
Decode the double matrix or array *x* from the base64 encoded string *s*.

The optional input parameter *dims* should be a vector containing the dimensions of the decoded array.

**See also:** [base64_encode], page 1058, [matlab.net.base64decode], page 1058, [matlab.net.base64encode], page 1058.

`b64_str = matlab.net.base64encode (in)`
Convert *in* to a base64 encoded string *b64_str*.

The input *in* can be a string or numeric vector. The output *b64_str* will be encoded according to RFC 4648.

**See also:** [matlab.net.base64decode], page 1058, [base64_decode], page 1058, [base64_encode], page 1058, [unicode2native], page 99.

`out_vec = matlab.net.base64decode (b64_str)`
Convert base64 encoded *b64_str* to uint8 vector *out_vec*.

The input *b64_str* must be a string vector. The output *out_vec* will be a uint8 vector that is decoded according to RFC 4648.

**See also:** [matlab.net.base64encode], page 1058, [base64_decode], page 1058, [base64_encode], page 1058, [native2unicode], page 99.

## 36.5 Controlling Subprocesses

Octave includes some high-level commands like `system` and `popen` for starting subprocesses. If you want to run another program to perform some task and then look at its output, you will probably want to use these functions.

Octave also provides several very low-level Unix-like functions which can also be used for starting subprocesses, but you should probably only use them if you can't find any way to do what you need with the higher-level functions.

`system ("string")`
`system ("string", return_output)`
`system ("string", return_output, type)`
`[status, output] = system (...)`

Execute a shell command specified by *string*.

If *system* is called with one or more output arguments, or if the optional argument *return_output* is true and the subprocess is started synchronously, then the output from the command is returned as a variable. Otherwise, if the subprocess is executed synchronously, its output is sent to the standard output. To send the output of a command executed with `system` through the pager, use a command like

```
[~, text] = system ("cmd");
more on;
disp (text);
```

or

```
more on;
printf ("%s\n", nthargout (2, "system", "cmd"));
```

If the optional argument *type* is `"async"`, the process is started in the background and the process ID of the child process is returned immediately. Otherwise, the child process is started and Octave waits until it exits. If the *type* argument is omitted, it defaults to the value `"sync"`.

The `system` function can return two values. The first is the exit status of the command and the second is any output from the command that was written to the standard output stream. For example,

```
[status, output] = system ("echo foo & exit 2");
```

will set the variable *output* to the string '`foo`', and the variable *status* to the integer '`2`'.

For commands run asynchronously, *status* is the process id of the command shell that is started to run the command.

The shell used for executing commands varies with operating system and is typically `/bin/sh` for UNIX systems and `cmd.exe` for Windows systems.

**See also:** [unix], page 1059, [dos], page 1060.

`unix ("command")`
`status = unix ("command")`
`[status, text] = unix ("command")`
`[...] = unix ("command", "-echo")`

Execute a system command if running under a Unix-like operating system, otherwise do nothing.

Octave waits for the external command to finish before returning the exit status of the program in *status* and any output in *text*.

When called with no output argument, or the `"-echo"` argument is given, then *text* is also sent to standard output.

**See also:** [dos], page 1060, [system], page 1058, [isunix], page 1074, [ismac], page 1074, [ispc], page 1074.

`dos ("command")`
`status = dos ("command")`
`[status, text] = dos ("command")`
`[...] = dos ("command", "-echo")`

Execute a system command if running under a Windows-like operating system, otherwise do nothing.

Octave waits for the external command to finish before returning the exit status of the program in *status* and any output in *text*.

When called with no output argument, or the `"-echo"` argument is given, then *text* is also sent to standard output.

**See also:** [unix], page 1059, [system], page 1058, [isunix], page 1074, [ismac], page 1074, [ispc], page 1074.

`open file`
`output = open (file)`

Open the file *file* in Octave or in an external application based on the file type as determined by the filename extension.

By default, recognized file types are

`.m` Open file in the editor. No *output* value is returned.

`.mat`
`octave-workspace`
Open the data file with `load`. If no return value *output* is requested, variables are loaded in the base workspace. Otherwise *output* will be a structure containing loaded data. See [load function], page 298.

`.ofig` Open the figure with `hgload`. See [hgload function], page 445.

`.fig, .ofig`
Load the figure

`.exe` Execute the program (on Windows systems only). No *output* value is returned.

Custom file extensions may also be handled if a function `openxxx`, where `xxx` is the extension, is found in the load path. The function must accept the file name as input. For example, in order to load `".dat"` data files in the base workspace, as is done by default for `".mat"` files, one may define `"opendat.m"` with the following contents:

```
function retval = opendat (fname)
  evalin ("base", sprintf ("load ('%s');", fname));
endfunction
```

Other file types are opened in the appropriate external application.

`output = perl (scriptfile)`
`output = perl (scriptfile, argument1, argument2, ...)`
`[output, status] = perl (...)`

Invoke Perl script *scriptfile*, possibly with a list of command line arguments.

Return output in *output* and optional status in *status*. If *scriptfile* is not an absolute filename it is searched for in the current directory and then in the Octave loadpath.

**See also:** [system], page 1058, [python], page 1061.

*output* = python (*scriptfile*)
*output* = python (*scriptfile*, *argument1*, *argument2*, ...)
[*output*, *status*] = python (...)

Invoke Python script *scriptfile*, possibly with a list of command line arguments.

Return output in *output* and optional status in *status*. If *scriptfile* is not an absolute filename it is searched for in the current directory and then in the Octave loadpath.

Programming Note: On UNIX systems, the script will be executed by `python3` and on Windows by `python`. You can override these defaults by setting the `PYTHON` environment variable, for example from within Octave using `setenv PYTHON /usr/local/bin/python3`.

**See also:** [system], page 1058, [perl], page 1060.

*fid* = popen (*command*, *mode*)

Start a process and create a pipe.

The name of the command to run is given by *command*. The argument *mode* may be

`"r"` The pipe will be connected to the standard output of the process, and open for reading.

`"w"` The pipe will be connected to the standard input of the process, and open for writing.

The file identifier corresponding to the input or output stream of the process is returned in *fid*.

For example:

```
fid = popen ("ls -ltr / | tail -3", "r");
while (ischar (s = fgets (fid)))
  fputs (stdout, s);
endwhile
```

```
-| drwxr-xr-x  33 root  root  3072 Feb 15 13:28 etc
-| drwxr-xr-x   3 root  root  1024 Feb 15 13:28 lib
-| drwxrwxrwt  15 root  root  2048 Feb 17 14:53 tmp
```

**See also:** [popen2], page 1061.

*status* = pclose (*fid*)

Close a file identifier *fid* that was opened by `popen`.

If successful, `fclose` returns 0, otherwise, it returns -1.

Programming Note: The function `fclose` may also be used for the same purpose.

**See also:** [fclose], page 313, [popen], page 1061.

[*in*, *out*, *pid*] = popen2 (*command*, *args*)

Start a subprocess with two-way communication.

The name of the process is given by *command*, and *args* is an array or cell array of strings containing options for the command.

The file identifiers for the input and output streams of the subprocess are returned in *in* and *out*. If execution of the command is successful, *pid* contains the process ID of the subprocess. Otherwise, *pid* is −1.

For example:

```
[in, out, pid] = popen2 ("sort", "-r");
fputs (in, "these\nare\nsome\nstrings\n");
fclose (in);
EAGAIN = errno ("EAGAIN");
done = false;
do
  s = fgets (out);
  if (ischar (s))
    fputs (stdout, s);
  elseif (errno () == EAGAIN)
    pause (0.1);
    fclear (out);
  else
    done = true;
  endif
until (done)
fclose (out);
waitpid (pid);
```

```
⊣ these
⊣ strings
⊣ some
⊣ are
```

Note that `popen2`, unlike `popen`, will not "reap" the child process. If you don't use `waitpid` to check the child's exit status, it will linger until Octave exits.

**See also:** [popen], page 1061, [waitpid], page 1064.

*val* = EXEC_PATH ()
*old_val* = EXEC_PATH (*new_val*)
*old_val* = EXEC_PATH (*new_val*, "*local*")

Query or set the internal variable that specifies a colon separated list of directories to append to the shell PATH when executing external programs.

The initial value of is taken from the environment variable `OCTAVE_EXEC_PATH`, but that value can be overridden by the command line argument `--exec-path PATH`.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [IMAGE_PATH], page 941, [OCTAVE_HOME], page 1074, [OCTAVE_EXEC_HOME], page 1074.

In most cases, the following functions simply decode their arguments and make the corresponding Unix system calls. For a complete example of how they can be used, look at the definition of the function `popen2`.

`[pid, msg] = fork ()`

Create a copy of the current process.

Fork can return one of the following values:

`> 0` You are in the parent process. The value returned from `fork` is the process id of the child process. You should probably arrange to wait for any child processes to exit by using `waitpid`.

`0` You are in the child process. You can call `exec` to start another process. If that fails, you should probably call `_Exit` to terminate the child.

`< 0` The call to `fork` failed for some reason. You must take evasive action. A system-dependent error message will be waiting in *msg*.

**See also:** [exec], page 1063, [_Exit], page 1063.

`_Exit ()`
`_Exit (status)`

Exit the currently running process with exit code *status*.

If called with no arguments, exit with status `0` indicating success.

The optional integer argument *status* specifies the exit code.

Programming Note: This function maps to the C++ function `quick_exit`. The calling process is stopped and any open file descriptors are closed. This is the correct function to call to end a child process started from Octave. The ordinary C library function 'exit' will not work.

**See also:** [fork], page 1063.

`[err, msg] = exec (file, args)`

Replace current process with a new process.

Calling `exec` without first calling `fork` will terminate your current Octave process and replace it with the program named by *file*. For example,

```
exec ("ls", "-l")
```

will run `ls` and return you to your shell prompt.

If successful, `exec` does not return. If `exec` does return, *err* will be nonzero, and *msg* will contain a system-dependent error message.

`[read_fd, write_fd, err, msg] = pipe ()`

Create a pipe and return the reading and writing ends of the pipe into *read_fd* and *write_fd* respectively.

If successful, *err* is 0 and *msg* is an empty string. Otherwise, *err* is nonzero and *msg* contains a system-dependent error message.

**See also:** [mkfifo], page 1040.

`[`*fid*`, `*msg*`] = dup2 (`*old*`, `*new*`)`

Duplicate a file descriptor.

If successful, *fid* is greater than zero and contains the new file ID. Otherwise, *fid* is negative and *msg* contains a system-dependent error message.

**See also:** [fopen], page 311, [fclose], page 313, [fcntl], page 1066.

`[`*pid*`, `*status*`, `*msg*`] = waitpid (`*pid*`, `*options*`)`

Wait for process *pid* to terminate.

The *pid* argument can be:

−1 Wait for any child process.

0 Wait for any child process whose process group ID is equal to that of the Octave interpreter process.

> 0 Wait for termination of the child process with ID *pid*.

The *options* argument can be a bitwise OR of zero or more of the following constants:

`0` Wait until signal is received or a child process exits (this is the default if the *options* argument is missing).

`WNOHANG` Do not hang if status is not immediately available.

`WUNTRACED`
Report the status of any child processes that are stopped, and whose status has not yet been reported since they stopped.

`WCONTINUE`
Return if a stopped child has been resumed by delivery of `SIGCONT`. This value may not be meaningful on all systems.

If the returned value of *pid* is greater than 0, it is the process ID of the child process that exited. If an error occurs, *pid* will be less than zero and *msg* will contain a system-dependent error message. The value of *status* contains additional system-dependent information about the subprocess that exited.

**See also:** [WCONTINUE], page 1064, [WCOREDUMP], page 1064, [WEXITSTATUS], page 1065, [WIFCONTINUED], page 1065, [WIFSIGNALED], page 1065, [WIFSTOPPED], page 1065, [WNOHANG], page 1066, [WSTOPSIG], page 1066, [WTERMSIG], page 1066, [WUNTRACED], page 1066.

*v* `= WCONTINUE ()`

Return the numerical value of the `WCONTINUE` macro.

`WCONTINUE` is the option argument that may be passed to `waitpid` to indicate that it should also return if a stopped child has been resumed by delivery of a `SIGCONT` signal.

**See also:** [waitpid], page 1064, [WNOHANG], page 1066, [WUNTRACED], page 1066.

*tf* `= WCOREDUMP (`*status*`)`

Given *status* from a call to `waitpid`, return true if the child produced a core dump.

This function should only be employed if `WIFSIGNALED` returned true. The macro used to implement this function is not specified in POSIX.1-2001 and is not available on some Unix implementations (e.g., AIX, SunOS).

**See also:** [waitpid], page 1064, [WIFEXITED], page 1065, [WEXITSTATUS], page 1065, [WIFSIGNALED], page 1065, [WTERMSIG], page 1066, [WIFSTOPPED], page 1065, [WSTOPSIG], page 1066, [WIFCONTINUED], page 1065.

*tf* = `WEXITSTATUS` (*status*)

Given *status* from a call to `waitpid`, return the exit status of the child.

This function should only be employed if `WIFEXITED` returned true.

**See also:** [waitpid], page 1064, [WIFEXITED], page 1065, [WIFSIGNALED], page 1065, [WTERMSIG], page 1066, [WCOREDUMP], page 1064, [WIFSTOPPED], page 1065, [WSTOPSIG], page 1066, [WIFCONTINUED], page 1065.

*tf* = `WIFCONTINUED` (*status*)

Given *status* from a call to `waitpid`, return true if the child process was resumed by delivery of `SIGCONT`.

**See also:** [waitpid], page 1064, [WIFEXITED], page 1065, [WEXITSTATUS], page 1065, [WIFSIGNALED], page 1065, [WTERMSIG], page 1066, [WCOREDUMP], page 1064, [WIFSTOPPED], page 1065, [WSTOPSIG], page 1066.

*tf* = `WIFSIGNALED` (*status*)

Given *status* from a call to `waitpid`, return true if the child process was terminated by a signal.

**See also:** [waitpid], page 1064, [WIFEXITED], page 1065, [WEXITSTATUS], page 1065, [WTERMSIG], page 1066, [WCOREDUMP], page 1064, [WIFSTOPPED], page 1065, [WSTOPSIG], page 1066, [WIFCONTINUED], page 1065.

*tf* = `WIFSTOPPED` (*status*)

Given *status* from a call to `waitpid`, return true if the child process was stopped by delivery of a signal.

This is only possible if the call was done using `WUNTRACED` or when the child is being traced (see ptrace(2)).

**See also:** [waitpid], page 1064, [WIFEXITED], page 1065, [WEXITSTATUS], page 1065, [WIFSIGNALED], page 1065, [WTERMSIG], page 1066, [WCOREDUMP], page 1064, [WSTOPSIG], page 1066, [WIFCONTINUED], page 1065.

*tf* = `WIFEXITED` (*status*)

Given *status* from a call to `waitpid`, return true if the child terminated normally.

**See also:** [waitpid], page 1064, [WEXITSTATUS], page 1065, [WIFSIGNALED], page 1065, [WTERMSIG], page 1066, [WCOREDUMP], page 1064, [WIFSTOPPED], page 1065, [WSTOPSIG], page 1066, [WIFCONTINUED], page 1065.

`v = WNOHANG ()`

Return the numerical value of the `WNOHANG` macro.

`WNOHANG` is the option argument that may be passed to `waitpid` to indicate that it should return its status immediately instead of waiting for a process to exit.

**See also:** [waitpid], page 1064, [WUNTRACED], page 1066, [WCONTINUE], page 1064.

`tf = WSTOPSIG (status)`

Given *status* from a call to `waitpid`, return the number of the signal which caused the child to stop.

This function should only be employed if `WIFSTOPPED` returned true.

**See also:** [waitpid], page 1064, [WIFEXITED], page 1065, [WEXITSTATUS], page 1065, [WIFSIGNALED], page 1065, [WTERMSIG], page 1066, [WCOREDUMP], page 1064, [WIFSTOPPED], page 1065, [WIFCONTINUED], page 1065.

`tf = WTERMSIG (status)`

Given *status* from a call to `waitpid`, return the number of the signal that caused the child process to terminate.

This function should only be employed if `WIFSIGNALED` returned true.

**See also:** [waitpid], page 1064, [WIFEXITED], page 1065, [WEXITSTATUS], page 1065, [WIFSIGNALED], page 1065, [WCOREDUMP], page 1064, [WIFSTOPPED], page 1065, [WSTOPSIG], page 1066, [WIFCONTINUED], page 1065.

`v = WUNTRACED ()`

Return the numerical value of the `WUNTRACED` macro.

`WUNTRACED` is the option argument that may be passed to `waitpid` to indicate that it should also return if the child process has stopped but is not traced via the `ptrace` system call

**See also:** [waitpid], page 1064, [WNOHANG], page 1066, [WCONTINUE], page 1064.

`fcntl (fid, request, arg)`
`[status, msg] = fcntl (fid, request, arg)`

Change the properties of the open file *fid*.

The following values may be passed as *request*:

`F_DUPFD` Return a duplicate file descriptor.

`F_GETFD` Return the file descriptor flags for *fid*.

`F_SETFD` Set the file descriptor flags for *fid*.

`F_GETFL` Return the file status flags for *fid*. The following codes may be returned (some of the flags may be undefined on some systems).

`O_RDONLY` Open for reading only.

`O_WRONLY` Open for writing only.

`O_RDWR` Open for reading and writing.

`O_APPEND` Append on each write.

`O_CREAT` Create the file if it does not exist.

`O_NONBLOCK`
Non-blocking mode.

`O_SYNC` Wait for writes to complete.

`O_ASYNC` Asynchronous I/O.

`F_SETFL` Set the file status flags for *fid* to the value specified by *arg*. The only flags that can be changed are `O_APPEND` and `O_NONBLOCK`.

If successful, *status* is 0 and *msg* is an empty string. Otherwise, *status* is -1 and *msg* contains a system-dependent error message.

**See also:** [fopen], page 311, [dup2], page 1064.

`kill (`*`pid`*`, `*`sig`*`)`
`[`*`status`*`, `*`msg`*`] = kill (`*`pid`*`, `*`sig`*`)`
Send signal *sig* to process *pid*.

If *pid* is positive, then signal *sig* is sent to *pid*.

If *pid* is 0, then signal *sig* is sent to every process in the process group of the current process.

If *pid* is -1, then signal *sig* is sent to every process except process 1.

If *pid* is less than -1, then signal *sig* is sent to every process in the process group -*pid*.

If *sig* is 0, then no signal is sent, but error checking is still performed.

If successful, *status* is 0 and *msg* is an empty string. Otherwise, *status* is -1 and *msg* contains a system-dependent error message.

`S = SIG ()`
Return a structure containing Unix signal names and their defined values.

## 36.6 Process, Group, and User IDs

`pgid = getpgrp ()`
Return the process group id of the current process.

`pid = getpid ()`
Return the process id of the current process.

**See also:** [getppid], page 1067.

`pid = getppid ()`
Return the process id of the parent process.

**See also:** [getpid], page 1067.

`euid = geteuid ()`
Return the effective user id of the current process.

**See also:** [getuid], page 1068.

`uid = getuid ()`
: Return the real user id of the current process.

  **See also:** [geteuid], page 1067.

`egid = getegid ()`
: Return the effective group id of the current process.

  **See also:** [getgid], page 1068.

`gid = getgid ()`
: Return the real group id of the current process.

  **See also:** [getegid], page 1068.

## 36.7 Environment Variables

*`val`* `= getenv ("`*`var`*`")`
: Return the value of the environment variable *var*.

  For example,

```
getenv ("PATH")
```

  returns a string containing the value of your path.

  **See also:** [setenv], page 1068, [unsetenv], page 1068, [isenv], page 1068.

*`tf`* `= isenv (`*`var`*`)`
: Return true if the variable *var* is an environment variable, and false otherwise.

  For example,

```
tf = isenv ("PATH")
```

  will typically return true on UNIX systems because `"PATH"` is an environment variable for UNIX.

  **See also:** [getenv], page 1068, [setenv], page 1068, [unsetenv], page 1068.

`setenv ("`*`var`*`", `*`value`*`)`
`setenv (`*`var`*`)`
`putenv (...)`
: Set the value of the environment variable *var* to *value*.

  If no *value* is specified then the variable will be assigned the null string.

  Programming Note: `putenv` is an alias for `setenv` and can be used interchangeably.

  **See also:** [unsetenv], page 1068, [getenv], page 1068, [isenv], page 1068.

`unsetenv ("`*`var`*`")`
*`status`* `= unsetenv ("`*`var`*`")`
: Delete the environment variable *var*.

  Return 0 if the variable was deleted, or did not exist, and -1 if an error occurred.

  **See also:** [setenv], page 1068, [getenv], page 1068, [isenv], page 1068.

*homedir* = get_home_directory ()

Return the current home directory.

On most systems, this is equivalent to `getenv ("HOME")`. On Windows systems, if the environment variable `HOME` is not set then it is equivalent to `fullfile (getenv ("HOMEDRIVE"), getenv ("HOMEPATH"))`

**See also:** [getenv], page 1068.

## 36.8 Current Working Directory

cd *dir*
cd
*old_dir* = cd
*old_dir* = cd (*dir*)
chdir ...

Change the current working directory to *dir*.

If called with no input or output arguments, the current directory is changed to the user's home directory (`"~"`).

For example,

```
cd ~/octave
```

changes the current working directory to `~/octave`. If the directory does not exist, an error message is printed and the working directory is not changed.

Programming Note: `chdir` is an alias for `cd` and can be used with all of the same calling formats.

Compatibility Note: When called with no arguments, MATLAB prints the present working directory rather than changing to the user's home directory.

**See also:** [pwd], page 1070, [mkdir], page 1039, [rmdir], page 1040, [dir], page 1070, [ls], page 1069.

ls
ls *filenames*
ls *options*
ls *options filenames*
*list* = ls (...)

List directory contents.

The `ls` function forwards to the `ls` command if it is available. It falls back to calling the native operating system's directory listing command. Available *options* may vary from system to system.

Filenames are subject to shell expansion if they contain any wildcard characters '`*`', '`?`', '`[]`'. If these wildcard characters are escaped with a backslash '`\`' (e.g., '`\*`') then they are not treated as wildcards, but instead as the corresponding literal character.

If the optional output *list* is requested then `ls` returns a character array with one row for each file/directory name.

Example usage on a UNIX-like system:

```
ls -l
     ⊣ total 12
     ⊣ -rw-r--r--   1 jwe  users  4488 Aug 19 04:02 foo.m
     ⊣ -rw-r--r--   1 jwe  users  1315 Aug 17 23:14 bar.m
```

**See also:** [dir], page 1070, [readdir], page 1044, [glob], page 1044, [what], page 156, [stat], page 1041, [filesep], page 1046, [ls_command], page 1070.

*val* = ls_command ()
*old_val* = ls_command (*new_val*)

Query or set the shell command used by Octave's `ls` command.

**See also:** [ls], page 1069.

dir
dir *directory*
[*list*] = dir (*directory*)

Display file listing for directory *directory*.

If *directory* is not specified then list the present working directory.

If a return value is requested, return a structure array with the fields

| | |
|---|---|
| name | File or directory name. |
| folder | Location of file or directory |
| date | Timestamp of file modification (string value). |
| bytes | File size in bytes. |
| isdir | True if name is a directory. |
| datenum | Timestamp of file modification as serial date number (double). |
| statinfo | Information structure returned from `stat`. |

If *directory* is a filename, rather than a directory, then return information about the named file. *directory* may also be a list rather than a single directory or file.

*directory* is subject to shell expansion if it contains any wildcard characters '`*`', '`?`', '`[]`'. If these wildcard characters are escaped with a backslash '`\`' (e.g., '`\*`') on a POSIX platform, then they are not treated as wildcards, but as the corresponding literal character. On Windows, it is not possible to escape wildcard characters because backslash '`\`' is treated as a file separator. On Windows, use `ls` for file or folder names that contain characters that would be treated as wildcards by `dir`.

Note that for symbolic links, `dir` returns information about the file that the symbolic link points to rather than the link itself. However, if the link points to a nonexistent file, `dir` returns information about the link.

**See also:** [ls], page 1069, [readdir], page 1044, [glob], page 1044, [what], page 156, [stat], page 1041, [lstat], page 1041.

*dir* = pwd ()

Return the current working directory.

**See also:** [cd], page 1069, [dir], page 1070, [ls], page 1069, [mkdir], page 1039, [rmdir], page 1040.

## 36.9 Password Database Functions

Octave's password database functions return information in a structure with the following fields.

`name` The user name.

`passwd` The encrypted password, if available.

`uid` The numeric user id.

`gid` The numeric group id.

`gecos` The GECOS field.

`dir` The home directory.

`shell` The initial shell.

In the descriptions of the following functions, this data structure is referred to as a *pw_struct*.

*pw_struct* = `getpwent` ()
: Return a structure containing an entry from the password database, opening it if necessary.

  Once the end of the data has been reached, `getpwent` returns 0.

  **See also:** [setpwent], page 1071, [endpwent], page 1071.

*pw_struct* = `getpwuid` (*uid*).
: Return a structure containing the first entry from the password database with the user ID *uid*.

  If the user ID does not exist in the database, `getpwuid` returns 0.

  **See also:** [getpwnam], page 1071.

*pw_struct* = `getpwnam` (*name*)
: Return a structure containing the first entry from the password database with the user name *name*.

  If the user name does not exist in the database, `getpwname` returns 0.

  **See also:** [getpwuid], page 1071.

[*status*, *msg*] = `setpwent` ()
: Return the internal pointer to the beginning of the password database.

  **See also:** [getpwent], page 1071, [endpwent], page 1071.

[*status*, *msg*] = `endpwent` ()
: Close the password database.

  **See also:** [getpwent], page 1071, [setpwent], page 1071.

## 36.10 Group Database Functions

Octave's group database functions return information in a structure with the following fields.

`name` The user name.

`passwd` The encrypted password, if available.

`gid` The numeric group id.

`mem` The members of the group.

In the descriptions of the following functions, this data structure is referred to as a *grp_struct*.

*grp_struct* = `getgrent ()`
: Return an entry from the group database, opening it if necessary.

  Once the end of data has been reached, `getgrent` returns 0.

  **See also:** [setgrent], page 1072, [endgrent], page 1072.

*grp_struct* = `getgrgid (`*gid*`)`.
: Return the first entry from the group database with the group ID *gid*.

  If the group ID does not exist in the database, `getgrgid` returns 0.

  **See also:** [getgrnam], page 1072.

*grp_struct* = `getgrnam (`*name*`)`
: Return the first entry from the group database with the group name *name*.

  If the group name does not exist in the database, `getgrnam` returns 0.

  **See also:** [getgrgid], page 1072.

[*status*, *msg*] = `setgrent ()`
: Return the internal pointer to the beginning of the group database.

  **See also:** [getgrent], page 1072, [endgrent], page 1072.

[*status*, *msg*] = `endgrent ()`
: Close the group database.

  **See also:** [getgrent], page 1072, [setgrent], page 1072.

## 36.11 System Information

`computer ()`
*comp* = `computer ()`
[*comp*, *maxsize*] = `computer ()`
[*comp*, *maxsize*, *endian*] = `computer ()`
*arch* = `computer ("`*arch*`")`
: Print or return a string of the form *cpu-vendor-os* that identifies the type of computer that Octave is running on.

If invoked with an output argument, the value is returned instead of printed. For example:

```
computer ()
   ⊣ x86_64-pc-linux-gnu

mycomp = computer ()
   ⇒ mycomp = x86_64-pc-linux-gnu
```

If two output arguments are requested, also return the maximum number of elements for an array. This will depend on whether Octave has been compiled with 32-bit or 64-bit index vectors.

If three output arguments are requested, also return the byte order of the current system as a character (`"B"` for big-endian or `"L"` for little-endian).

If the argument `"arch"` is specified, return a string indicating the architecture of the computer on which Octave is running.

Results may be different if Octave was invoked with the –traditional option.

**See also:** [isunix], page 1074, [ismac], page 1074, [ispc], page 1074.

`[`*uts*`, `*err*`, `*msg*`] = uname ()`

Return system information in the structure.

For example:

```
uname ()
   ⇒ {
        sysname = x86_64
        nodename = segfault
        release = 2.6.15-1-amd64-k8-smp
        version = Linux
        machine = #2 SMP Thu Feb 23 04:57:49 UTC 2006
     }
```

If successful, *err* is 0 and *msg* is an empty string. Otherwise, *err* is nonzero and *msg* contains a system-dependent error message.

*n* = nproc ()
*n* = nproc (*query*)

Return the current number of available (logical) processors.

This returns the number of logical processors. For processors with hyperthreading, this is larger than the number of physical cores.

If called with the optional argument *query*, modify how processors are counted as follows:

`all` total number of processors.

`current` processors available to the current process.

`overridable`
same as `current`, but overridable through the `OMP_NUM_THREADS` environment variable.

*tf* = `ispc` ()
: Return true if Octave is running on a Windows system and false otherwise.

  **See also:** [isunix], page 1074, [ismac], page 1074.

*tf* = `isunix` ()
: Return true if Octave is running on a Unix-like system and false otherwise.

  **See also:** [ismac], page 1074, [ispc], page 1074.

*tf* = `ismac` ()
: Return true if Octave is running on a Mac OS X system and false otherwise.

  **See also:** [isunix], page 1074, [ispc], page 1074.

*tf* = `isieee` ()
: Return true if your computer *claims* to conform to the IEEE 754 standard for floating point calculations.

  No actual tests are performed.

*tf* = `isdeployed` ()
: Return true if the current program has been compiled and is running separately from the Octave interpreter and false if it is running in the Octave interpreter.

  Currently, this function always returns false in Octave.

*tf* = `isstudent` ()
: Return true if running in the student edition of MATLAB.

  `isstudent` always returns false in Octave.

  **See also:** [false], page 67.

*dir* = `OCTAVE_HOME` ()
: Return the name of the top-level Octave installation directory.

  OCTAVE_HOME corresponds to the configuration variable *prefix*.

  **See also:** [EXEC_PATH], page 1062, [IMAGE_PATH], page 941, [OCTAVE_EXEC_HOME], page 1074.

*dir* = `OCTAVE_EXEC_HOME` ()
: Return the name of the top-level Octave installation directory for architecture-dependent files.

  If not specified separately, the value is the same as OCTAVE_HOME. OCTAVE_EXEC_HOME corresponds to the configuration variable *exec_prefix*.

  **See also:** [EXEC_PATH], page 1062, [IMAGE_PATH], page 941, [OCTAVE_HOME], page 1074.

*dir* = `matlabroot` ()
: Return the name of the top-level Octave installation directory.

  This is an alias for the function `OCTAVE_HOME` provided for compatibility.

  **See also:** [OCTAVE_HOME], page 1074.

`cfg_dir = user_config_dir ()`
Return the (platform-specific) directory for user configuration.

**See also:** [user_data_dir], page 1075.

`data_dir = user_data_dir ()`
Return the (platform-specific) directory for user data.

**See also:** [user_config_dir], page 1075.

*`verstr`* `= OCTAVE_VERSION ()`
Return the version number of Octave as a string.

**See also:** [ver], page 1075, [version], page 1075.

`v = version ()`
`[v, d] = version ()`
`v = version (`*`feature`*`)`
Get version information for Octave.

If called without input argument, the first return value *v* gives the version number of Octave as a string. The second return value *d* holds the release date as a string.

The following options can be passed for *feature*:

`"-date"` for the release date of the running build,

`"-description"`
for a description of the release (always an empty string),

`"-release"`
for the name of the running build (always an empty string),

`"-java"` for version information of the Java VM,

`"-fftw"` for version information for the linked FFTW,

`"-blas"` for version information for the linked BLAS,

`"-lapack"`
for version information for the linked LAPACK.

`"-hgid"` the mercurial ID of the sources used to build Octave.

The information returned for the `"-blas"` and `"-lapack"` options might be unreliable. It might report which library was linked in when Octave was built instead of which library is currently used.

The variant with no input and output argument is an alias for the function `OCTAVE_VERSION` provided for compatibility.

**See also:** [OCTAVE_VERSION], page 1075, [ver], page 1075.

`ver`
`ver` *Octave*
`ver` *package*

`v = ver (...)`

Display a header containing the current Octave version number, license string, and operating system. The header is followed by a list of installed packages, versions, and installation directories.

Use the package name *package* or Octave to query a specific component.

When called with an output argument, return a vector of structures describing Octave and each installed package. Each structure includes the following fields.

| | |
|---|---|
| `Name` | Package name. |
| `Version` | Version of the package. |
| `Release` | Release of the package (currently unused, defaults to `[]`). |
| `Date` | Date that the version was released. |

**See also:** [version], page 1075, [usejava], page 1147, [pkg], page 1084.

*tf* = `compare_versions (`*v1*, *v2*, *operator*`)`

Compare two version strings using the given *operator*.

This function assumes that versions *v1* and *v2* are arbitrarily long strings made of numeric and period characters possibly followed by an arbitrary string (e.g., `"1.2.3"`, `"0.3"`, `"0.1.2+"`, or `"1.2.3.4-test1"`).

The version is first split into numeric and character portions and then the parts are padded to be the same length (i.e., `"1.1"` would be padded to be `"1.1.0"` when being compared with `"1.1.1"`, and separately, the character parts of the strings are padded with nulls).

The operator can be any logical operator from the set

- `"=="` equal
- `"<"` less than
- `"<="` less than or equal to
- `">"` greater than
- `">="` greater than or equal to
- `"!="`, `"~="` not equal

Note that version `"1.1-test2"` will compare as greater than `"1.1-test10"`. Also, since the numeric part is compared first, `"a"` compares less than `"1a"` because the second string starts with a numeric part even though `double ("a")` is greater than `double ("1")`.

*tf* = `verLessThan (`*package*, *version*`)`

Return true if the installed version of the package is less than *version*.

*package* is the name of the package to check. Use `"Octave"` as the *package* to check the version of Octave itself.

*version* is the version to compare it to. A version is a string in the format accepted by `compare_versions`: an arbitrarily long string composed of numeric and period characters, possibly followed by an arbitrary string (e.g., `"1.2.3"`, `"0.3"`, `"0.1.2+"`, or `"1.2.3.4-test1"`).

Examples:

```
tf = verLessThan ("Octave", "5")
⇒ tf = 0

tf = verLessThan ("io", "2.4.12")
⇒ ...

if (! verLessThan ("Octave", "5"))
  ## ... use new Octave 5 features ...
endif
```

**See also:** [compare_versions], page 1076, [version], page 1075, [ver], page 1075, [pkg], page 1084.

`license`
`license` *inuse*
`license` *inuse* *feature*
`license` ("*inuse*")
*license_struct* = `license` ("*inuse*")
*license_struct* = `license` ("*inuse*", *feature*)
*status* = `license` ("*test*", *feature*)
*status* = `license` ("*checkout*", *feature*)
[*status*, *errmsg*] = `license` ("*checkout*", *feature*)

Get license information for Octave and Octave packages.

GNU Octave is free software distributed under the GNU General Public License (GPL), and a license manager makes no sense. This function is provided only for MATLAB compatibility.

When called with no extra input arguments, it returns the Octave license, otherwise the first input defines the operation mode and must be one of the following strings: `inuse`, `test`, and `checkout`. The optional *feature* argument can either be `"octave"` (core) or the name of an Octave package.

`"inuse"` Print a list of loaded features, i.e., "octave" and the list of loaded packages. If an output is requested, it returns a struct array with the fields `"feature"`, and `"user"`.

`"test"` Return true if the specified *feature* is installed, false otherwise.

An optional third argument `"enable"` or `"disable"` is accepted but ignored.

`"checkout"`

Return true if the specified *feature* is installed, false otherwise. An optional second output will have an error message if a package is not installed.

**See also:** [pkg], page 1084, [ver], page 1075, [version], page 1075.

`memory` ()
[*userdata*, *systemdata*] = `memory` ()

Display or return information about the memory usage of Octave.

If the function is called without output arguments, a table with an overview of the current memory consumption is displayed.

The output argument *userdata* is a structure with the following fields containing data for the Octave process:

`MaxPossibleArrayBytes`
: Maximum size for an array to be allocated. Be aware that this includes *all* physical memory and swap space. Allocating that amount of memory might result in system instability, data corruption, and/or file system corruption. Note that depending on the platform (32-bit systems), the largest contiguous memory block might further limit the maximum possible allocatable array. This check is not currently implemented.

`MemAvailableAllArrays`
: The total size of available memory in bytes.

`ram_available_all_arrays`
: The maximum size for an array that can be allocated in physical memory (excluding swap space). Note that depending on the platform (32-bit systems), the largest contiguous memory block might further limit the maximum possible allocatable array. This check is not currently implemented.

`MemUsedMATLAB`
`mem_used_octave`
: The memory (including swap space) currently used by Octave in bytes.

`ram_used_octave`
: The physical memory (excluding swap space) currently used by Octave in bytes.

The output argument *systemdata* is a nested structure with the following fields containing information about the system's memory:

`PhysicalMemory.Available`
: The currently available physical memory in bytes.

`PhysicalMemory.Total`
: The total physical memory in bytes.

`SystemMemory.Available`
: The currently available memory (including swap space) in bytes.

`SystemMemory.Total`
: The total memory (including swap space) in bytes.

`VirtualAddressSpace.Available`
: The currently available virtual address space in bytes.

`VirtualAddressSpace.Total`
: The total virtual address space in bytes.

Example #1 : print formatted table of memory usage

```
memory ()
⇒
System    RAM: 3934008 KiB,  swap: 4087804 KiB
Octave    RAM:  170596 KiB,  virt: 1347944 KiB
Free      RAM: 1954940 KiB,  swap: 4087804 KiB
Available RAM: 2451948 KiB, total: 6042744 KiB
```

Example #2 : return structs with memory usage information

```
[userdata, systemdata] = memory ()
⇒
 userdata =

   scalar structure containing the fields:

     MaxPossibleArrayBytes = 6.1622e+09
     MemAvailableAllArrays = 6.1622e+09
     ram_available_all_arrays = 2.4883e+09
     MemUsedMATLAB = 1.3825e+09
     mem_used_octave = 1.3825e+09
     ram_used_octave = 1.7824e+08

 systemdata =

   scalar structure containing the fields:

     PhysicalMemory =

       scalar structure containing the fields:

         Available = 2.4954e+09
         Total = 4.0284e+09

     SystemMemory =

       scalar structure containing the fields:

         Available = 6.6813e+09
         Total = 8.2143e+09

     VirtualAddressSpace =

       scalar structure containing the fields:

         Available = 2.8147e+14
         Total = 2.8147e+14
```

Programming Note: This function is implemented for Linux and Windows only.

**See also:** [computer], page 1072, [getpid], page 1067, [getrusage], page 1080, [nproc], page 1073, [uname], page 1073.

*procstats* = `getrusage ()`

Return a structure containing a number of statistics about the current Octave process.

Not all fields are available on all systems. If it is not possible to get CPU time statistics, the CPU time slots are set to zero. Other missing data are replaced by NaN. The list of possible fields is:

`idrss` Unshared data size.

`inblock` Number of block input operations.

`isrss` Unshared stack size.

`ixrss` Shared memory size.

`majflt` Number of major page faults.

`maxrss` Maximum data size.

`minflt` Number of minor page faults.

`msgrcv` Number of messages received.

`msgsnd` Number of messages sent.

`nivcsw` Number of involuntary context switches.

`nsignals` Number of signals received.

`nswap` Number of swaps.

`nvcsw` Number of voluntary context switches.

`oublock` Number of block output operations.

`stime` A structure containing the system CPU time used. The structure has the elements `sec` (seconds) `usec` (microseconds).

`utime` A structure containing the user CPU time used. The structure has the elements `sec` (seconds) `usec` (microseconds).

*value* = `winqueryreg (`*rootkey*`, `*subkey*`, `*valuename*`)`
*value* = `winqueryreg (`*rootkey*`, `*subkey*`)`
*names* = `winqueryreg ("name", `*rootkey*`, `*subkey*`)`

Query names or value from the Windows registry.

On Windows, return the value of the registry key *subkey* from the root key *rootkey*. You can specify the name of the queried registry value with the optional argument *valuename*. Otherwise, if called with only two arguments or *valuename* is empty, then the default value of *subkey* is returned. If the registry value is of type `"REG_DWORD"` then *value* is of class int32. If the value is of the type `"REG_SZ"` or `"REG_EXPAND_SZ"` a string is returned.

If the first argument is `"name"`, a cell array of strings with the names of the values at that key is returned.

The variable *rootkey* must be a string with a valid root key identifier:

HKCR
HKEY_CLASSES_ROOT
HKEY_CURRENT_CONFIG
HKCU
HKEY_CURRENT_USER
HKLM
HKEY_LOCAL_MACHINE
HKU
HKEY_USERS
HKEY_PERFORMANCE_DATA

Examples:

Get a list of value names at the key `'HKCU\Environment'`:

```
valuenames = winqueryreg ("name", "HKEY_CURRENT_USER", ...
                          "Environment");
```

For each *valuenames*, display the value:

```
for k = 1:numel (valuenames)
  val = winqueryreg ("HKEY_CURRENT_USER", "Environment", ...
                     valuenames{k});
  str = sprintf ("%s = %s", valuenames{k}, num2str (val));
  disp (str);
endfor
```

On non-Windows platforms this function fails with an error.

## 36.12 Hashing Functions

It is often necessary to find if two strings or files are identical. This might be done by comparing them character by character and looking for differences. However, this can be slow, and so comparing a hash of the string or file can be a rapid way of finding if the files differ.

Another use of the hashing function is to check for file integrity. The user can check the hash of the file against a known value and find if the file they have is the same as the one that the original hash was produced with.

Octave supplies the `hash` function to calculate hash values of strings and files, the latter in combination with the `fileread` function. The `hash` function supports many commonly used hash methods.

*hashval* = hash ("*hashfcn*", *str*)

Calculate the hash value of the string *str* using the hash function *hashfcn*.

The available hash functions are given in the table below.

'MD2' Message-Digest Algorithm 2 (RFC 1319).

'MD4' Message-Digest Algorithm 4 (RFC 1320).

'MD5' Message-Digest Algorithm 5 (RFC 1321).

'SHA1' Secure Hash Algorithm 1 (RFC 3174)

'SHA224' Secure Hash Algorithm 2 (224 Bits, RFC 3874)

'SHA256' Secure Hash Algorithm 2 (256 Bits, RFC 6234)

'SHA384' Secure Hash Algorithm 2 (384 Bits, RFC 6234)

'SHA512' Secure Hash Algorithm 2 (512 Bits, RFC 6234)

To calculate for example the MD5 hash value of the string `"abc"` the `hash` function is called as follows:

```
hash ("md5", "abc")
     ⊣ ans = 900150983cd24fb0d6963f7d28e17f72
```

For the same string, the SHA-1 hash value is calculated with:

```
hash ("sha1", "abc")
     ⊣ ans = a9993e364706816aba3e25717850c26c9cd0d89d
```

And to compute the hash value of a file, e.g., `file = "file.txt"`, call `hash` in combination with the `fileread`:

```
hash ("md5", fileread (file));
```

# 37 Packages

Since Octave is Free Software users are encouraged to share their programs with others. To aid this sharing Octave supports the installation of extra packages maintained by the Octave community at https://packages.octave.org.

## 37.1 Background Information

Packages can be installed globally (i.e., for all users of the system) or locally (i.e., for the current user only).

Global packages are installed by default in a system-wide location. This is usually a subdirectory of the folder where Octave itself is installed. Therefore, Octave needs write access to this folder to install global packages, which is usually only available when Octave is run with administrative privileges, such as when run as root (or superuser) on Unix-like systems, or run with elevated privileges ("Run as administrator") on Windows.

In contrast, local packages are installed by default in the user's home directory (or user profile on Windows) and are only available to that specific user. Usually, they can be installed without administrative privileges.

When Octave is running with administrative privileges, `pkg` will install packages to the global package location by default. Otherwise, packages will be installed to the local location by default. The user can override this default installation location with optional arguments (`-local` or `-global`) as described below. The currently used default package installation location can be queried with `pkg prefix`.

For global and local packages, there are separate databases holding the information about the installed packages. If some package is installed globally as well as locally, the local installation takes precedence over ("shadows") the global one. Which (global or local) package installation is used can also be manipulated by using prefixes and/or using the '`local_list`' input argument. Using these mechanisms, several different releases of the same package can be installed side by side as well (but cannot be loaded simultaneously).

Packages might depend on external software and/or other packages. To be able to install such packages, these dependencies should be installed beforehand. A package that depends on other package(s) can still be installed using the `-nodeps` flag. The effects of unsatisfied dependencies on external software—like libraries—depends on the individual package.

Packages must be loaded before they can be used. When loading a package, Octave performs the following tasks:

1. If the package depends on other packages (and `pkg load` is called without the `-nodeps` option), the package is not loaded immediately. Instead, those dependencies are loaded first (recursively if needed).
2. When all dependencies are satisfied, the package's subdirectories are added to the search path.

This load order leads to functions that are provided by dependencies being potentially shadowed by functions of the same name that are provided by top-level packages.

Each time, a package is added to the search path, initialization script(s) for the package are automatically executed if they are provided by the package.

## 37.2 Installing and Removing Packages

Assuming a package is available in the file `image-1.0.0.tar.gz` it can be installed from the Octave prompt with the command

```
pkg install image-1.0.0.tar.gz
```

If the package is installed successfully nothing will be printed on the prompt, but if a warning or error occurred during installation it will be reported. It is possible to install several packages at once by writing several package file names after the `pkg install` command. If a different version of the package is already installed it will be removed prior to installing the new package. This makes it easy to upgrade and downgrade the version of a package, but makes it impossible to have several versions of the same package installed at once.

To see which packages are installed type

```
pkg list
 ⊣ Package Name  | Version | Installation directory
 ⊣ --------------+---------+----------------------
 ⊣         image *|   1.0.0 | /home/jwe/octave/image-1.0.0
```

In this case, version 1.0.0 of the `image` package is installed. The `'*'` character next to the package name shows that the image package is loaded and ready for use.

It is possible to remove a package from the system using the `pkg uninstall` command like this

```
pkg uninstall image
```

If the package is removed successfully nothing will be printed in the prompt, but if a warning or error occurred it will be reported. It should be noted that the package file used for installation is not needed for removal, and that only the package name as reported by `pkg list` should be used when removing a package. It is possible to remove several packages at once by writing several package names after the `pkg uninstall` command.

To minimize the amount of code duplication between packages, it is possible that one package depends on another one. If a package depends on another, it will check if that package is installed during installation. If it is not, an error will be reported and the package will not be installed. This behavior can be disabled by passing the `-nodeps` flag to the `pkg install` command

```
pkg install -nodeps my_package_with_dependencies.tar.gz
```

Since the installed package expects its dependencies to be installed it may not function correctly. Because of this it is not recommended to disable dependency checking.

`pkg` *`command`*
`pkg` *`command pkg_name`*
`pkg` *`command pkg_name1 pkg_name2 ...`*
`pkg` *`command option1 ... pkg_name1 ...`*
`[`*`out1, ...`*`] = pkg (`*`command, ...`*`)`

Manage or query packages (groups of add-on functions) for Octave.

Depending on the value of *command* and on the number of requested return arguments, `pkg` can be used to perform several tasks. Possible values for *command* are:

'`search`' Search for packages having the specified search terms in the Octave Packages index. This may require an internet connection and the cURL library.

```
pkg search foo bar baz
```

shows packages whose descriptions contain all the search terms.

Search terms are case-insensitive and can be regular expressions as well. For example,

```
pkg search "[aeiou]{4,}"
```

shows all packages whose descriptions have four or more consecutive vowels.

The option `-all` as in

```
pkg search -all
```

shows *all* packages available on the Octave Packages index (not only those that can be installed with the `pkg` command).

By default, the Octave Packages index is downloaded only once per Octave session in order to improve performance and to reduce the load on the index server during repeated searches. An update of the index can be enforced with the option `-refresh`:

```
pkg search -all -refresh foo
```

If an output variable is provided, as in

```
mypackages = pkg ("search", "foo")
```

then `pkg search` returns only those package names matching the search term(s) *and* which can be installed with `pkg install`.

'`install`' Install named packages. For example, each of the following commands:

```
pkg install pkgname
pkg install 'pkgname-1.0.0.tar.gz'
pkg install 'https://somewebsite.org/pkgname-1.0.0.tar.gz'
```

might install the package `pkgname`. The sequence for the resolution of the arguments after `pkg install` is:

1. If `pkgname` is a local file, Octave installs it.
2. Otherwise, if `pkgname` resembles a URL, Octave downloads and installs it.
3. Otherwise, Octave queries the Octave Packages index online for a package named `pkgname`, and if found, downloads and installs its latest version. If the Octave Packages index cannot be reached, a locally cached version of the index might be used. The Octave Packages index is only downloaded once per Octave session (unless an index refresh is requested with the option `-refresh`).

Online access requires an internet connection and the cURL library.

*No support*: the GNU Octave community is not responsible for any installed packages. For support or for reporting bugs, you need to contact

the maintainers of the installed package directly (run `pkg describe` on the package to get information).

The following options are accepted for `pkg install`:

`-nodeps` Disable dependency checking. With this option it is possible to install a package even when it depends on another package which is not installed on the system. **Use this option with care.**

`-local` A local installation (package available only to current user) is forced, even if Octave is being run with administrative privileges.

`-global` Force a global installation (package available to all users), even if Octave is not being run with administrative privileges. The user must have write access to the global package store.

`-verbose` Print the output of all commands as they are performed.

`-refresh` Force an update of the cached package database from the online index. By default, Octave tries to download the Octave Packages index only once per Octave session. This option forces refreshing the index even if it was already downloaded in the current Octave session. The option `-refresh` applies to `install`, `search`, and `update` commands.

`-forge` The option `-forge` forces an installation of a package from the Octave Packages index. It cannot be used to install packages from a URL or a local file. It implies `-refresh`.

If the package tarball contains a `configure` script, it is run during the installation of the package. If it contains a `Makefile`, the command `make` is used by default. It is possible to override the default command with the environment variable `MAKE`. Some Octave packages might require GNU `make` which may be present under a different name such as `gmake`.

'`update`' Check installed Octave packages against the latest version on the Octave Packages index and update any outdated items. Updated packages are installed either globally or locally depending on whether Octave is running with elevated privileges. This requires an internet connection and the cURL library.

Options for the install command and the names of individual packages to be checked for updates may be specified as a list following the update command. If the option `-local` or `-global` is specified, `pkg update` limits the update check to the local or global installed packages, and installs updates in that same context. For example,

Update all packages:

```
pkg update
```

Update all local packages:

```
pkg update -local
```

Update certain packages, ignore dependencies, max verbosity:

```
pkg update -verbose -nodeps image signal geometry
```

Updates for multiple packages are sorted alphabetically and not checked for dependencies affected by installation order. If a dependency order related `pkg update` failure occurs, use `pkg update -nodeps` to ignore dependencies, or `pkg install <package_name>` to update individual packages manually.

'`uninstall`'
: Uninstall named packages. For example,

```
pkg uninstall image
```

removes the `image` package from the system. If another installed package depends on the `image` package an error is issued. The package removal can be forced by using the `-nodeps` option.

Depending on whether Octave is being run with administrative privileges only global or local packages are removed by default. See `-global` or `-local` to override that default behavior.

'`load`'
: Add named packages to the path. After loading a package it is possible to use the functions provided by the package. For example,

```
pkg load image
```

adds the `image` package to the path.

Note: When loading a package, `pkg` automatically tries to load any unloaded dependencies as well, unless the `-nodeps` flag has been specified. For example,

```
pkg load signal
```

adds the `signal` package and also tries to load its dependency: the `control` package. Be aware that use of the `-nodeps` option can adversely affect package functionality.

'`unload`'
: Remove named packages from the path. After unloading a package it is no longer possible to use the functions provided by the package. Trying to unload a package that other loaded packages still depend on causes an error; no packages are unloaded in this case. A package can be forcibly unloaded with the `-nodeps` flag, at the risk of breaking dependent packages that are still loaded.

'`list`'
: Show the list of currently installed packages. For example,

```
pkg list
```

will produce a short report with the package name, version, and installation directory for each installed package. Supply a package name to limit reporting to a particular package. For example:

```
pkg list image
```

If a single return argument is requested then `pkg` returns a cell array where each element is a structure with information on a single package.

```
installed_packages = pkg ("list")
```

If two output arguments are requested `pkg` splits the list of installed packages into those which were installed by the current user (local packages), and those which were installed by the system administrator (global packages).

```
[user_packages, system_packages] = pkg ("list")
```

The option `-forge` lists all packages available on the Octave Packages index that can be installed with the `pkg` command. This requires an internet connection and the cURL library. This also updates the locally cached file of the Octave Packages index. Calling

```
octave_packages = pkg ("list", "-forge")
```

is identical to calling

```
octave_packages = pkg ("search", "-refresh", "-all")
```

'`describe`'

Show a short description of installed packages. With the option `"-verbose"` also list functions provided by the package. For example,

```
pkg describe -verbose
```

describes all installed packages and the functions they provide.

The result can be limited to a set of packages:

```
## describe control and signal packages
pkg describe control signal
```

If one output is requested a cell of structure containing the description and list of functions of each package is returned as output rather than printed on screen:

```
desc = pkg ("describe", "secs1d", "image")
```

If any of the requested packages is not installed, `pkg` returns an error, unless a second output is requested:

```
[desc, flag] = pkg ("describe", "secs1d", "image")
```

*flag* takes one of the values `"Not installed"`, `"Loaded"`, or `"Not loaded"` for each of the named packages.

'`prefix`' Set the installation prefix directory. For example,

```
pkg prefix ~/my_octave_packages
```

sets the installation prefix to `~/my_octave_packages`.

Packages are installed in this directory.

Giving an output argument returns the current installation prefix. For example:

```
pfx = pkg ("prefix")
```

The location in which to install the architecture dependent files can be independently specified with an additional argument. For example:

```
pkg prefix ~/my_octave_packages ~/my_arch_dep_pkgs
```

'local_list'
: Set the file in which to look for information on locally installed packages. Locally installed packages are those that are typically available only to the current user. For example:

```
pkg local_list ~/.octave_packages
```

Get the current value of local_list with

```
pkg local_list
```

'global_list'
: Set the file in which to look for information on globally installed packages. Globally installed packages are those that are typically available to all users. For example:

```
pkg global_list /usr/share/octave/site/api-v59/octave_packages
```

Get the current value of global_list with

```
pkg global_list
```

'build'
: Build a binary form of a package or packages. The binary file produced is itself an Octave package that can be installed normally with `pkg`. For example:

```
pkg build builddir image-1.0.0.tar.gz ...
```

where `builddir` is a directory where the binary packages are built. The options `-verbose` and `-nodeps` are respected, while all other options are ignored.

'rebuild'
: Rebuild the package database from the installed global and local packages. This can be used in cases where the package database has been corrupted.

'test'
: Perform the built-in self tests contained in all functions provided by the named packages. For example:

```
pkg test image
```

**See also:** [ver], page 1075, [news], page 22.

## 37.3 Using Packages

By default installed packages are not available from the Octave prompt, but it is possible to control this using the `pkg load` and `pkg unload` commands. The functions from a package can be added to the Octave path by typing

```
pkg load package_name
```

where `package_name` is the name of the package to be added to the path.

In much the same way a package can be removed from the Octave path by typing

```
pkg unload package_name
```

## 37.4 Administrating Packages

It is possible to make both per-user (local) and system-wide (global) installations of a package. If the user performing the installation is `root` (or Administrator with elevated privileges on Windows), the packages by default install in a system-wide directory that defaults to

*OCTAVE_HOME*/share/octave/packages/. If the user is not root (or Octave is running without elevated privileges), packages are installed locally. The default installation directory for local packages is *user_data_dir*/octave/*OCTAVE_API_VERSION*/packages. Packages will be installed in a subdirectory of the installation directory that will be named after the package. It is possible to change the installation directory by using the pkg prefix command:

```
pkg prefix new_installation_directory
```

The current installation directory can be retrieved by typing

```
current_installation_directory = pkg ("prefix")
```

The package manager stores some information about the installed packages in configuration files. For per-user (local) packages, this information is stored in the file *user_config_dir*/octave/*OCTAVE_API_VERSION*/octave_packages by default. For system-wide (global) installations, it is stored in *OCTAVE_HOME*/share/octave/octave_packages. The path to the per-user file can be changed with the pkg local_list command:

```
pkg local_list /path/to/new_file
```

For system-wide installations, this can be changed in the same way using the pkg global_list command. If these commands are called without a new path, the current path will be returned. To retain these settings between sessions, they can be set in one of the startup files, see Section 2.1.2 [Startup Files], page 19.

## 37.5 Creating Packages

Internally a package is simply a gzipped tar file that contains a top level directory of any given name. This directory will in the following be referred to as package and may contain the following files:

package/CITATION
: This is am optional file describing instructions on how to cite the package for publication. It will be displayed verbatim by the function citation.

package/COPYING
: This is a required file containing the license of the package. No restrictions is made on the license in general. If however the package contains dynamically linked functions the license must be compatible with the GNU General Public License.

package/DESCRIPTION
: This is a required file containing information about the package. See Section 37.5.1 [The DESCRIPTION File], page 1092, for details on this file.

package/ChangeLog
: This is an optional file describing all the changes made to the package source files.

package/INDEX
: This is an optional file describing the functions provided by the package. If this file is not given then one with be created automatically from the functions in the package and the Categories keyword in the DESCRIPTION file. See Section 37.5.2 [The INDEX File], page 1094, for details on this file.

`package/NEWS`
: This is an optional file describing all user-visible changes worth mentioning. As this file increases on size, old entries can be moved into `package/ONEWS`.

`package/ONEWS`
: This is an optional file describing old entries from the `NEWS` file.

`package/PKG_ADD`
: An optional file that includes commands that are run when the package is added to the users path. Note that `PKG_ADD` directives in the source code of the package will also be added to this file by the Octave package manager. Note that symbolic links are to be avoided in packages, as symbolic links do not exist on some file systems, and so a typical use for this file is the replacement of the symbolic link

```
ln -s foo.oct bar.oct
```

  with an autoload directive like

```
autoload ('bar', which ('foo'));
```

  See Section 37.5.3 [PKG_ADD and PKG_DEL Directives], page 1095, for details on `PKG_ADD` directives.

`package/PKG_DEL`
: An optional file that includes commands that are run when the package is removed from the users path. Note that `PKG_DEL` directives in the source code of the package will also be added to this file by the Octave package manager. See Section 37.5.3 [PKG_ADD and PKG_DEL Directives], page 1095, for details on `PKG_DEL` directives.

`package/pre_install.m`
: This is an optional function that is run prior to the installation of a package. This function is called with a single argument, a struct with fields names after the data in the `DESCRIPTION`, and the paths where the package functions will be installed.

`package/post_install.m`
: This is an optional function that is run after the installation of a package. This function is called with a single argument, a struct with fields names after the data in the `DESCRIPTION`, and the paths where the package functions were installed.

`package/on_uninstall.m`
: This is an optional function that is run prior to the removal of a package. This function is called with a single argument, a struct with fields names after the data in the `DESCRIPTION`, the paths where the package functions are installed, and whether the package is currently loaded.

Besides the above mentioned files, a package can also contain one or more of the following directories:

`package/inst`
: An optional directory containing any files that are directly installed by the package. Typically this will include any `m`-files.

`package/src`
: An optional directory containing code that must be built prior to the packages installation. The Octave package manager will execute `./configure` in this directory if this script exists, and will then call `make` if a file `Makefile` exists in this directory. `make install` will however not be called. The environment variables `MKOCTFILE`, `OCTAVE_CONFIG`, and `OCTAVE` will be set to the full paths of the programs `mkoctfile`, `octave-config`, and `octave`, respectively, of the correct version when `configure` and `make` are called. If a file called `FILES` exists all files listed there will be copied to the `inst` directory, so they also will be installed. If the `FILES` file doesn't exist, `src/*.m` and `src/*.oct` will be copied to the `inst` directory.

`package/doc`
: An optional directory containing documentation for the package. The files in this directory will be directly installed in a sub-directory of the installed package for future reference.

  If the Octave GUI is running, the optional Qt help file `package/doc/pkg-name.qch` is added to the documentation browser while loading the package. The help file is removed from the documentation browser when the package is unloaded.

`package/bin`
: An optional directory containing files that will be added to the Octave `EXEC_PATH` when the package is loaded. This might contain external scripts, etc., called by functions within the package.

### 37.5.1 The DESCRIPTION File

The `DESCRIPTION` file contains various information about the package, such as its name, author, and version. This file has a very simple format

- Lines starting with '`#`' are comments.
- Lines starting with a blank character are continuations from the previous line.
- Everything else is of the form `NameOfOption: ValueOfOption`.

The following is a simple example of a `DESCRIPTION` file

```
Name: The name of my package
Version: 1.0.0
Date: 2007-18-04
Author: The name (and possibly email) of the package author.
Maintainer: The name (and possibly email) of the current
 package maintainer.
Title: The title of the package
Description: A short description of the package.  If this
 description gets too long for one line it can continue
 on the next by adding a space to the beginning of the
 following lines.
License: GPLv3+
```

The package manager currently recognizes the following keywords

`Name` Name of the package.

`Version` Version of the package. A package version is typically digits separated by dots but may also contain '`+`', '`-`', '`~`', and alphanumeric characters (in the `"C"` locale). For example, `"2.1.0+"` could indicate a development version of a package. Versions are compared using [compare_versions], page 1076.

`Date` Date of last update.

`Author` Original author of the package.

`Maintainer`
Maintainer of the package.

`Title` A one line description of the package.

`Description`
A one paragraph description of the package.

`Categories`
Optional keyword describing the package (if no `INDEX` file is given this is mandatory).

`Problems` Optional list of known problems.

`Url` Optional URL to the homepage or repository related to the package.

`Tracker` Optional URL to the bug tracker related to the package. It is highly encouraged that maintainers utilize a dedicated tracker for reporting issues related to the package's functionality in order to keep Octave's bug tracker at Savannah.org less bloated and solely for core Octave bug reports.

`Depends` A list of other Octave packages that this package depends on. This can include dependencies on particular versions, with the following format:

```
Depends: package (>= 1.0.0)
```

Possible operators are `<`, `<=`, `==`, `>=` or `>`. If the part of the dependency in `()` is missing, any version of the package is acceptable. Multiple dependencies can be defined as a comma separated list. This can be used to define a range of versions of a particular package:

```
Depends: package (>= 1.0.0), package (< 1.5.0)
```

It is also possible to depend on particular versions of Octave core:

```
Depends: octave (>= 3.8.0)
```

`License` An optional short description of the used license (e.g., GPL version 3 or newer). This is optional since the file `COPYING` is mandatory.

`SystemRequirements`
These are the external install dependencies of the package and are not checked by the package manager. This is here as a hint to the distribution packager. They follow the same conventions as the `Depends` keyword.

`BuildRequires`
These are the external build dependencies of the package and are not checked by the package manager. This is here as a hint to the distribution packager. They follow the same conventions as the `Depends` keyword. Note that in general,

packaging systems such as `rpm` or `deb` autoprobe the install dependencies from the build dependencies, and therefore a `BuildRequires` dependency usually removes the need for a `SystemRequirements` dependency.

The developer is free to add additional arguments to the `DESCRIPTION` file for their own purposes. One further detail to aid the packager is that the `SystemRequirements` and `BuildRequires` keywords can have a distribution dependent section, and the automatic build process will use these. An example of the format of this is

```
BuildRequires: libtermcap-devel [Mandriva] libtermcap2-devel
```

where the first package name will be used as a default and if the RPMs are built on a Mandriva distribution, then the second package name will be used instead.

### 37.5.2 The INDEX File

The optional `INDEX` file provides a categorical view of the functions in the package. This file has a very simple format

- Lines beginning with '`#`' are comments.
- The first non-comment line should look like this

  ```
  toolbox >> Toolbox name
  ```

- Lines beginning with an alphabetical character indicates a new category of functions.
- Lines starting with a white space character indicate that the function names on the line belong to the last mentioned category.

The format can be summarized with the following example:

```
# A comment
toolbox >> Toolbox name
Category Name 1
 function1 function2 function3
 function4
Category Name 2
 function2 function5
```

If you wish to refer to a function that users might expect to find in your package but is not there, providing a work around or pointing out that the function is available elsewhere, you can use:

```
fn = workaround description
```

This workaround description will not appear when listing functions in the package with `pkg describe` but they will be published in the HTML documentation online. Workaround descriptions can use any HTML markup, but keep in mind that it will be enclosed in a bold-italic environment. For the special case of:

```
fn = use <code>alternate expression</code>
```

the bold-italic is automatically suppressed. You will need to use `<code>` even in references:

```
fn = use <a href="someothersite.html"><code>fn</code></a>
```

Sometimes functions are only partially compatible, in which case you can list the non-compatible cases separately. To refer to another function in the package, use `<f>fn</f>`. For example:

```
eig (a, b) = use <f>qz</f>
```

Since sites may have many missing functions, you can define a macro rather than typing the same link over and again.

```
$id = expansion
```

defines the macro id. You can use `$id` anywhere in the description and it will be expanded. For example:

```
$TSA = see <a href="link_to_spctools">SPC Tools</a>
arcov = $TSA <code>armcv</code>
```

id is any string of letters, numbers and `_`.

### 37.5.3 PKG_ADD and PKG_DEL Directives

If the package contains files called `PKG_ADD` or `PKG_DEL` the commands in these files will be executed when the package is added or removed from the users path. In some situations such files are a bit cumbersome to maintain, so the package manager supports automatic creation of such files. If a source file in the package contains a `PKG_ADD` or `PKG_DEL` directive they will be added to either the `PKG_ADD` or `PKG_DEL` files.

In `m`-files a `PKG_ADD` directive looks like this

```
## PKG_ADD: some_octave_command
```

Such lines should be added before the `function` keyword. In C++ files a `PKG_ADD` directive looks like this

```
// PKG_ADD: some_octave_command
```

In both cases `some_octave_command` should be replaced by the command that should be placed in the `PKG_ADD` file. `PKG_DEL` directives work in the same way, except the `PKG_ADD` keyword is replaced with `PKG_DEL` and the commands get added to the `PKG_DEL` file.

### 37.5.4 Missing Components

If a package relies on a component, such as another Octave package, that may not be present it may be useful to install a function which informs users what to do when a particular component is missing. The function must be written by the package maintainer and registered with Octave using `missing_component_hook`.

*val* = missing_component_hook ()
*old_val* = missing_component_hook (*new_val*)
*old_val* = missing_component_hook (*new_val*, "*local*")

Query or set the internal variable that specifies the function to call when a component of Octave is missing.

This can be useful for packagers that may split the Octave installation into multiple sub-packages, for example, to provide a hint to users for how to install the missing components.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

The hook function is expected to be of the form

*fcn* (*component*)

Octave will call *fcn* with the name of the function that requires the component and a string describing the missing component. The hook function should return an error message to be displayed.

**See also:** [missing_function_hook], page 1201.

# Appendix A External Code Interface

"The sum of human wisdom is not contained in any one language"

— Ezra Pound

Octave is a fantastic language for solving many problems in science and engineering. However, it is not the only computer language and there are times when you may want to use code written in other languages. Good reasons for doing so include: 1) not re-inventing the wheel; existing function libraries which have been thoroughly tested and debugged or large scale simulation codebases are a good example, 2) accessing unique capabilities of a different language; for example the well-known regular expression functions of Perl (but don't do that because `regexp` already exists in Octave).

Performance should generally **not** be a reason for using compiled extensions. Although compiled extensions can run faster, particularly if they replace a loop in Octave code, this is almost never the best path to take. First, there are many techniques to speed up Octave performance while remaining within the language. Second, Octave is a high-level language that makes it easy to perform common mathematical tasks. Giving that up means shifting the focus from solving the real problem to solving a computer programming problem. It means returning to low-level constructs such as pointers, memory management, mathematical overflow/underflow, etc. Because of the low level nature, and the fact that the compiled code is executed outside of Octave, there is the very real possibility of crashing the interpreter and losing work.

Before going further, you should first determine if you really need to bother writing code outside of Octave.

- Can I get the same functionality using the Octave scripting language alone?

  Even when a function already exists outside the language, it may be better to simply reproduce the behavior in an m-file rather than attempt to interface to the outside code.

- Is the code thoroughly optimized for Octave?

  If performance is an issue you should always start with the in-language techniques for getting better performance. Chief among these is vectorization (see Chapter 19 [Vectorization and Faster Code Execution], page 689) which not only makes the code concise and more understandable but improves performance (10X-100X). If loops must be used, make sure that the allocation of space for variables takes place outside the loops using an assignment to a matrix of the right size, or zeros.

- Does the code make as much use as possible of existing built-in library routines?

  These routines are highly optimized and many do not carry the overhead of being interpreted.

- Does writing a dynamically linked function represent a useful investment of your time, relative to staying in Octave?

  It will take time to learn Octave's interface for external code and there will inevitably be issues with tools such as compilers.

With that said, Octave offers a versatile interface for including chunks of compiled code as dynamically linked extensions. These dynamically linked functions can be called from the interpreter in the same manner as any ordinary function. The interface is bi-directional and

external code can call Octave functions (like `plot`) which otherwise might be very difficult to develop.

The interface is centered around supporting the languages C++, C, and Fortran. Octave itself is written in C++ and can call external C++/C code through its native oct-file interface. The C language is also supported through the mex-file interface for compatibility with MATLAB. Fortran code is easiest to reach through the oct-file interface.

Because many other languages provide C or C++ APIs it is relatively simple to build bridges between Octave and other languages. This is also a way to bridge to hardware resources which often have device drivers written in C.

## A.1 Oct-Files

### A.1.1 Getting Started with Oct-Files

Oct-files are pieces of C++ code that have been compiled with the Octave API into a dynamically loadable object. They take their name from the file which contains the object which has the extension `.oct`.

Finding a C++ compiler, using the correct switches, adding the right include paths for header files, etc. is a difficult task. Octave automates this by providing the `mkoctfile` command with which to build oct-files. The command is available from within Octave or at the shell command line.

`mkoctfile` [*-options*] *file ...*
`[`*output*`, `*status*`] = mkoctfile (...)`

The `mkoctfile` function compiles source code written in C, C++, or Fortran. Depending on the options used with `mkoctfile`, the compiled code can be called within Octave or can be used as a stand-alone application.

`mkoctfile` can be called from the shell prompt or from the Octave prompt. Calling it from the Octave prompt simply delegates the call to the shell prompt. Any output is stored in the *output* variable and the exit status in the *status* variable. If called with no outputs and the compilation fails then Octave will emit an error. If the programmer requests *output* or *status*, however, Octave will merely issue a warning and it is the programmer's responsibility to verify the command was successful.

`mkoctfile` accepts the following options, all of which are optional except for the filename of the code you wish to compile:

'`-I DIR`' Add the include directory DIR to compile commands.

'`-D DEF`' Add the definition DEF to the compiler call.

'`-l LIB`' Add the library LIB to the link command.

'`-L DIR`' Add the library directory DIR to the link command.

'`-M`'
'`--depend`'
Generate dependency files (.d) for C and C++ source files.

'`-R DIR`' Add the run-time path to the link command.

‘`-Wl,...`’ Pass options to the linker like `"-Wl,-rpath=..."`. The quotes are needed since commas are interpreted as command separators.

‘`-W...`’ Pass options to the assembler like `"-Wa,OPTION"`.

‘`-c`’ Compile but do not link.

‘`-g`’ Enable debugging options for compilers.

‘`-o FILE`’
‘`--output FILE`’
Output filename. Default extension is `.oct` (or `.mex` if ‘`--mex`’ is specified) unless linking a stand-alone executable.

‘`-p VAR`’
‘`--print VAR`’
Print configuration variable VAR. There are three categories of variables:

Octave configuration variables that users may override with environment variables. These are used in commands that `mkoctfile` executes.

| | |
|---|---|
| ALL_CFLAGS | INCLUDEDIR |
| ALL_CXXFLAGS | LAPACK_LIBS |
| ALL_FFLAGS | LDFLAGS |
| ALL_LDFLAGS | LD_STATIC_FLAG |
| BLAS_LIBS | LIBDIR |
| CC | LIBOCTAVE |
| CFLAGS | LIBOCTINTERP |
| CPICFLAG | LIBOCTMEX |
| CPPFLAGS | OCTAVE_LINK_OPTS |
| CXX | OCTINCLUDEDIR |
| CXXFLAGS | OCTAVE_LIBS |
| CXXLD | OCTAVE_LINK_DEPS |
| CXXPICFLAG | OCTLIBDIR |
| DL_LDFLAGS | OCT_LINK_DEPS |
| F77 | OCT_LINK_OPTS |
| F77_INTEGER8_FLAG | RDYNAMIC_FLAG |
| FFLAGS | SPECIAL_MATH_LIB |
| FPICFLAG | XTRA_CFLAGS |
| INCFLAGS | XTRA_CXXFLAGS |

Octave configuration variables as above, but currently unused by `mkoctfile`.

```
AR
DEPEND_EXTRA_SED_PATTERN
DEPEND_FLAGS
FFTW3F_LDFLAGS
FFTW3F_LIBS
FFTW3_LDFLAGS
FFTW3_LIBS
FFTW_LIBS
FLIBS
LIBS
RANLIB
READLINE_LIBS
```

Octave configuration variables that are provided for informational purposes only. Except for 'OCTAVE_HOME' and 'OCTAVE_EXEC_HOME', users may not override these variables.

If OCTAVE_HOME or OCTAVE_EXEC_HOME are set in the environment, then other variables are adjusted accordingly with OCTAVE_HOME or OCTAVE_EXEC_HOME substituted for the original value of the directory specified by the --prefix or --exec-prefix options that were used when Octave was configured.

```
API_VERSION                 LOCALARCHLIBDIR
ARCHLIBDIR                  LOCALFCNFILEDIR
BINDIR                      LOCALOCTFILEDIR
CANONICAL_HOST_TYPE         LOCALSTARTUPFILEDIR
DATADIR                     LOCALVERARCHLIBDIR
DATAROOTDIR                 LOCALVERFCNFILEDIR
DEFAULT_PAGER               LOCALVEROCTFILEDIR
EXEEXT                      MAN1DIR
FCNFILEDIR                  MAN1EXT
IMAGEDIR                    MANDIR
INCLUDEDIR                  OCTAVE_EXEC_HOME
INFODIR                     OCTAVE_HOME
INFOFILE                    OCTDATADIR
LIBDIR                      OCTDOCDIR
LIBEXECDIR                  OCTFILEDIR
LOCALAPIARCHLIBDIR          OCTFONTSDIR
LOCALAPIFCNFILEDIR          OCTINCLUDEDIR
LOCALAPIOCTFILEDIR          OCTLIBDIR
LOCALAPIPKGDIR              STARTUPFILEDIR
                            VERSION
```

'--link-stand-alone'
: Link a stand-alone executable file.

'--mex'
: Assume creation of a MEX file. Set the default output extension to .mex. Link to liboctmex instead of liboctinterp and liboctave.

‘`-s`’
‘`--strip`’ Strip the output file.

‘`-v`’
‘`--verbose`’
Echo commands as they are executed.

‘`file`’ The file to compile or link. Recognized file types are:

```
.c    C source
.cc   C++ source
.cp   C++ source
.cpp  C++ source
.CPP  C++ source
.cxx  C++ source
.c++  C++ source
.C    C++ source
.f    Fortran source (fixed form)
.F    Fortran source (fixed form)
.f90  Fortran source (free form)
.F90  Fortran source (free form)
.o    object file
.a    library file
```

Consider the following short example which introduces the basics of writing a C++ function that can be linked to Octave.

```
#include <octave/oct.h>

DEFUN_DLD (helloworld, args, nargout,
           "Hello World Help String")
{
  octave_stdout << "Hello World has "
                << args.length () << " input arguments and "
                << nargout << " output arguments.\n";

  // Return empty matrices for any outputs
  octave_value_list retval (nargout);
  for (int i = 0; i < nargout; i++)
    retval(i) = octave_value (Matrix ());

  return retval;
}
```

The first critical line is `#include <octave/oct.h>` which makes available most of the definitions necessary for a C++ oct-file. Note that `octave/oct.h` is a C++ header and cannot be directly `#include`’ed in a C source file, nor any other language.

Included by `oct.h` is a definition for the macro `DEFUN_DLD` which creates a dynamically loaded function. This macro takes four arguments:

1. The function name as it will be seen in Octave,

2. The list of arguments to the function of type `octave_value_list`,
3. The number of output arguments, which can be—and often is—omitted if not used, and
4. The string to use for the help text of the function.

The return type of functions defined with `DEFUN_DLD` is always `octave_value_list`.

There are a couple of important considerations in the choice of function name. First, it must be a valid Octave function name and so must be a sequence of letters, digits, and underscores not starting with a digit. Second, as Octave uses the function name to define the filename it attempts to find the function in, the function name in the `DEFUN_DLD` macro must match the filename of the oct-file. Therefore, the above function should be in a file `helloworld.cc`, and would be compiled to an oct-file using the command

```
mkoctfile helloworld.cc
```

This will create a file called `helloworld.oct` that is the compiled version of the function. It should be noted that it is perfectly acceptable to have more than one `DEFUN_DLD` function in a source file. However, there must either be a symbolic link to the oct-file for each of the functions defined in the source code with the `DEFUN_DLD` macro or the `autoload` (see Section 11.10 [Function Files], page 224) function should be used.

The rest of the function shows how to find the number of input arguments, how to print through the Octave pager, and how to return from the function. After compiling this function as above, an example of its use is

```
helloworld (1, 2, 3)
⊣ Hello World has 3 input arguments and 0 output arguments.
```

Subsequent sections show how to use specific classes from Octave's core internals. Base classes like `dMatrix` (a matrix of double values) are found in the directory `liboctave/array`. The definitive reference for how to use a particular class is the header file itself. However, it is often enough simply to study the examples in the manual in order to be able to use a class.

### A.1.2 Matrices and Arrays in Oct-Files

Octave supports a number of different array and matrix classes, the majority of which are based on the `Array` class. The exception are the sparse matrix types discussed separately below. There are three basic matrix types:

`Matrix` A double precision matrix class defined in `dMatrix.h`

`ComplexMatrix`
: A complex matrix class defined in `CMatrix.h`

`BoolMatrix`
: A boolean matrix class defined in `boolMatrix.h`

These are the basic two-dimensional matrix types of Octave. In addition there are a number of multi-dimensional array types including

`NDArray` A double precision array class defined in `dNDArray.h`

`ComplexNDarray`
: A complex array class defined in `CNDArray.h`

`boolNDArray`
: A boolean array class defined in `boolNDArray.h`

`int8NDArray`
`int16NDArray`
`int32NDArray`
`int64NDArray`
: 8, 16, 32, and 64-bit signed array classes defined in `int8NDArray.h`, `int16NDArray.h`, etc.

`uint8NDArray`
`uint16NDArray`
`uint32NDArray`
`uint64NDArray`
: 8, 16, 32, and 64-bit unsigned array classes defined in `uint8NDArray.h`, `uint16NDArray.h`, etc.

There are several basic ways of constructing matrices or multi-dimensional arrays. Using the class `Matrix` as an example one can

- Create an empty matrix or array with the empty constructor. For example:

  ```
  Matrix a;
  ```

  This can be used for all matrix and array types.
- Define the dimensions of the matrix or array with a dim_vector which has the same characteristics as the vector returned from `size`. For example:

  ```
  dim_vector dv (2, 3);  // 2 rows, 3 columns
  Matrix a (dv);
  ```

  This can be used for all matrix and array types.
- Define the number of rows and columns in the matrix. For example:

  ```
  Matrix a (2, 2)
  ```

  This constructor can **only** be used with matrix types.

These types all share a number of basic methods and operators. Many bear a resemblance to functions that exist in the interpreter. A selection of useful methods include

`T& operator ()` (*octave_idx_type*) [Method]
`T& elem` (*octave_idx_type*) [Method]
: The `()` operator or `elem` method allow the values of the matrix or array to be read or set. These methods take a single argument, which is of type `octave_idx_type`, that is the index into the matrix or array. Additionally, the matrix type allows two argument versions of the `()` operator and `elem` method, giving the row and column index of the value to get or set.

Note that these functions do significant error checking and so in some circumstances the user might prefer to access the data of the array or matrix directly through the `rwdata` method discussed below.

`octave_idx_type numel ()` *const* [Method]
: The total number of elements in the matrix or array.

`size_t byte_size () const` [Method]
: The number of bytes used to store the matrix or array.

`dim_vector dims () const` [Method]
: The dimensions of the matrix or array in value of type `dim_vector`.

`int ndims () const` [Method]
: The number of dimensions of the matrix or array. Matrices are always 2-D, but arrays can be N-dimensional.

`void resize (const dim_vector&)` [Method]
`void resize (nrows, ncols)` [Method]
: A method taking either an argument of type `dim_vector`, or, in the case of a matrix, two arguments of type `octave_idx_type` defining the number of rows and columns in the matrix.

`T * rwdata ()` [Method]
: This method returns a pointer to the underlying data of the matrix or array so that it can be manipulated directly, either within Octave or by an external library.

Operators such as `+`, `-`, or `*` can be used on the majority of the matrix and array types. In addition there are a number of methods that are of interest only for matrices such as `transpose`, `hermitian`, `solve`, etc.

The typical way to extract a matrix or array from the input arguments of `DEFUN_DLD` function is as follows

```
#include <octave/oct.h>

DEFUN_DLD (addtwomatrices, args, , "Add A to B")
{
  if (args.length () != 2)
    print_usage ();

  NDArray A = args(0).array_value ();
  NDArray B = args(1).array_value ();

  return octave_value (A + B);
}
```

To avoid segmentation faults causing Octave to abort, this function explicitly checks that there are sufficient arguments available before accessing these arguments. It then obtains two multi-dimensional arrays of type `NDArray` and adds these together. Note that the `array_value` method is called without using the `is_matrix_type` method. If an error occurs when attempting to extract the value, Octave will print a message and throw an exception. The reason to prefer this coding structure is that the arguments might be a type which is not an `NDArray`, but for which it would make sense to convert them to one. The `array_value` method allows this conversion to be performed transparently when possible. If you need to catch errors like this, and perform some kind of cleanup or other operation, you can catch the `octave_execution_error` exception.

`A + B`, operating on two `NDArray` objects returns an `NDArray`, which is cast to an `octave_value` on the return from the function. An example of the use of this demonstration function is

```
addtwomatrices (ones (2, 2), eye (2, 2))
      ⇒  2  1
         1  2
```

A list of the basic `Matrix` and `Array` types, the methods to extract these from an `octave_value`, and the associated header file is listed below.

| **Type** | **Function** | **Source Code** |
|---|---|---|
| RowVector | row_vector_value | dRowVector.h |
| ComplexRowVector | complex_row_vector_value | CRowVector.h |
| ColumnVector | column_vector_value | dColVector.h |
| ComplexColumnVector | complex_column_vector_value | CColVector.h |
| Matrix | matrix_value | dMatrix.h |
| ComplexMatrix | complex_matrix_value | CMatrix.h |
| boolMatrix | bool_matrix_value | boolMatrix.h |
| charMatrix | char_matrix_value | chMatrix.h |
| NDArray | array_value | dNDArray.h |
| ComplexNDArray | complex_array_value | CNDArray.h |
| boolNDArray | bool_array_value | boolNDArray.h |
| charNDArray | char_array_value | charNDArray.h |
| int8NDArray | int8_array_value | int8NDArray.h |
| int16NDArray | int16_array_value | int16NDArray.h |
| int32NDArray | int32_array_value | int32NDArray.h |
| int64NDArray | int64_array_value | int64NDArray.h |
| uint8NDArray | uint8_array_value | uint8NDArray.h |
| uint16NDArray | uint16_array_value | uint16NDArray.h |
| uint32NDArray | uint32_array_value | uint32NDArray.h |
| uint64NDArray | uint64_array_value | uint64NDArray.h |

### A.1.3 Character Strings in Oct-Files

A character string in Octave is just a special `Array` class. Consider the example:

```
#include <octave/oct.h>

DEFUN_DLD (stringdemo, args, , "String Demo")
{
  if (args.length () != 1)
    print_usage ();

  octave_value_list retval;

  charMatrix ch = args(0).char_matrix_value ();

  retval(1) = octave_value (ch, '\'');  // Single Quote String
```

```
  octave_idx_type nr = ch.rows ();

  for (octave_idx_type i = 0; i < nr / 2; i++)
    {
      std::string tmp = ch.row_as_string (i);

      ch.insert (ch.row_as_string (nr-i-1).c_str (), i, 0);
      ch.insert (tmp.c_str (), nr-i-1, 0);
    }

  retval(0) = octave_value (ch, '"');  // Double Quote String

  return retval;
}
```

An example of the use of this function is

```
s0 = ["First String"; "Second String"];
[s1,s2] = stringdemo (s0)
⇒ s1 = Second String
       First String

⇒ s2 = First String
       Second String

typeinfo (s2)
⇒ sq_string
typeinfo (s1)
⇒ string
```

One additional complication of strings in Octave is the difference between single quoted and double quoted strings. To find out if an `octave_value` contains a single or double quoted string use one of the predicate tests shown below.

```
if (args(0).is_sq_string ())
  octave_stdout << "First argument is a single quoted string\n";
else if (args(0).is_dq_string ())
  octave_stdout << "First argument is a double quoted string\n";
```

Note, however, that both types of strings are represented by the `charNDArray` type, and so when assigning to an `octave_value`, the type of string should be specified. For example:

```
octave_value_list retval;
charNDArray ch;
...
// Create single quoted string
retval(1) = octave_value (ch);   // default constructor is sq_string
            OR
retval(1) = octave_value (ch, '\'');  // explicitly create sq_string

// Create a double quoted string
retval(0) = octave_value (ch, '"');
```

### A.1.4 Cell Arrays in Oct-Files

Octave's cell type is also available from within oct-files. A cell array is just an `Array` of `octave_value`s, and thus each element of the cell array can be treated like any other `octave_value`. A simple example is

```
#include <octave/oct.h>
#include <octave/Cell.h>

DEFUN_DLD (celldemo, args, , "Cell Demo")
{
  if (args.length () != 1)
    print_usage ();

  Cell c = args(0).cell_value ();

  octave_value_list retval;
  retval.resize (c.numel ());    // faster code by pre-declaring size

  for (octave_idx_type i = 0; i < c.numel (); i++)
    {
      retval(i) = c(i);          // using operator syntax
      //retval(i) = c.elem (i);  // using method syntax
    }

  return retval;
}
```

Note that cell arrays are used less often in standard oct-files and so the `Cell.h` header file must be explicitly included. The rest of the example extracts the `octave_value`s one by one from the cell array and returns them as individual output arguments. For example:

```
[b1, b2, b3] = celldemo ({1, [1, 2], "test"})
⇒
b1 =  1
b2 =

   1   2

b3 = test
```

### A.1.5 Structures in Oct-Files

A structure in Octave is a map between a number of fields represented and their values. The Standard Template Library `map` class is used, with the pair consisting of a `std::string` and an Octave `Cell` variable.

A simple example demonstrating the use of structures within oct-files is

```
#include <octave/oct.h>
#include <octave/ov-struct.h>
```

```
DEFUN_DLD (structdemo, args, , "Struct Demo")
{
  if (args.length () != 2)
    print_usage ();

  if (! args(0).isstruct ())
    error ("structdemo: ARG1 must be a struct");

  octave_scalar_map arg0 = args(0).scalar_map_value ();
  //octave_map arg0 = args(0).map_value ();

  if (! args(1).is_string ())
    error ("structdemo: ARG2 must be a character string");

  std::string arg1 = args(1).string_value ();

  octave_value tmp = arg0.contents (arg1);
  //octave_value tmp = arg0.contents (arg1)(0);

  if (! tmp.is_defined ())
    error ("structdemo: struct does not have a field named '%s'\n",
           arg1.c_str ());

  octave_scalar_map st;

  st.assign ("selected", tmp);

  return octave_value (st);
}
```

An example of its use is

```
x.a = 1; x.b = "test"; x.c = [1, 2];
structdemo (x, "b")
⇒ selected = test
```

The example above specifically uses the `octave_scalar_map` class which is for representing a single struct. For structure arrays, the `octave_map` class is used instead. The commented code shows how the demo could be modified to handle a structure array. In that case, the `contents` method returns a `Cell` which may have more than one element. Therefore, to obtain the underlying `octave_value` in the single struct example we would write

```
octave_value tmp = arg0.contents (arg1)(0);
```

where the trailing `(0)` is the `()` operator on the `Cell` object. If this were a true structure array with multiple elements we could iterate over the elements using the `()` operator.

Structures are a relatively complex data container and there are more functions available in `oct-map.h` which make coding with them easier than relying on just `contents`.

### A.1.6 Sparse Matrices in Oct-Files

There are three classes of sparse objects that are of interest to the user.

`SparseMatrix`
: A double precision sparse matrix class

`SparseComplexMatrix`
: A complex sparse matrix class

`SparseBoolMatrix`
: A boolean sparse matrix class

All of these classes inherit from the `Sparse<T>` template class, and so all have similar capabilities and usage. The `Sparse<T>` class was based on Octave's `Array<T>` class and users familiar with Octave's `Array` classes will be comfortable with the use of the sparse classes.

The sparse classes will not be entirely described in this section, due to their similarity with the existing `Array` classes. However, there are a few differences due the nature of sparse objects, and these will be described. First, although it is fundamentally possible to have N-dimensional sparse objects, the Octave sparse classes do not allow them at this time; All instances of the sparse classes **must** be 2-dimensional. This means that `SparseMatrix` is actually more similar to Octave's `Matrix` class than it is to the `NDArray` class.

#### A.1.6.1 Array and Sparse Class Differences

The number of elements in a sparse matrix is considered to be the number of nonzero elements, rather than the product of the dimensions. Therefore,

```
SparseMatrix sm;
...
int nnz = sm.nelem ();
```

returns the number of nonzero elements (like the interpreter function `nnz`). If the user really requires the number of elements in the matrix, including the nonzero elements, they should use `numel` rather than `nelem`. Note that for very large matrices, where the product of the two dimensions is larger than the representation of an unsigned int, `numel` can overflow. An example is `speye (1e6)` which will create a matrix with a million rows and columns, but only a million nonzero elements. In this case, the number of rows multiplied by the number of columns is more than two hundred times the maximum value that can be represented by an unsigned 32-bit int. The use of `numel` should, therefore, be avoided unless it is known that it will not overflow.

Extreme care is also required when using the `elem` method or the `()` operator which perform essentially the same function. The reason is that if a sparse object is non-const, then Octave will assume that a request for a zero element in a sparse matrix is in fact a request to create this element so it can be filled. Therefore, a piece of code like

```
SparseMatrix sm;
...
for (int j = 0; j < nc; j++)
  for (int i = 0; i < nr; i++)
    std::cerr << " (" << i << "," << j << "): " << sm(i,j) << "\n";
```

is a great way of turning a sparse matrix into a dense one, and a very slow way at that since it reallocates the sparse object for each zero element in the matrix.

A simple way of preventing the above from happening is to create a temporary constant version of the sparse matrix. Note that only the container for the sparse matrix will be copied, while the actual representation of the data will be shared between the two versions of the sparse matrix; This is not a costly operation. The example above, re-written to prevent sparse-to-dense conversion, is

```
SparseMatrix sm;
...
const SparseMatrix tmp (sm);
for (int j = 0; j < nc; j++)
  for (int i = 0; i < nr; i++)
    std::cerr << " (" << i << "," << j << "): " << tmp(i,j) << "\n";
```

Finally, because the sparse types aren't represented by a contiguous block of memory, the `rwdata` method of `Array<T>` is not available. It is, however, replaced by three separate methods `ridx`, `cidx`, and `data`, that access the raw compressed column format that Octave sparse matrices are stored in. These methods can be used in a manner similar to `elem` to allow the matrix to be accessed or filled. However, it is up to the user to respect the sparse matrix compressed column format or the matrix will become corrupted.

### A.1.6.2 Creating Sparse Matrices in Oct-Files

There are two useful strategies for creating a sparse matrix. The first is to create three vectors representing the row index, column index, and data values, and from these create the matrix. The second alternative is to create a sparse matrix with the appropriate amount of space, and then fill in the values. Both techniques have their advantages and disadvantages.

Below is an example of creating a small sparse matrix using the first technique

```
int nz, nr, nc;
nz = 4, nr = 3, nc = 4;

ColumnVector ridx (nz);
ColumnVector cidx (nz);
ColumnVector data (nz);

ridx(0) = 1; cidx(0) = 1; data(0) = 1;
ridx(1) = 2; cidx(1) = 2; data(1) = 2;
ridx(2) = 2; cidx(2) = 4; data(2) = 3;
ridx(3) = 3; cidx(3) = 4; data(3) = 4;
SparseMatrix sm (data, ridx, cidx, nr, nc);
```

which creates the matrix given in section Section 22.1.1 [Storage of Sparse Matrices], page 727. Note that the compressed matrix format is not used at the time of the creation of the matrix itself, but is used internally.

As discussed in the chapter on Sparse Matrices, the values of the sparse matrix are stored in increasing column-major ordering. Although the data passed by the user need not respect this requirement, pre-sorting the data will significantly speed up creation of the sparse matrix.

The disadvantage of this technique for creating a sparse matrix is that there is a brief time when two copies of the data exist. For extremely memory constrained problems this may not be the best technique for creating a sparse matrix.

The alternative is to first create a sparse matrix with the desired number of nonzero elements and then later fill those elements in. Sample code:

```
int nz, nr, nc;
nz = 4, nr = 3, nc = 4;
SparseMatrix sm (nr, nc, nz);
sm(0,0) = 1; sm(0,1) = 2; sm(1,3) = 3; sm(2,3) = 4;
```

This creates the same matrix as previously. Again, although not strictly necessary, it is significantly faster if the sparse matrix is created and the elements are added in column-major ordering. The reason for this is that when elements are inserted at the end of the current list of known elements then no element in the matrix needs to be moved to allow the new element to be inserted; Only the column indices need to be updated.

There are a few further points to note about this method of creating a sparse matrix. First, it is possible to create a sparse matrix with fewer elements than are actually inserted in the matrix. Therefore,

```
int nr, nc;
nr = 3, nc = 4;
SparseMatrix sm (nr, nc, 0);
sm(0,0) = 1; sm(0,1) = 2; sm(1,3) = 3; sm(2,3) = 4;
```

is perfectly valid. However, it is a very bad idea because as each new element is added to the sparse matrix the matrix needs to request more space and reallocate memory. This is an expensive operation that will significantly slow this means of creating a sparse matrix. It is possible to create a sparse matrix with excess storage, so having *nz* greater than 4 in this example is also valid. The disadvantage is that the matrix occupies more memory than strictly needed.

Of course, it is not always possible to know the number of nonzero elements prior to filling a matrix. For this reason the additional unused storage of a sparse matrix can be removed after its creation with the `maybe_compress` function. In addition to deallocating unused storage, `maybe_compress` can also remove zero elements from the matrix. The removal of zero elements from the matrix is controlled by setting the argument of the `maybe_compress` function to be `true`. However, the cost of removing the zeros is high because it implies re-sorting the elements. If possible, it is better for the user to avoid adding the unnecessary zeros in the first place. An example of the use of `maybe_compress` is

```
int nz, nr, nc;
nz = 6, nr = 3, nc = 4;

SparseMatrix sm1 (nr, nc, nz);
sm1(0,0) = 1; sm1(0,1) = 2; sm1(1,3) = 3; sm1(2,3) = 4;
sm1.maybe_compress ();   // No zero elements were added

SparseMatrix sm2 (nr, nc, nz);
sm2(0,0) = 1; sm2(0,1) = 2; sm(0,2) = 0; sm(1,2) = 0;
sm1(1,3) = 3; sm1(2,3) = 4;
sm2.maybe_compress (true);  // Zero elements were added
```

The use of the `maybe_compress` function should be avoided if possible as it will slow the creation of the matrix.

A third means of creating a sparse matrix is to work directly with the data in compressed row format. An example of this advanced technique might be

```
octave_value arg;
...
int nz, nr, nc;
nz = 6, nr = 3, nc = 4;   // Assume we know the max # nz
SparseMatrix sm (nr, nc, nz);
Matrix m = arg.matrix_value ();

int ii = 0;
sm.cidx (0) = 0;
for (int j = 1; j < nc; j++)
  {
    for (int i = 0; i < nr; i++)
      {
        double tmp = m(i,j);
        if (tmp != 0.)
          {
            sm.data(ii) = tmp;
            sm.ridx(ii) = i;
            ii++;
          }
      }
    sm.cidx(j+1) = ii;
 }
sm.maybe_compress ();  // If don't know a priori the final # of nz.
```

which is probably the most efficient means of creating a sparse matrix.

Finally, it may sometimes arise that the amount of storage initially created is insufficient to completely store the sparse matrix. Therefore, the method `change_capacity` exists to reallocate the sparse memory. The above example would then be modified as

```
octave_value arg;
...
```

```
int nz, nr, nc;
nz = 6, nr = 3, nc = 4;   // Guess the number of nz elements
SparseMatrix sm (nr, nc, nz);
Matrix m = arg.matrix_value ();

int ii = 0;
sm.cidx (0) = 0;
for (int j = 1; j < nc; j++)
  {
    for (int i = 0; i < nr; i++)
      {
        double tmp = m(i,j);
        if (tmp != 0.)
          {
            if (ii == nz)
              {
                nz += 2;   // Add 2 more elements
                sm.change_capacity (nz);
              }
            sm.data(ii) = tmp;
            sm.ridx(ii) = i;
            ii++;
          }
      }
    sm.cidx(j+1) = ii;
 }
sm.maybe_compress ();  // If don't know a priori the final # of nz.
```

Note that both increasing and decreasing the number of nonzero elements in a sparse matrix is expensive as it involves memory reallocation. Also because parts of the matrix, though not its entirety, exist as old and new copies at the same time, additional memory is needed. Therefore, if possible avoid changing capacity.

### A.1.6.3 Using Sparse Matrices in Oct-Files

Most of the same operators and functions for sparse matrices that are available from the Octave interpreter are also available within oct-files. The basic means of extracting a sparse matrix from an `octave_value`, and returning it as an `octave_value`, can be seen in the following example.

```
octave_value_list retval;

SparseMatrix sm = args(0).sparse_matrix_value ();
SparseComplexMatrix scm = args(1).sparse_complex_matrix_value ();
SparseBoolMatrix sbm = args(2).sparse_bool_matrix_value ();
...
retval(2) = sbm;
retval(1) = scm;
retval(0) = sm;
```

The conversion to an `octave_value` is handled by the sparse `octave_value` constructors, and so no special care is needed.

### A.1.7 Accessing Global Variables in Oct-Files

Global variables allow variables in the global scope to be accessed. Global variables can be accessed within oct-files by using the support functions `global_varval` and `global_assign` from the current interpreter's symbol table. Both functions take as first argument a string representing the variable name to be obtained or assigned. The second argument of `global_assign` is the value to be assigned. An example of the use of these two functions is

```
#include <octave/oct.h>
#include <octave/interpreter.h>

DEFMETHOD_DLD (globaldemo, interp, args, , "Global Demo")
{
  if (args.length () != 1)
    print_usage ();

  octave_value retval;

  std::string s = args(0).string_value ();

  octave::symbol_table& symtab = interp.get_symbol_table ();

  octave_value tmp = symtab.global_varval (s);

  if (tmp.is_defined ())
    retval = tmp;
  else
    retval = "Global variable not found";

  symtab.global_assign ("a", 42.0);

  return retval;
}
```

An example of its use is

```
global a b
b = 10;
globaldemo ("b")
⇒ 10
globaldemo ("c")
⇒ "Global variable not found"
num2str (a)
⇒ 42
```

### A.1.8 Calling Octave Functions from Oct-Files

There is often a need to be able to call another Octave function from within an oct-file, and there are many examples of such within Octave itself. For example, the `quad` function is an oct-file that calculates the definite integral by quadrature over a user-supplied function.

There are also many ways in which a function could be given as input. It might be passed as one of

1. Function Handle
2. Anonymous Function Handle
3. String

The code below demonstrates all four methods of passing a function to an oct-file.

```
#include <octave/oct.h>
#include <octave/parse.h>

DEFMETHOD_DLD (funcdemo, interp, args, nargout, "Function Demo")
{
  int nargin = args.length ();

  if (nargin < 2)
    print_usage ();

  octave_value_list newargs;

  for (octave_idx_type i = nargin - 1; i > 0; i--)
    newargs(i-1) = args(i);

  octave_value_list retval;

  if (args(0).is_function_handle () || args(0).is_inline_function ()
      || args(0).is_string ())
    retval = interp.feval (args(0), newargs, nargout);
  else
    error ("funcdemo: INPUT must be string, inline, or function handle");

  return retval;
}
```

The first input to the demonstration code is a user-supplied function and the remaining arguments are all passed to the function.

```
funcdemo (@sin, 1)
⇒ 0.84147
funcdemo (@(x) sin (x), 1)
⇒ 0.84147
funcdemo ("sin", 1)
⇒ 0.84147
funcdemo (@atan2, 1, 1)
⇒ 0.78540
```

When the user function is passed as a string the treatment of the function is different. In some cases it is necessary to have the user supplied function as an `octave_function` object. In that case the string argument can be used to create a temporary function as demonstrated below.

```
std::octave fcn_name = unique_symbol_name ("__fcn__");
std::string fcode = "function y = ";
fcode.append (fcn_name);
fcode.append ("(x) y = ");
fcn = extract_function (args(0), "funcdemo", fcn_name,
                        fcode, "; endfunction");
...
if (fcn_name.length ())
  clear_function (fcn_name);
```

There are two important things to know in this case. First, the number of input arguments to the user function is fixed, and in the above example is a single argument. Second, to avoid leaving the temporary function in the Octave symbol table it should be cleared after use. Also, by convention all internal function names begin and end with the character sequence '`__`'.

### A.1.9 Calling External Code from Oct-Files

Linking external C code to Octave is relatively simple, as the C functions can easily be called directly from C++. One possible issue is that the declarations of the external C functions may need to be explicitly defined as C functions to the compiler. If the declarations of the external C functions are in the header `foo.h`, then the tactic to ensure that the C++ compiler treats these declarations as C code is

```
#ifdef __cplusplus
extern "C"
{
#endif
#include "foo.h"
#ifdef __cplusplus
}  /* end extern "C" */
#endif
```

When calling functions that are implemented in Fortran code, some peculiarities have to be taken into account. Symbol names in Fortran are case-insensitive, and depending on the used Fortran compiler, function names are either exported with all lowercase or with all uppercase characters. Additionally, some compilers append none, one or two underscores "`_`" at the end of exported function names. This is called "name-mangling".

Octave supplies macros that allow writing code that automatically handles the name-mangling for a number of different Fortran compilers. These macros are `F77_FUNC` and `F77_FUNC_`. The former should be used for Fortran functions that do not contain any underscores in their name. The latter should be used for Fortran functions with underscores in their names. Both macros take two arguments: The first is the Fortran function name in all lowercase characters. The second is the same Fortran function name in all uppercase characters.

Additionally to the name-mangling, different compilers are using different calling conventions for some types. Octave defines the following preprocessor macros to allow writing code that can be used with different Fortran calling conventions.

Note that we don't attempt to handle Fortran functions, we always use subroutine wrappers for them and pass the return value as an extra argument.

Use the following macros to pass character strings from C to Fortran:

```
F77_CHAR_ARG(x)
F77_CONST_CHAR_ARG(x)
F77_CXX_STRING_ARG(x)
F77_CHAR_ARG_LEN(l)
F77_CHAR_ARG_DECL
F77_CONST_CHAR_ARG_DECL
F77_CHAR_ARG_LEN_DECL
```

Use the following macros to write C-language functions that accept Fortran-style character strings:

```
F77_CHAR_ARG_DEF(s, len)
F77_CONST_CHAR_ARG_DEF(s, len)
F77_CHAR_ARG_LEN_DEF(len)
F77_CHAR_ARG_USE(s)
F77_CHAR_ARG_LEN_USE(s, len)
```

Use the following macros for Fortran types in C++ code:

`F77_INT4` Equivalent to Fortran `INTEGER*4` type

`F77_DBLE` Equivalent to Fortran `DOUBLE PRECISION` type

`F77_REAL` Equivalent to Fortran `REAL` type

`F77_CMPLX`
: Equivalent to Fortran `COMPLEX` type

`F77_DBLE_CMPLX`
: Equivalent to Fortran `DOUBLE COMPLEX` type

`F77_LOGICAL`
: Equivalent to Fortran `LOGICAL` type

`F77_RET_T`
: Return type of a C++ function that acts like a Fortran subroutine.

Use the following macros to return from C-language functions that are supposed to act like Fortran subroutines. `F77_NORETURN` is intended to be used as the last statement of such a function that has been tagged with a `"noreturn"` attribute.

```
F77_RETURN(retval)
F77_NORETURN(retval)
```

The underlying Fortran code should use the `XSTOPX` function to replace the Fortran `STOP` function. `XSTOPX` uses the Octave exception handler to treat failing cases in the Fortran code explicitly. Note that Octave supplies its own replacement BLAS `XERBLA` function, which uses `XSTOPX`.

The following example shows the inclusion of a Fortran function in an oct-file, where the C++ wrapper is

```
#include <octave/oct.h>
#include <octave/f77-fcn.h>

extern "C"
{
  F77_RET_T
  F77_FUNC (fortransub, FORTRANSUB)
    (const F77_INT&, F77_DBLE*, F77_CHAR_ARG_DECL F77_CHAR_ARG_LEN_DECL);
}

DEFUN_DLD (fortrandemo, args, , "Fortran Demo")
{
  if (args.length () != 1)
    print_usage ();

  NDArray a = args(0).array_value ();

  double *av = a.rwdata ();
  octave_idx_type na = a.numel ();

  OCTAVE_LOCAL_BUFFER (char, ctmp, 128);

  F77_FUNC (fortransub, FORTRANSUB)
            (na, av, ctmp F77_CHAR_ARG_LEN (128));

  return ovl (a, std::string (ctmp));
}
```

and the Fortran function is

```
      subroutine fortransub (n, a, s)
      implicit none
      character*(*) s
      real*8 a(*)
      integer*4 i, n, ioerr
      do i = 1, n
        if (a(i) .eq. 0d0) then
          call xstopx ('fortransub: divide by zero')
        else
          a(i) = 1d0 / a(i)
        endif
      enddo
      write (unit = s, fmt = '(a,i3,a,a)', iostat = ioerr)
     $       'There are ', n,
     $       ' values in the input vector', char(0)
      if (ioerr .ne. 0) then
```

```
      call xstopx ('fortransub: error writing string')
    endif
    return
    end
```

This example demonstrates most of the features needed to link to an external Fortran function, including passing arrays and strings, as well as exception handling. Both the Fortran and C++ files need to be compiled in order for the example to work.

```
mkoctfile fortrandemo.cc fortransub.f
[b, s] = fortrandemo (1:3)
⇒
  b = 1.00000   0.50000   0.33333
  s = There are   3 values in the input vector
[b, s] = fortrandemo (0:3)
error: fortrandemo: fortransub: divide by zero
```

### A.1.10 Allocating Local Memory in Oct-Files

Allocating memory within an oct-file might seem easy, as the C++ new/delete operators can be used. However, in that case great care must be taken to avoid memory leaks. The preferred manner in which to allocate memory for use locally is to use the `OCTAVE_LOCAL_BUFFER` macro. An example of its use is

```
OCTAVE_LOCAL_BUFFER (double, tmp, len)
```

that returns a pointer `tmp` of type `double *` of length `len`.

In this case, Octave itself will worry about reference counting and variable scope and will properly free memory without programmer intervention.

### A.1.11 Input Parameter Checking in Oct-Files

Because oct-files are compiled functions they open up the possibility of crashing Octave through careless function calls or memory faults. It is quite important that each and every function have a sufficient level of parameter checking to ensure that Octave behaves well.

The minimum requirement, as previously discussed, is to check the number of input arguments before using them to avoid referencing a nonexistent argument. However, in some cases this might not be sufficient as the underlying code imposes further constraints. For example, an external function call might be undefined if the input arguments are not integers, or if one of the arguments is zero, or if the input is complex and a real value was expected. Therefore, oct-files often need additional input parameter checking.

There are several functions within Octave that can be useful for the purposes of parameter checking. These include the methods of the `octave_value` class like `is_real_matrix`, `is_numeric_type`, etc. (see `ov.h`). Often, with a knowledge of the Octave m-file language, you can guess at what the corresponding C++ routine will. In addition there are some more specialized input validation functions of which a few are demonstrated below.

```
#include <octave/oct.h>

DEFUN_DLD (paramdemo, args, nargout, "Parameter Check Demo")
{
  if (args.length () != 1)
```

```
    print_usage ();

  NDArray m = args(0).array_value ();

  double min_val = -10.0;
  double max_val = 10.0;

  octave_stdout << "Properties of input array:\n";

  if (m.any_element_is_negative ())
    octave_stdout << "  includes negative values\n";

  if (m.any_element_is_inf_or_nan ())
    octave_stdout << "  includes Inf or NaN values\n";

  if (m.any_element_not_one_or_zero ())
    octave_stdout << "  includes other values than 1 and 0\n";

  if (m.all_elements_are_int_or_inf_or_nan ())
    octave_stdout << "  includes only int, Inf or NaN values\n";

  if (m.all_integers (min_val, max_val))
    octave_stdout << "  includes only integers in [-10,10]\n";

  return octave_value_list ();
}
```

An example of its use is:

```
paramdemo ([1, 2, NaN, Inf])
⇒ Properties of input array:
     includes Inf or NaN values
     includes other values than 1 and 0
     includes only int, Inf or NaN values
```

### A.1.12 Exception and Error Handling in Oct-Files

Another important feature of Octave is its ability to react to the user typing `Control-C` during extended calculations. This ability is based on the C++ exception handler, where memory allocated by the C++ new/delete methods is automatically released when the exception is treated. When writing an oct-file which may run for a long time the programmer must periodically use the macro `OCTAVE_QUIT`, in order to allow Octave to check and possibly respond to a user typing `Control-C`. For example:

```
for (octave_idx_type i = 0; i < a.nelem (); i++)
  {
    OCTAVE_QUIT;
    b.elem (i) = 2. * a.elem (i);
  }
```

The presence of the `OCTAVE_QUIT` macro in the inner loop allows Octave to detect and acknowledge a `Control-C` key sequence. Without this macro, the user must either wait for the oct-file function to return before the interrupt is processed, or the user must press `Control-C` three times which will force Octave to exit completely.

The `OCTAVE_QUIT` macro does impose a very small performance penalty; For loops that are known to be small it may not make sense to include `OCTAVE_QUIT`.

When creating an oct-file that uses an external library, the function might spend a significant portion of its time in the external library. It is not generally possible to use the `OCTAVE_QUIT` macro in this case. The alternative code in this case is

```
BEGIN_INTERRUPT_IMMEDIATELY_IN_FOREIGN_CODE;
... some code that calls a "foreign" function ...
END_INTERRUPT_IMMEDIATELY_IN_FOREIGN_CODE;
```

The disadvantage of this is that if the foreign code allocates any memory internally, then this memory might be lost during an interrupt, without being deallocated. Therefore, ideally Octave itself should allocate any memory that is needed by the foreign code, with either the `rwdata` method or the `OCTAVE_LOCAL_BUFFER` macro.

The Octave `unwind_protect` mechanism (see Section 10.8 [The unwind_protect Statement], page 202) can also be used in oct-files. In conjunction with the exception handling of Octave, it ensures that certain recovery code is always run even if an exception occurs. An example of the use of this mechanism is

```
#include <octave/oct.h>
#include <octave/unwind-prot.h>

void
my_err_handler (const char *fmt, ...)
{
  // Do nothing!!
}

void
my_err_with_id_handler (const char *id, const char *fmt, ...)
{
  // Do nothing!!
}

DEFUN_DLD (unwinddemo, args, nargout, "Unwind Demo")
{
  if (args.length () < 2)
    print_usage ();

  NDArray a = args(0).array_value ();
  NDArray b = args(1).array_value ();

  // Create unwind_action objects.  At the end of the enclosing scope,
  // destructors for these objects will call the given functions with
```

```
  // the specified arguments.

  octave::unwind_action restore_warning_handler
    (set_liboctave_warning_handler, current_liboctave_warning_handler);

  octave::unwind_action restore_warning_with_id_handler
    (set_liboctave_warning_with_id_handler,
     current_liboctave_warning_with_id_handler);

  set_liboctave_warning_handler (my_err_handler);
  set_liboctave_warning_with_id_handler (my_err_with_id_handler);

  return octave_value (quotient (a, b));
}
```

As can be seen in the example:

```
unwinddemo (1, 0)
⇒ Inf
1 / 0
⇒ warning: division by zero
   Inf
```

The warning for division by zero (and in fact all warnings) are disabled in the `unwinddemo` function.

### A.1.13 Documentation and Testing of Oct-Files

The documentation for an oct-file is contained in the fourth string parameter of the `DEFUN_DLD` macro. This string can be formatted in the same manner as the help strings for user functions, however there are some issues that are particular to the formatting of help strings within oct-files.

The major issue is that the help string will typically be longer than a single line of text, and so the formatting of long multi-line help strings needs to be taken into account. There are several possible solutions, but the most common is illustrated in the following example,

```
DEFUN_DLD (do_what_i_want, args, nargout,
  "-*- texinfo -*-\n\
@deftypefn {} {} do_what_i_say (@var{n})\n\
A function that does what the user actually wants rather\n\
than what they requested.\n\
@end deftypefn")
{
...
}
```

where each line of text is terminated by `\n\` which is an embedded newline in the string together with a C++ string continuation character. Note that the final `\` must be the last character on the line.

Octave also includes the ability to embed test and demonstration code for a function within the code itself (see Appendix B [Test and Demo Functions], page 1151). This can

be used from within oct-files (or in fact any file) with certain provisos. First, the test and demo functions of Octave look for `%!` as the first two characters of a line to identify test and demonstration code. This is a requirement for oct-files as well. In addition, the test and demonstration code must be wrapped in a comment block to avoid it being interpreted by the compiler. Finally, the Octave test and demonstration code must have access to the original source code of the oct-file—not just the compiled code—as the tests are stripped from the compiled code. An example in an oct-file might be

```
/*
%!assert (sin ([1,2]), [sin(1),sin(2)])
%!error (sin ())
%!error (sin (1,1))
*/
```

## A.2 Mex-Files

Octave includes an interface to allow legacy mex-files to be compiled and used with Octave. This interface can also be used to share compiled code between Octave and MATLAB users. However, as mex-files expose MATLAB's internal API, and the internal structure of Octave is different, a mex-file can never have the same performance in Octave as the equivalent oct-file. In particular, to support the manner in which variables are passed to mex functions there are a significant number of additional copies of memory blocks when invoking or returning from a mex-file function. For this reason, it is recommended that any new code be written with the oct-file interface previously discussed.

### A.2.1 Getting Started with Mex-Files

The basic command to build a mex-file is either `mkoctfile --mex` or `mex`. The first command can be used either from within Octave or from the command line. To avoid issues with MATLAB's own `mex` command, the use of the command `mex` is limited to within Octave. Compiled mex-files have the extension `.mex`.

`mex` [*-options*] *file* . . .
`status = mex (...)`

Compile source code written in C, C++, or Fortran, to a MEX file.

*status* is the return status of the `mkoctfile` function.

If the compilation fails, and the output argument is not requested, an error is raised. If the programmer requests *status*, however, Octave will merely issue a warning and it is the programmer's responsibility to verify the command was successful.

This is equivalent to `mkoctfile --mex [-options] file`.

**See also:** [mkoctfile], page 1098, [mexext], page 1123.

*ext* `= mexext ()`

Return the filename extension used for MEX files.

Programming Note: Octave uses the extension `mex` for all MEX files regardless of the operating system (Linux, Windows, Apple) or the bit-width (32-bit or 64-bit) of the hardware.

**See also:** [mex], page 1123.

Consider the following short example:

```
#include "mex.h"

void
mexFunction (int nlhs, mxArray *plhs[],
             int nrhs, const mxArray *prhs[])
{
  mexPrintf ("Hello, World!\n");

  mexPrintf ("I have %d inputs and %d outputs\n", nrhs, nlhs);

  /* Return empty matrices for any outputs */
  int i;
  for (i = 0; i < nlhs; i++)
    plhs[i] = mxCreateDoubleMatrix (0, 0, mxREAL);
}
```

The first line `#include "mex.h"` makes available all of the definitions necessary for a mex-file. One important difference between Octave and MATLAB is that the header file `"matrix.h"` is implicitly included through the inclusion of `"mex.h"`. This is necessary to avoid a conflict with the Octave file `"Matrix.h"` for operating systems and compilers that don't distinguish between filenames in upper and lower case.

The entry point into the mex-file is defined by `mexFunction`. The function takes four arguments:

1. The number of return arguments (# of left-hand side args).
2. An array of pointers to return arguments.
3. The number of input arguments (# of right-hand side args).
4. An array of pointers to input arguments.

Note that the function name definition is not explicitly included in `mexFunction` and so there can only be a single `mexFunction` entry point per file. Instead, the name of the function as seen in Octave is determined by the name of the mex-file itself minus the extension. If the above function is in the file `myhello.c`, it can be compiled with

```
mkoctfile --mex myhello.c
```

which creates a file `myhello.mex`. The function can then be run from Octave as

```
myhello (1,2,3)
⇒ Hello, World!
⇒ I have 3 inputs and 0 outputs
```

It should be noted that the mex-file contains no help string. To document mex-files, there should exist an m-file in the same directory as the mex-file itself. Taking the above as an example, there would need to be a file `myhello.m` which might contain the text

```
%MYHELLO Simple test of the functionality of a mex-file.
```

In this case, the function that will be executed within Octave will be given by the mex-file, while the help string will come from the m-file. This can also be useful to allow a sample implementation of the mex-file within the Octave language itself for testing purposes.

Although there cannot be multiple entry points in a single mex-file, one can use the `mexFunctionName` function to determine what name the mex-file was called with. This can be used to alter the behavior of the mex-file based on the function name. For example, if

```
#include "mex.h"

void
mexFunction (int nlhs, mxArray *plhs[],
             int nrhs, const mxArray *prhs[])
{
  const char *nm;

  nm = mexFunctionName ();
  mexPrintf ("You called function: %s\n", nm);
  if (strcmp (nm, "myfunc") == 0)
    mexPrintf ("This is the principal function\n", nm);

  return;
}
```

is in the file `myfunc.c`, and is compiled with

```
mkoctfile --mex myfunc.c
ln -s myfunc.mex myfunc2.mex
```

then as can be seen by

```
myfunc ()
⇒ You called function: myfunc
    This is the principal function
myfunc2 ()
⇒ You called function: myfunc2
```

the behavior of the mex-file can be altered depending on the function's name.

Although the user should only include `mex.h` in their code, Octave declares additional functions, typedefs, etc., available to the user to write mex-files in the headers `mexproto.h` and `mxarray.h`.

### A.2.2 Working with Matrices and Arrays in Mex-Files

The basic mex type of all variables is `mxArray`. Any object, such as a matrix, cell array, or structure, is stored in this basic type. `mxArray` serves essentially the same purpose as the `octave_value` class in oct-files in that it acts as a container for all the more specialized types.

The `mxArray` structure contains at a minimum, the name of the variable it represents, its dimensions, its type, and whether the variable is real or complex. It can also contain a number of additional fields depending on the type of the `mxArray`. There are a number of functions to create `mxArray` structures, including `mxCreateDoubleMatrix`, `mxCreateCellArray`, `mxCreateSparse`, and the generic `mxCreateNumericArray`.

The basic function to access the data in an array is `mxGetPr`. Because the mex interface assumes that real and imaginary parts of a complex array are stored separately, there is an equivalent function `mxGetPi` that gets the imaginary part. Both of these functions

are only for use with double precision matrices. The generic functions `mxGetData` and `mxGetImagData` perform the same operation for all matrix types. For example:

```
mxArray *m;
mwSize *dims;
UINT32_T *pr;

dims = (mwSize *) mxMalloc (2 * sizeof (mwSize));
dims[0] = 2; dims[1] = 2;
m = mxCreateNumericArray (2, dims, mxUINT32_CLASS, mxREAL);
pr = (UINT32_T *) mxGetData (m);
```

There are also the functions `mxSetPr`, etc., that perform the inverse, and set the data of an array to use the block of memory pointed to by the argument of `mxSetPr`.

Note the type `mwSize` used above, and also `mwIndex`, are defined as the native precision of the indexing in Octave on the platform on which the mex-file is built. This allows both 32- and 64-bit platforms to support mex-files. `mwSize` is used to define array dimensions and the maximum number or elements, while `mwIndex` is used to define indexing into arrays.

An example that demonstrates how to work with arbitrary real or complex double precision arrays is given by the file `mypow2.c` shown below.

```
#include "mex.h"

void
mexFunction (int nlhs, mxArray *plhs[],
             int nrhs, const mxArray *prhs[])
{
  mwSize n;
  mwIndex i;
  double *vri, *vro;

  if (nrhs != 1 || ! mxIsDouble (prhs[0]))
    mexErrMsgTxt ("ARG1 must be a double matrix");

  n = mxGetNumberOfElements (prhs[0]);
  plhs[0] = mxCreateNumericArray (mxGetNumberOfDimensions (prhs[0]),
                                  mxGetDimensions (prhs[0]),
                                  mxGetClassID (prhs[0]),
                                  mxIsComplex (prhs[0]));
  vri = mxGetPr (prhs[0]);
  vro = mxGetPr (plhs[0]);

  if (mxIsComplex (prhs[0]))
    {
      double *vii, *vio;
      vii = mxGetPi (prhs[0]);
      vio = mxGetPi (plhs[0]);

      for (i = 0; i < n; i++)
```

```
          {
            vro[i] = vri[i] * vri[i] - vii[i] * vii[i];
            vio[i] = 2 * vri[i] * vii[i];
          }
      }
    else
      {
        for (i = 0; i < n; i++)
          vro[i] = vri[i] * vri[i];
      }
}
```

An example of its use is

```
b = randn (4,1) + 1i * randn (4,1);
all (b.^2 == mypow2 (b))
⇒ 1
```

The example above uses the functions `mxGetDimensions`, `mxGetNumberOfElements`, and `mxGetNumberOfDimensions` to work with the dimensions of multi-dimensional arrays. The functions `mxGetM`, and `mxGetN` are also available to find the number of rows and columns in a 2-D matrix (MxN matrix).

### A.2.3 Character Strings in Mex-Files

As mex-files do not make the distinction between single and double quoted strings that Octave does, there is perhaps less complexity in the use of strings and character matrices. An example of their use that parallels the demo in `stringdemo.cc` is given in the file `mystring.c`, as shown below.

```
#include <string.h>
#include "mex.h"

void
mexFunction (int nlhs, mxArray *plhs[],
             int nrhs, const mxArray *prhs[])
{
  mwSize m, n;
  mwIndex i, j;
  mxChar *pi, *po;

  if (nrhs != 1 || ! mxIsChar (prhs[0])
      || mxGetNumberOfDimensions (prhs[0]) > 2)
    mexErrMsgTxt ("ARG1 must be a char matrix");

  m = mxGetM (prhs[0]);
  n = mxGetN (prhs[0]);
  pi = mxGetChars (prhs[0]);
  plhs[0] = mxCreateNumericMatrix (m, n, mxCHAR_CLASS, mxREAL);
  po = mxGetChars (plhs[0]);

  for (j = 0; j < n; j++)
    for (i = 0; i < m; i++)
      po[j*m + m - 1 - i] = pi[j*m + i];
}
```

An example of its expected output is

```
mystring (["First String"; "Second String"])
⇒ Second String
   First String
```

Other functions in the mex interface for handling character strings are `mxCreateString`, `mxArrayToString`, and `mxCreateCharMatrixFromStrings`. In a mex-file, a character string is considered to be a vector rather than a matrix. This is perhaps an arbitrary distinction as the data in the `mxArray` for the matrix is consecutive in any case.

### A.2.4 Cell Arrays with Mex-Files

One can perform exactly the same operations on Cell arrays in mex-files as in oct-files. An example that duplicates the function of the `celldemo.cc` oct-file in a mex-file is given by `mycell.c` as shown below.

```
#include "mex.h"

void
mexFunction (int nlhs, mxArray *plhs[],
             int nrhs, const mxArray *prhs[])
{
  mwSize n;
  mwIndex i;

  if (nrhs != 1 || ! mxIsCell (prhs[0]))
    mexErrMsgTxt ("ARG1 must be a cell");

  n = mxGetNumberOfElements (prhs[0]);
  n = (n > nlhs ? nlhs : n);

  for (i = 0; i < n; i++)
    plhs[i] = mxDuplicateArray (mxGetCell (prhs[0], i));
}
```

The output is identical to the oct-file version as well.

```
[b1, b2, b3] = mycell ({1, [1, 2], "test"})
⇒
b1 =  1
b2 =

   1   2

b3 = test
```

Note in the example the use of the `mxDuplicateArray` function. This is needed as the `mxArray` pointer returned by `mxGetCell` might be deallocated. The inverse function to `mxGetCell`, used for setting Cell values, is `mxSetCell` and is defined as

```
void mxSetCell (mxArray *ptr, int idx, mxArray *val);
```

Finally, to create a cell array or matrix, the appropriate functions are

```
mxArray *mxCreateCellArray (int ndims, const int *dims);
mxArray *mxCreateCellMatrix (int m, int n);
```

### A.2.5 Structures with Mex-Files

The basic function to create a structure in a mex-file is `mxCreateStructMatrix` which creates a structure array with a two dimensional matrix, or `mxCreateStructArray`.

```
mxArray *mxCreateStructArray (int ndims, int *dims,
                              int num_keys,
                              const char **keys);
mxArray *mxCreateStructMatrix (int rows, int cols,
                               int num_keys,
                               const char **keys);
```

Accessing the fields of the structure can then be performed with `mxGetField` and `mxSetField` or alternatively with the `mxGetFieldByNumber` and `mxSetFieldByNumber` functions.

```
mxArray *mxGetField (const mxArray *ptr, mwIndex index,
                     const char *key);
mxArray *mxGetFieldByNumber (const mxArray *ptr,
                             mwIndex index, int key_num);
void mxSetField (mxArray *ptr, mwIndex index,
                 const char *key, mxArray *val);
void mxSetFieldByNumber (mxArray *ptr, mwIndex index,
                         int key_num, mxArray *val);
```

A difference between the oct-file interface to structures and the mex-file version is that the functions to operate on structures in mex-files directly include an `index` over the elements of the arrays of elements per `field`; Whereas, the oct-file structure includes a Cell Array per field of the structure.

An example that demonstrates the use of structures in a mex-file can be found in the file `mystruct.c` shown below.

```
#include "mex.h"

void
mexFunction (int nlhs, mxArray *plhs[],
             int nrhs, const mxArray *prhs[])
{
  int i;
  mwIndex j;
  mxArray *v;
  const char *keys[] = { "this", "that" };

  if (nrhs != 1 || ! mxIsStruct (prhs[0]))
    mexErrMsgTxt ("ARG1 must be a struct");

  for (i = 0; i < mxGetNumberOfFields (prhs[0]); i++)
    for (j = 0; j < mxGetNumberOfElements (prhs[0]); j++)
      {
        mexPrintf ("field %s(%d) = ", mxGetFieldNameByNumber (prhs[0], i), j);
        v = mxGetFieldByNumber (prhs[0], j, i);
        mexCallMATLAB (0, NULL, 1, &v, "disp");
      }
```

```
  v = mxCreateStructMatrix (2, 2, 2, keys);

  mxSetFieldByNumber (v, 0, 0, mxCreateString ("this1"));
  mxSetFieldByNumber (v, 0, 1, mxCreateString ("that1"));
  mxSetFieldByNumber (v, 1, 0, mxCreateString ("this2"));
  mxSetFieldByNumber (v, 1, 1, mxCreateString ("that2"));
  mxSetFieldByNumber (v, 2, 0, mxCreateString ("this3"));
  mxSetFieldByNumber (v, 2, 1, mxCreateString ("that3"));
  mxSetFieldByNumber (v, 3, 0, mxCreateString ("this4"));
  mxSetFieldByNumber (v, 3, 1, mxCreateString ("that4"));

  if (nlhs)
    plhs[0] = v;
}
```

An example of the behavior of this function within Octave is then

```
a(1).f1 = "f11"; a(1).f2 = "f12";
a(2).f1 = "f21"; a(2).f2 = "f22";
b = mystruct (a);
⇒  field f1(0) = f11
    field f1(1) = f21
    field f2(0) = f12
    field f2(1) = f22
b
⇒ 2x2 struct array containing the fields:

     this
     that

b(3)
⇒ scalar structure containing the fields:

     this = this3
     that = that3
```

### A.2.6 Sparse Matrices with Mex-Files

The Octave format for sparse matrices is identical to the mex format in that it is a compressed column sparse format. Also, in both implementations sparse matrices are required to be two-dimensional. The only difference of importance to the programmer is that the real and imaginary parts of the matrix are stored separately.

The mex-file interface, in addition to using `mxGetM`, `mxGetN`, `mxSetM`, `mxSetN`, `mxGetPr`, `mxGetPi`, `mxSetPr`, and `mxSetPi`, also supplies the following functions.

```
mwIndex *mxGetIr (const mxArray *ptr);
mwIndex *mxGetJc (const mxArray *ptr);
mwSize mxGetNzmax (const mxArray *ptr);

void mxSetIr (mxArray *ptr, mwIndex *ir);
void mxSetJc (mxArray *ptr, mwIndex *jc);
void mxSetNzmax (mxArray *ptr, mwSize nzmax);
```

`mxGetNzmax` gets the maximum number of elements that can be stored in the sparse matrix. This is not necessarily the number of nonzero elements in the sparse matrix. `mxGetJc` returns an array with one additional value than the number of columns in the sparse matrix. The difference between consecutive values of the array returned by `mxGetJc` define the number of nonzero elements in each column of the sparse matrix. Therefore,

```
mwSize nz, n;
mwIndex *Jc;
mxArray *m;
...
n = mxGetN (m);
Jc = mxGetJc (m);
nz = Jc[n];
```

returns the actual number of nonzero elements stored in the matrix in `nz`. As the arrays returned by `mxGetPr` and `mxGetPi` only contain the nonzero values of the matrix, we also need a pointer to the rows of the nonzero elements, and this is given by `mxGetIr`. A complete example of the use of sparse matrices in mex-files is given by the file `mysparse.c` shown below.

```
#include "mex.h"

void
mexFunction (int nlhs, mxArray *plhs[],
             int nrhs, const mxArray *prhs[])
{
  mwSize m, n, nz;
  mxArray *v;
  mwIndex i;
  double *pr, *pi;
  double *pr2, *pi2;
  mwIndex *ir, *jc;
  mwIndex *ir2, *jc2;

  if (nrhs != 1 || ! mxIsSparse (prhs[0]))
    mexErrMsgTxt ("ARG1 must be a sparse matrix");

  m = mxGetM (prhs[0]);
  n = mxGetN (prhs[0]);
  nz = mxGetNzmax (prhs[0]);

  if (mxIsComplex (prhs[0]))
    {
      mexPrintf ("Matrix is %d-by-%d complex sparse matrix", m, n);
      mexPrintf (" with %d elements\n", nz);

      pr = mxGetPr (prhs[0]);
      pi = mxGetPi (prhs[0]);
      ir = mxGetIr (prhs[0]);
```

```
    jc = mxGetJc (prhs[0]);

    i = n;
    while (jc[i] == jc[i-1] && i != 0) i--;

    mexPrintf ("last nonzero element (%d, %d) = (%g, %g)\n",
               ir[nz-1]+ 1, i, pr[nz-1], pi[nz-1]);

    v = mxCreateSparse (m, n, nz, mxCOMPLEX);
    pr2 = mxGetPr (v);
    pi2 = mxGetPi (v);
    ir2 = mxGetIr (v);
    jc2 = mxGetJc (v);

    for (i = 0; i < nz; i++)
      {
        pr2[i] = 2 * pr[i];
        pi2[i] = 2 * pi[i];
        ir2[i] = ir[i];
      }
    for (i = 0; i < n + 1; i++)
      jc2[i] = jc[i];

    if (nlhs > 0)
      plhs[0] = v;
  }
else if (mxIsLogical (prhs[0]))
  {
    mxLogical *pbr, *pbr2;
    mexPrintf ("Matrix is %d-by-%d logical sparse matrix", m, n);
    mexPrintf (" with %d elements\n", nz);

    pbr = mxGetLogicals (prhs[0]);
    ir = mxGetIr (prhs[0]);
    jc = mxGetJc (prhs[0]);

    i = n;
    while (jc[i] == jc[i-1] && i != 0) i--;
    mexPrintf ("last nonzero element (%d, %d) = %d\n",
               ir[nz-1]+ 1, i, pbr[nz-1]);

    v = mxCreateSparseLogicalMatrix (m, n, nz);
    pbr2 = mxGetLogicals (v);
    ir2 = mxGetIr (v);
    jc2 = mxGetJc (v);

    for (i = 0; i < nz; i++)
```

```
        {
          pbr2[i] = pbr[i];
          ir2[i] = ir[i];
        }
      for (i = 0; i < n + 1; i++)
        jc2[i] = jc[i];

      if (nlhs > 0)
        plhs[0] = v;
    }
  else
    {
      mexPrintf ("Matrix is %d-by-%d real sparse matrix", m, n);
      mexPrintf (" with %d elements\n", nz);

      pr = mxGetPr (prhs[0]);
      ir = mxGetIr (prhs[0]);
      jc = mxGetJc (prhs[0]);

      i = n;
      while (jc[i] == jc[i-1] && i != 0) i--;
      mexPrintf ("last nonzero element (%d, %d) = %g\n",
                 ir[nz-1]+ 1, i, pr[nz-1]);

      v = mxCreateSparse (m, n, nz, mxREAL);
      pr2 = mxGetPr (v);
      ir2 = mxGetIr (v);
      jc2 = mxGetJc (v);

      for (i = 0; i < nz; i++)
        {
          pr2[i] = 2 * pr[i];
          ir2[i] = ir[i];
        }
      for (i = 0; i < n + 1; i++)
        jc2[i] = jc[i];

      if (nlhs > 0)
        plhs[0] = v;
    }
}
```

A sample usage of `mysparse` is

```
sm = sparse ([1, 0; 0, pi]);
mysparse (sm)
⇒
Matrix is 2-by-2 real sparse matrix with 2 elements
last nonzero element (2, 2) = 3.14159
```

### A.2.7 Calling Other Functions in Mex-Files

It is possible to call other Octave functions from within a mex-file using `mexCallMATLAB`. An example of the use of `mexCallMATLAB` can be see in the example below.

```
#include "mex.h"

void
mexFunction (int nlhs, mxArray *plhs[],
             int nrhs, const mxArray *prhs[])
{
  char *str;

  mexPrintf ("Starting file myfeval.mex\n");

  mexPrintf ("I have %d inputs and %d outputs\n", nrhs, nlhs);

  if (nrhs < 1 || ! mxIsChar (prhs[0]))
    mexErrMsgTxt ("ARG1 must be a function name");

  str = mxArrayToString (prhs[0]);

  mexPrintf ("I'm going to call the function %s\n", str);

  if (nlhs == 0)
    nlhs = 1;  // Octave's automatic 'ans' variable

  /* Cast prhs just to get rid of 'const' qualifier and stop compile warning */
  mexCallMATLAB (nlhs, plhs, nrhs-1, (mxArray**)prhs+1, str);

  mxFree (str);
}
```

If this code is in the file `myfeval.c`, and is compiled to `myfeval.mex`, then an example of its use is

```
a = myfeval ("sin", 1)
⇒ Starting file myfeval.mex
   I have 2 inputs and 1 outputs
   I'm going to call the interpreter function sin
   a =  0.84147
```

Note that it is not possible to use function handles within a mex-file.

## A.3 Standalone Programs

The libraries Octave uses itself can be utilized in standalone applications. These applications then have access, for example, to the array and matrix classes, as well as to all of the Octave algorithms. The following C++ program, uses class Matrix from `liboctave.a` or `liboctave.so`.

```
#include <iostream>
```

```
#include <octave/oct.h>

int
main ()
{
  std::cout << "Hello Octave world!\n";

  int n = 2;
  Matrix a_matrix = Matrix (n, n);

  for (octave_idx_type i = 0; i < n; i++)
    for (octave_idx_type j = 0; j < n; j++)
      a_matrix(i,j) = (i + 1) * 10 + (j + 1);

  std::cout << a_matrix;

  return 0;
}
```

mkoctfile can be used to build a standalone application with a command like

```
$ mkoctfile --link-stand-alone standalone.cc -o standalone
$ ./standalone
Hello Octave world!
  11 12
  21 22
$
```

Note that the application `standalone` will be dynamically linked against the Octave libraries and any Octave support libraries. The above allows the Octave math libraries to be used by an application. It does not, however, allow the script files, oct-files, or built-in functions of Octave to be used by the application. To do that, the Octave interpreter needs to be initialized first. An example of how to do this can then be seen in the code

```
#include <iostream>
#include <octave/oct.h>
#include <octave/octave.h>
#include <octave/parse.h>
#include <octave/interpreter.h>

int
main ()
{
  // Create interpreter.

  octave::interpreter interpreter;

  try
    {
      // Inhibit reading history file by calling
```

```
//
//   interpreter.initialize_history (false);

// Set custom load path here if you wish by calling
//
//   interpreter.initialize_load_path (false);

// Perform final initialization of interpreter, including
// executing commands from startup files by calling
//
//   interpreter.initialize ();
//
//   if (! interpreter.is_initialized ())
//     {
//       std::cerr << "Octave interpreter initialization failed!"
//                 << std::endl;
//       exit (1);
//     }
//
// You may skip this step if you don't need to do anything
// between reading the startup files and telling the interpreter
// that you are ready to execute commands.

// Tell the interpreter that we're ready to execute commands:

int status = interpreter.execute ();

if (status != 0)
  {
    std::cerr << "creating embedded Octave interpreter failed!"
              << std::endl;
    return status;
  }

octave_idx_type n = 2;
octave_value_list in;

for (octave_idx_type i = 0; i < n; i++)
  in(i) = octave_value (5 * (i + 2));

octave_value_list out = octave::feval ("gcd", in, 1);

if (out.length () > 0)
  std::cout << "GCD of ["
            << in(0).int_value ()
            << ", "
            << in(1).int_value ()
```

```
                    << "] is " << out(0).int_value ()
                    << std::endl;
        else
          std::cout << "invalid\n";
      }
    catch (const octave::exit_exception& ex)
      {
        std::cerr << "Octave interpreter exited with status = "
                  << ex.exit_status () << std::endl;
      }
    catch (const octave::execution_exception&)
      {
        std::cerr << "error encountered in Octave evaluator!" << std::endl;
      }

    return 0;
  }
```

which, as before, is compiled and run as a standalone application with

```
$ mkoctfile --link-stand-alone embedded.cc -o embedded
$ ./embedded
GCD of [10, 15] is 5
$
```

It is worth re-iterating that, if only built-in functions are to be called from a C++ standalone program then it does not need to initialize the interpreter. The general rule is that for a built-in function named `function_name` in the interpreter, there will be a C++ function named `Ffunction_name` (note the prepended capital `F`) accessible in the C++ API. The declarations for all built-in functions are collected in the header file `builtin-defun-decls.h`. This feature should be used with care as the list of built-in functions can change. No guarantees can be made that a function that is currently a built-in won't be implemented as a `.m` file or as a dynamically linked function in the future. An example of how to call built-in functions from C++ can be seen in the code

```
#include <iostream>
#include <octave/oct.h>
#include <octave/builtin-defun-decls.h>

int
main ()
{
  int n = 2;
  Matrix a_matrix = Matrix (n, n);

  for (octave_idx_type i = 0; i < n; i++)
    for (octave_idx_type j = 0; j < n; j++)
      a_matrix(i,j) = (i + 1) * 10 + (j + 1);

  std::cout << "This is a matrix:" << std::endl
```

```
            << a_matrix            << std::endl;

  octave_value_list in;
  in(0) = a_matrix;

  octave_value_list out = octave::Fnorm (in, 1);
  double norm_of_the_matrix = out(0).double_value ();

  std::cout << "This is the norm of the matrix:" << std::endl
            << norm_of_the_matrix                << std::endl;

  return 0;
}
```

which is compiled and run as a standalone application with

```
$ mkoctfile --link-stand-alone standalonebuiltin.cc -o standalonebuiltin
$ ./standalonebuiltin
This is a matrix:
 11 12
 21 22

This is the norm of the matrix:
34.4952
$
```

## A.4 Java Interface

The Java Interface is designed for calling Java functions from within Octave. If you want to do the reverse, and call Octave from within Java, try a library like `joPas` (http://jopas.sourceforge.net).

### A.4.1 Making Java Classes Available

Java finds classes by searching a *classpath* which is a list of Java archive files and/or directories containing class files. In Octave the *classpath* is composed of two parts:

- the *static classpath* is initialized once at startup of the JVM, and
- the *dynamic classpath* which can be modified at runtime.

Octave searches the *static classpath* first, and then the *dynamic classpath*. Classes appearing in the *static classpath*, as well as in the *dynamic classpath*, will therefore be found in the *static classpath* and loaded from this location. Classes which will be used frequently, or must be available to all users, should be added to the *static classpath*. The *static classpath* is populated once from the contents of a plain text file named `javaclasspath.txt` (or `classpath.txt` historically) when the Java Virtual Machine starts. This file contains one line for each individual classpath to be added to the *static classpath*. These lines can identify directories containing class files, or Java archives with complete class file hierarchies. Comment lines starting with a '`#`' or a '`%`' character are ignored.

The search rules for the file `javaclasspath.txt` (or `classpath.txt`) are:

- First, Octave tries to locate it in the current directory (where Octave was started from). If such a file is found, it is read and defines the initial *static classpath*. Thus, it is possible to define a static classpath on a 'per Octave invocation' basis.
- Next, Octave searches in the user's home directory. If a file `javaclasspath.txt` exists here, its contents are appended to the static classpath (if any). Thus, it is possible to build an initial static classpath on a 'per user' basis.
- Finally, Octave looks for a `javaclasspath.txt` in the m-file directory where Octave Java functions live. This is where the function `javaclasspath.m` resides, usually something like `OCTAVE_HOME/share/octave/OCTAVE_VERSION/m/java/`. You can find this directory by executing the command

  ```
  which javaclasspath
  ```

  If this file exists here, its contents are also appended to the *static classpath*. Note that the archives and class directories defined in this last step will affect all users.

Classes which are used only by a specific script should be placed in the *dynamic classpath*. This portion of the classpath can be modified at runtime using the `javaaddpath` and `javarmpath` functions.

Example:

```
octave> base_path = "C:/Octave/java_files";

octave> # add two JAR archives to the dynamic classpath
octave> javaaddpath ([base_path, "/someclasses.jar"]);
octave> javaaddpath ([base_path, "/moreclasses.jar"]);

octave> # check the dynamic classpath
octave> p = javaclasspath;
octave> disp (p{1});
C:/Octave/java_files/someclasses.jar
octave> disp (p{2});
C:/Octave/java_files/moreclasses.jar

octave> # remove the first element from the classpath
octave> javarmpath ([base_path, "/someclasses.jar"]);
octave> p = javaclasspath;
octave> disp (p{1});
C:/Octave/java_files/moreclasses.jar

octave> # provoke an error
octave> disp (p{2});
error: A(I): Index exceeds matrix dimension.
```

Another way to add files to the *dynamic classpath* exclusively for your user account is to use the file `.octaverc` which is stored in your home directory. All Octave commands in this file are executed each time you start a new instance of Octave. The following example adds the directory `octave` to Octave's search path and the archive `myclasses.jar` in this directory to the Java search path.

```
# contents of .octaverc:
addpath ("~/octave");
javaaddpath ("~/octave/myclasses.jar");
```

### A.4.2 How to use Java from within Octave

The function [javaObject], page 1143, creates Java objects. In fact it invokes the public constructor of the class with the given name and with the given parameters.

The following example shows how to invoke the constructors `BigDecimal(double)` and `BigDecimal(String)` of the builtin Java class `java.math.BigDecimal`.

```
javaObject ("java.math.BigDecimal",  1.001 );
javaObject ("java.math.BigDecimal", "1.001");
```

Note that parameters of the Octave type `double` are implicitly converted into the Java type `double` and the Octave type (array of) `char` is converted into the java type `String`. A Java object created by [javaObject], page 1143, is never automatically converted into an Octave type but remains a Java object. It can be assigned to an Octave variable.

```
a = 1.001;
b = javaObject ("java.math.BigDecimal", a);
```

Using [isjava], page 1143, it is possible to check whether a variable is a Java object and its class can be determined as well. In addition to the previous example:

```
isjava (a)
⇒ ans = 0
class (a)
⇒ ans = double
isjava (b)
⇒ ans = 1
class (b)
⇒ ans = java.math.BigDecimal
```

The example above can be carried out using only Java objects:

```
a = javaObject ("java.lang.Double", 1.001);
b = javaObject ("java.math.BigDecimal", a);

isjava (a)
⇒ ans = 1
class (a)
⇒ ans = java.lang.Double
isjava (b)
⇒ ans = 1
class (b)
⇒ ans = java.math.BigDecimal
```

One can see, that even a `java.lang.Double` is not converted to an Octave `double`, when created by [javaObject], page 1143. But ambiguities might arise, if the Java classes `java.lang.Double` or `double` are parameters of a method (or a constructor). In this case they can be converted into one another, depending on the context.

Via [javaObject], page 1143, one may create all kinds of Java objects but arrays. The latter are created through [javaArray], page 1143.

It is possible to invoke public member methods on Java objects in Java syntax:

```
a.toString
⇒ ans = 1.001
b.toString
⇒ ans = 1.000999999999999889865...
```

The second result may be surprising, but simply comes from the fact, that `1.001` cannot exactly be represented as `double`, due to rounding. Note that unlike in Java, in Octave methods without arguments can be invoked with and without parentheses `()`.

Currently it is not possible to invoke static methods with a Java like syntax from within Octave. Instead, one has to use the function [javaMethod], page 1145, as in the following example:

```
java.math.BigDecimal.valueOf(1.001);                    # does not work
javaMethod ("valueOf", "java.math.BigDecimal", 1.001);  # workaround
```

As mentioned before, method and constructor parameters are converted automatically between Octave and Java types, if appropriate. For functions this is also true with return values, whereas for constructors this is not.

It is also possible to access public fields of Java objects from within Octave using Java syntax, with the limitation of static fields:

```
java.math.BigDecimal.ONE;                  # does not work
java_get ("java.math.BigDecimal", "ONE");  # workaround
```

Accordingly, with [java_set], page 1144, the value of a field can be set. Note that only public Java fields are accessible from within Octave.

The following example indicates that in Octave empty brackets `[]` represent Java's `null` value and how Java exceptions are represented.

```
javaObject ("java.math.BigDecimal", []);
⇒ error: [java] java.lang.NullPointerException
```

It is not recommended to represent Java's `null` value by empty brackets `[]`, because `null` has no type whereas `[]` has type `double`.

In Octave it is possible to provide limited Java reflection by listing the public fields and methods of a Java object, both static or not.

```
fieldnames (<Java object>)
methods (<Java object>)
```

Finally, an examples is shown how to access the stack trace from within Octave, where the function [debug_java], page 1148, is used to set and to get the current debug state. In debug mode, the Java error and the stack trace are displayed.

```
debug_java (true)  # use "false" to omit display of stack trace
debug_java ()
⇒ ans = 1
javaObject ("java.math.BigDecimal", "1") ...
  .divide (javaObject ("java.math.BigDecimal", "0"))
```

### A.4.3 Set up the JVM

In order to execute Java code Octave creates a Java Virtual Machine (JVM). By default the version of the JVM is used that was detected during configuration on Unix-like systems or that is pointed to from the registry keys at `HKEY_LOCAL_MACHINE\SOFTWARE\JavaSoft\JRE` or `HKEY_LOCAL_MACHINE\SOFTWARE\JavaSoft\Java Runtime Environment` on Windows. The default path to the JVM can be overridden by setting the environment variable `JAVA_HOME` to the path where the JVM is installed. On Windows that might be, for example, `C:\Program Files\Java\jre-10.0.2`. Make sure that you select a directory that contains the JVM with a bit-ness that matches Octave's.

The JVM is only loaded once per Octave session. Thus, to change the used version of the JVM, you might have to re-start Octave. To check which version of the JVM is currently being used, run `version -java`.

The JVM allocates a fixed amount of initial memory and may expand this pool up to a fixed maximum memory limit. The default values depend on the Java version (see [javamem], page 1147). The memory pool is shared by all Java objects running in the JVM. This strict memory limit is intended mainly to avoid runaway applications inside web browsers or in enterprise servers which can consume all memory and crash the system. When the maximum memory limit is hit, Java code will throw exceptions so that applications will fail or behave unexpectedly.

You can specify options for the creation of the JVM inside a file named `java.opts`. This is a text file where enter you enter lines containing `-X` and `-D` options that are then passed to the JVM during initialization.

The directory where the Java options file is located is specified by the environment variable `OCTAVE_JAVA_DIR`. If unset the directory where `javaclasspath.m` resides is used instead (typically `OCTAVE_HOME/share/octave/OCTAVE_VERSION/m/java/`). You can find this directory by executing

```
which javaclasspath
```

The `-X` options allow you to increase the maximum amount of memory available to the JVM. The following example allows up to 256 Megabytes to be used by adding the following line to the `java.opts` file:

```
-Xmx256m
```

The maximum possible amount of memory depends on your system. On a Windows system with 2 Gigabytes main memory you should be able to set this maximum to about 1 Gigabyte.

If your application requires a large amount of memory from the beginning, you can also specify the initial amount of memory allocated to the JVM. Adding the following line to the `java.opts` file starts the JVM with 64 Megabytes of initial memory:

```
-Xms64m
```

For more details on the available `-X` options of your Java Virtual Machine issue the command '`java -X`' at the operating system command prompt and consult the Java documentation.

The `-D` options can be used to define system properties which can then be used by Java classes inside Octave. System properties can be retrieved by using the `getProperty()`

methods of the `java.lang.System` class. The following example line defines the property *MyProperty* and assigns it the string `12.34`.

```
-DMyProperty=12.34
```

The value of this property can then be retrieved as a string by a Java object or in Octave:

```
octave> javaMethod ("getProperty", "java.lang.System", "MyProperty");
ans = 12.34
```

### A.4.4 Java Interface Functions

The following functions are the core of the Java Interface. They provide a way to create a Java object, get and set its data fields, and call Java methods which return results to Octave.

*jobj* = javaObject (*classname*)
*jobj* = javaObject (*classname*, *arg1*, ...)

Create a Java object of class *classsname*, by calling the class constructor with the arguments *arg1*, ...

The first example below creates an uninitialized object, while the second example supplies an initial argument to the constructor.

```
x = javaObject ("java.lang.StringBuffer")
x = javaObject ("java.lang.StringBuffer", "Initial string")
```

**See also:** [javaMethod], page 1145, [javaArray], page 1143.

*jary* = javaArray (*classname*, *sz*)
*jary* = javaArray (*classname*, *m*, *n*, ...)

Create a Java array of size *sz* with elements of class *classname*.

*classname* may be a Java object representing a class or a string containing the fully qualified class name. The size of the object may also be specified with individual integer arguments *m*, *n*, etc.

The generated array is uninitialized. All elements are set to null if *classname* is a reference type, or to a default value (usually 0) if *classname* is a primitive type.

Sample code:

```
jary = javaArray ("java.lang.String", 2, 2);
jary(1,1) = "Hello";
```

**See also:** [javaObject], page 1143.

There are many different variable types in Octave, but only ones created through `javaObject` can use Java functions. Before using Java with an unknown object the type can be checked with `isjava`.

*tf* = isjava (*x*)

Return true if *x* is a Java object.

**See also:** [class], page 41, [typeinfo], page 41, [isa], page 41, [javaObject], page 1143.

Once an object has been created it is natural to find out what fields the object has, and to read (get) and write (set) them.

In Octave the `fieldnames` function for structures has been overloaded to return the fields of a Java object. For example:

```
dobj = javaObject ("java.lang.Double", pi);
fieldnames (dobj)
⇒
{
  [1,1] = public static final double java.lang.Double.POSITIVE_INFINITY
  [1,2] = public static final double java.lang.Double.NEGATIVE_INFINITY
  [1,3] = public static final double java.lang.Double.NaN
  [1,4] = public static final double java.lang.Double.MAX_VALUE
  [1,5] = public static final double java.lang.Double.MIN_NORMAL
  [1,6] = public static final double java.lang.Double.MIN_VALUE
  [1,7] = public static final int java.lang.Double.MAX_EXPONENT
  [1,8] = public static final int java.lang.Double.MIN_EXPONENT
  [1,9] = public static final int java.lang.Double.SIZE
  [1,10] = public static final java.lang.Class java.lang.Double.TYPE
}
```

The analogy of objects with structures is carried over into reading and writing object fields. To read a field the object is indexed with the '`.`' operator from structures. This is the preferred method for reading fields, but Octave also provides a function interface to read fields with `java_get`. An example of both styles is shown below.

```
dobj = javaObject ("java.lang.Double", pi);
dobj.MAX_VALUE
⇒  1.7977e+308
java_get ("java.lang.Float", "MAX_VALUE")
⇒  3.4028e+38
```

*val* = java_get (*obj*, *name*)

Get the value of the field *name* of the Java object *obj*.

For static fields, *obj* can be a string representing the fully qualified name of the corresponding class.

When *obj* is a regular Java object, structure-like indexing can be used as a shortcut syntax. For instance, the following two statements are equivalent

```
java_get (x, "field1")
x.field1
```

**See also:** [java_set], page 1144, [javaMethod], page 1145, [javaObject], page 1143.

*obj* = java_set (*obj*, *name*, *val*)

Set the value of the field *name* of the Java object *obj* to *val*.

For static fields, *obj* can be a string representing the fully qualified named of the corresponding Java class.

When *obj* is a regular Java object, structure-like indexing can be used as a shortcut syntax. For instance, the following two statements are equivalent

```
java_set (x, "field1", val)
x.field1 = val
```

**See also:** [java_get], page 1144, [javaMethod], page 1145, [javaObject], page 1143.

To see what functions can be called with an object use `methods`. For example, using the previously created *dobj*:

```
methods (dobj)
⇒
Methods for class java.lang.Double:
boolean equals(java.lang.Object)
java.lang.String toString(double)
java.lang.String toString()
...
```

To call a method of an object the same structure indexing operator '`.`' is used. Octave also provides a functional interface to calling the methods of an object through `javaMethod`. An example showing both styles is shown below.

```
dobj = javaObject ("java.lang.Double", pi);
dobj.equals (3)
⇒  0
javaMethod ("equals", dobj, pi)
⇒  1
```

*ret* = javaMethod (*methodname*, *obj*)
*ret* = javaMethod (*methodname*, *obj*, *arg1*, ...)

Invoke the method *methodname* on the Java object *obj* with the arguments *arg1*, ...

For static methods, *obj* can be a string representing the fully qualified name of the corresponding class.

When *obj* is a regular Java object, structure-like indexing can be used as a shortcut syntax. For instance, the two following statements are equivalent

```
ret = javaMethod ("method1", x, 1.0, "a string")
ret = x.method1 (1.0, "a string")
```

`javaMethod` returns the result of the method invocation.

**See also:** [methods], page 974, [javaObject], page 1143.

The following three functions are used to display and modify the class path used by the Java Virtual Machine. This is entirely separate from Octave's `PATH` variable and is used by the JVM to find the correct code to execute.

javaclasspath ()
*dpath* = javaclasspath ()
[*dpath*, *spath*] = javaclasspath ()
*clspath* = javaclasspath (*what*)

Return the class path of the Java virtual machine in the form of a cell array of strings.

If called with no inputs:

- If no output is requested, the dynamic and static classpaths are printed to the standard output.

- If one output value *dpath* is requested, the result is the dynamic classpath.
- If two output values*dpath* and *spath* are requested, the first variable will contain the dynamic classpath and the second will contain the static classpath.

If called with a single input parameter *what*:

`"-dynamic"`
: Return the dynamic classpath.

`"-static"`
: Return the static classpath.

`"-all"` Return both the static and dynamic classpath in a single cellstr.

**See also:** [javaaddpath], page 1146, [javarmpath], page 1146.

`javaaddpath (`*`clspath`*`)`
`javaaddpath (`*`clspath1`*`, ...)`
`javaaddpath ({`*`clspath1`*`, ...})`
`javaaddpath (..., "-`*`end`*`")`

Add *clspath* to the beginning of the dynamic class path of the Java virtual machine.

*clspath* may either be a directory where `.class` files are found, or a `.jar` file containing Java classes. Multiple paths may be added at once by specifying additional arguments, or by using a cell array of strings.

If the final argument is `"-end"`, append the new element to the end of the current classpath.

**See also:** [javarmpath], page 1146, [javaclasspath], page 1145.

`javarmpath (`*`clspath`*`)`
`javarmpath (`*`clspath1`*`, ...)`
`javarmpath ({`*`clspath1`*`, ...})`

Remove *clspath* from the dynamic class path of the Java virtual machine.

*clspath* may either be a directory where `.class` files are found, or a `.jar` file containing Java classes. Multiple paths may be removed at once by specifying additional arguments, or by using a cell array of strings.

**See also:** [javaaddpath], page 1146, [javaclasspath], page 1145.

The following functions provide information and control over the interface between Octave and the Java Virtual Machine.

*`msg`* `= javachk (`*`feature`*`)`
*`msg`* `= javachk (`*`feature`*`, `*`caller`*`)`

Check for the presence of the Java *feature* in the current session. Return an error structure if *feature* is not available, not enabled, or not recognized.

Possible recognized features are:

`"awt"` Abstract Window Toolkit for GUIs.

`"desktop"`
: Interactive desktop is running.

`"jvm"` Java Virtual Machine.

`"swing"` Swing components for lightweight GUIs.

If *feature* is not supported, a scalar struct with field `"message"` and `"identifier"` is returned. The field `"message"` contains an error message mentioning *feature* as well as the optional user-specified *caller*. This structure is suitable for passing in to the `error` function.

If *feature* is supported and available, an empty struct array is returned with fields `"message"` and `"identifier"`.

`javachk` determines if specific Java features are available in an Octave session. This function is provided for scripts which may alter their behavior based on the availability of Java or specific Java runtime features.

Compatibility Note: The feature `"desktop"` is never available since Octave has no Java-based desktop.

**See also:** [usejava], page 1147, [error], page 255.

*tf* = `usejava` (*feature*)

Return true if the Java element *feature* is available.

Possible features are:

`"awt"` Abstract Window Toolkit for GUIs.

`"desktop"`
Interactive desktop is running.

`"jvm"` Java Virtual Machine.

`"swing"` Swing components for lightweight GUIs.

`usejava` determines if specific Java features are available in an Octave session. This function is provided for scripts which may alter their behavior based on the availability of Java. The feature `"desktop"` always returns `false` as Octave has no Java-based desktop. Other features may be available if Octave was compiled with the Java Interface and Java is installed.

**See also:** [javachk], page 1146.

`javamem` ()
*jmem* = `javamem` ()

Show the current memory usage of the Java virtual machine (JVM) and run the garbage collector.

When no return argument is given the info is printed to the screen. Otherwise, the output cell array *jmem* contains Maximum, Total, and Free memory (in bytes).

All Java-based routines are run in the JVM's shared memory pool, a dedicated and separate part of memory claimed by the JVM from your computer's total memory (which comprises physical RAM and virtual memory / swap space on hard disk).

The maximum allowable memory usage can be configured using the file `java.opts`. The directory where this file resides is determined by the environment variable

`OCTAVE_JAVA_DIR`. If unset, the directory where `javaaddpath.m` resides is used instead (typically `OCTAVE_HOME/share/octave/OCTAVE_VERSION/m/java/`).

`java.opts` is a plain text file with one option per line. The default initial memory size and default maximum memory size (which are both system dependent) can be overridden like so:

-Xms64m

-Xmx512m

(in megabytes in this example). You can adapt these values to your own requirements if your system has limited available physical memory or if you get Java memory errors.

`"Total memory"` is what the operating system has currently assigned to the JVM and depends on actual and active memory usage. `"Free memory"` is self-explanatory. During operation of Java-based Octave functions the amount of Total and Free memory will vary, due to Java's own cleaning up and your operating system's memory management.

*val* = java_matrix_autoconversion ()
*old_val* = java_matrix_autoconversion (*new_val*)
*old_val* = java_matrix_autoconversion (*new_val*, "*local*")

Query or set the internal variable that controls whether Java arrays are automatically converted to Octave matrices.

The default value is false.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [java_unsigned_autoconversion], page 1148, [debug_java], page 1148.

*val* = java_unsigned_autoconversion ()
*old_val* = java_unsigned_autoconversion (*new_val*)
*old_val* = java_unsigned_autoconversion (*new_val*, "*local*")

Query or set the internal variable that controls how integer classes are converted when `java_matrix_autoconversion` is enabled.

When enabled, Java arrays of class Byte or Integer are converted to matrices of class uint8 or uint32 respectively. The default value is true.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [java_matrix_autoconversion], page 1148, [debug_java], page 1148.

*val* = debug_java ()
*old_val* = debug_java (*new_val*)
*old_val* = debug_java (*new_val*, "*local*")

Query or set the internal variable that determines whether extra debugging information regarding the initialization of the JVM and any Java exceptions is printed.

When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

**See also:** [java_matrix_autoconversion], page 1148, [java_unsigned_autoconversion], page 1148.

# Appendix B Test and Demo Functions

Octave includes a number of functions to allow the integration of testing and demonstration code in the source code of the functions themselves.

## B.1 Test Functions

`test` *name*
`test` *name quiet*|*normal*|*verbose*
`test` ("*name*", "*quiet*|*normal*|*verbose*", *fid*)
`test` ("*name*", "*quiet*|*normal*|*verbose*", *fname*)
*success* = `test` (...)
[*n*, *nmax*, *nxfail*, *nbug*, *nskip*, *nrtskip*, *nregression*] = `test` (...)
[*code*, *idx*] = `test` ("*name*", "*grabdemo*")
`test` ([], "*explain*", *fid*)
`test` ([], "*explain*", *fname*)

Perform built-in self-tests from the first file in the loadpath matching *name*.

`test` can be called in either command or functional form. The exact operation of test is determined by a combination of mode (interactive or batch), reporting level (`"quiet"`, `"normal"`, `"verbose"`), and whether a logfile or summary output variable is used.

The default mode when `test` is called from the command line is interactive. In this mode, tests will be run until the first error is encountered, or all tests complete successfully. In batch mode, all tests are run regardless of any failures, and the results are collected for reporting. Tests which require user interaction, i.e., demo blocks, are never run in batch mode.

Batch mode is enabled by either 1) specifying a logfile using the third argument *fname* or *fid*, or 2) requesting an output argument such as *success*, *n*, etc.

The optional second argument determines the amount of output to generate and which types of tests to run. The default value is `"normal"`. Requesting an output argument will suppress printing the final summary message and any intermediate warnings, unless verbose reporting is enabled.

`"quiet"` Print a summary message when all tests pass, or print an error with the results of the first bad test when a failure occurs. Don't run tests which require user interaction.

`"normal"` Display warning messages about skipped tests or failing xtests during test execution. Print a summary message when all tests pass, or print an error with the results of the first bad test when a failure occurs. Don't run tests which require user interaction.

`"verbose"`
Display tests before execution. Print all warning messages. In interactive mode, run all tests including those which require user interaction.

The optional third input argument specifies a logfile where results of the tests should be written. The logfile may be a character string (*fname*) or an open file descriptor

ID (*fid*). To enable batch processing, but still print the results to the screen, use `stdout` for *fid*.

When called with just a single output argument *success*, `test` returns true if all of the tests were successful. If called with more than one output argument then the number of successful tests (*n*), the total number of tests in the file (*nmax*), the number of xtest failures (*nxfail*), the number of tests failed due known bugs (*nbug*), the number of tests skipped due to missing features (*nskip*), the number of tests skipped due to run-time conditions (*nrtskip*), and the number of regressions (*nregression*) are returned.

Example

```
test sind
⇒
PASSES 5 out of 5 tests

[n, nmax] = test ("sind")
⇒
n =  5
nmax =  5
```

Additional Calling Syntaxes

If the second argument is the string `"grabdemo"`, the contents of any built-in demo blocks are extracted but not executed. The text for all code blocks is concatenated and returned as *code* with *idx* being a vector of positions of the ends of each demo block. For an easier way to extract demo blocks from files, See [example], page 1161.

If the second argument is `"explain"` then *name* is ignored and an explanation of the line markers used in `test` output reports is written to the file specified by *fname* or *fid*.

**See also:** [assert], page 1157, [fail], page 1159, [demo], page 1160, [example], page 1161, [error], page 255.

`test` scans the named script file looking for lines which start with the identifier '`%!`'. The prefix is stripped off and the rest of the line is processed through the Octave interpreter. If the code generates an error, then the test is said to fail.

Since `eval()` will stop at the first error it encounters, you must divide your tests up into blocks, with anything in a separate block evaluated separately. Blocks are introduced by valid keywords like `test`, `function`, or `assert` immediately following '`%!`'. A block is defined by indentation as in Python. Lines beginning with '`%!<whitespace>`' are part of the preceding block.

For example:

```
%!test error ("this test fails!")
%!test "test doesn't fail.  it doesn't generate an error"
```

When a test fails, you will see something like:

```
  ***** test error ("this test fails!")
!!!!! test failed
this test fails!
```

Generally, to test if something works, you want to assert that it produces a correct value. A real test might look something like

```
%!test
%! a = [1, 2, 3; 4, 5, 6]; B = [1; 2];
%! expect = [ a ; 2*a ];
%! get = kron (b, a);
%! if (any (size (expect) != size (get)))
%!   error ("wrong size: expected %d,%d but got %d,%d",
%!          size (expect), size (get));
%! elseif (any (any (expect != get)))
%!   error ("didn't get what was expected.");
%! endif
```

To make the process easier, use the `assert` function. For example, with `assert` the previous test is reduced to:

```
%!test
%! a = [1, 2, 3; 4, 5, 6]; b = [1; 2];
%! assert (kron (b, a), [ a; 2*a ]);
```

`assert` can accept a tolerance so that you can compare results absolutely or relatively. For example, the following all succeed:

```
%!test assert (1+eps, 1, 2*eps)           # absolute error
%!test assert (100+100*eps, 100, -2*eps)  # relative error
```

You can also do the comparison yourself, but still have assert generate the error:

```
%!test assert (isempty ([]))
%!test assert ([1, 2; 3, 4] > 0)
```

Because `assert` is so frequently used alone in a test block, there is a shorthand form:

```
%!assert (...)
```

which is equivalent to:

```
%!test assert (...)
```

Occasionally a block of tests will depend on having optional functionality in Octave. Before testing such blocks the availability of the required functionality must be checked. A `%!testif HAVE_XXX` block will only be run if Octave was compiled with functionality '`HAVE_XXX`'. For example, the sparse single value decomposition, `svds()`, depends on having the ARPACK library. All of the tests for `svds` begin with

```
%!testif HAVE_ARPACK
```

Review `config.h` or `__octave_config_info__ ("build_features")` to see some of the possible values to check.

Sometimes during development there is a test that should work but is known to fail. You still want to leave the test in because when the final code is ready the test should pass, but you may not be able to fix it immediately. To avoid unnecessary bug reports for these known failures, mark the block with `xtest` rather than `test`:

```
%!xtest assert (1==0)
%!xtest fail ("success=1", "error")
```

In this case, the test will run and any failure will be reported. However, testing is not aborted and subsequent test blocks will be processed normally. Another use of `xtest` is for statistical tests which should pass most of the time but are known to fail occasionally.

Each block is evaluated in its own function environment, which means that variables defined in one block are not automatically shared with other blocks. If you do want to share variables, then you must declare them as `shared` before you use them. For example, the following declares the variable *a*, gives it an initial value (default is empty), and then uses it in several subsequent tests.

```
%!shared a
%! a = [1, 2, 3; 4, 5, 6];
%!assert (kron ([1; 2], a), [ a; 2*a ])
%!assert (kron ([1, 2], a), [ a, 2*a ])
%!assert (kron ([1,2; 3,4], a), [ a,2*a; 3*a,4*a ])
```

You can share several variables at the same time:

```
%!shared a, b
```

Modifications to shared variables persist from one test to the next **only** if the test succeeds. Thus, if one test modifies a shared variable, later tests cannot know which value of the shared variable to expect because the pass/fail status of earlier tests is unknown. For this reason, it is not recommended to modify shared variables in tests.

You can also share test functions:

```
%!function a = fn (b)
%!  a = 2*b;
%!endfunction
%!assert (fn(2), 4)
```

Note that all previous variables and values are lost when a new shared block is declared.

Remember that `%!function` begins a new block and that `%!endfunction` ends this block. Be aware that until a new block is started, lines starting with '`%!<space>`' will be discarded as comments. The following is nearly identical to the example above, but does nothing.

```
%!function a = fn (b)
%!  a = 2*b;
%!endfunction
%! assert (fn(2), 4)
```

Because there is a space after '`%!`' the `assert` statement does not begin a new block and this line is treated as a comment.

Error and warning blocks are like test blocks, but they only succeed if the code generates an error. You can check the text of the error is correct using an optional regular expression `<pattern>`. For example:

```
%!error <passes!> error ("this test passes!")
```

If the code doesn't generate an error, the test fails. For example:

```
%!error "this is an error because it succeeds."
```

produces

```
  ***** error "this is an error because it succeeds."
!!!!! test failed: no error
```

It is important to automate the tests as much as possible, however some tests require user interaction. These can be isolated into demo blocks, which if you are in batch mode, are only run when called with `demo` or the `verbose` option to `test`. The code is displayed before it is executed. For example,

```
%!demo
%! t = [0:0.01:2*pi]; x = sin (t);
%! plot (t, x);
%! # you should now see a sine wave in your figure window
```

produces

```
funcname example 1:
 t = [0:0.01:2*pi]; x = sin (t);
 plot (t, x);
 # you should now see a sine wave in your figure window

Press <enter> to continue:
```

Note that demo blocks cannot use any shared variables. This is so that they can be executed by themselves, ignoring all other tests.

If you want to temporarily disable a test block, put `#` in place of the block type. This creates a comment block which is echoed in the log file but not executed. For example:

```
%!#demo
%! t = [0:0.01:2*pi]; x = sin (t);
%! plot (t, x);
%! # you should now see a sine wave in your figure window
```

The following trivial code snippet provides examples for the use of fail, assert, error, and xtest:

```
function output = must_be_zero (input)
  if (input != 0)
    error ("Nonzero input!")
  endif
  output = input;
endfunction

%!fail ("must_be_zero (1)")
%!assert (must_be_zero (0), 0)
%!error <Nonzero> must_be_zero (1)
%!xtest error ("This code generates an error")
```

When putting this in a file `must_be_zero.m`, and running the test, we see

```
test must_be_zero verbose

⇒
>>>>> /path/to/must_be_zero.m
***** fail ("must_be_zero (1)")
***** assert (must_be_zero (0), 0)
***** error <Nonzero> must_be_zero (1)
***** xtest error ("This code generates an error")
!!!!! known failure
This code generates an error
PASSES 3 out of 4 tests (1 expected failure)
```

## Block type summary

`%!test`
`%!test <MESSAGE>`

Check that entire block is correct. If `<MESSAGE>` is present, the test block is interpreted as for `xtest`.

`%!testif HAVE_XXX`
`%!testif HAVE_XXX, HAVE_YYY, ...`
`%!testif ...; RUNTIME_COND`
`%!testif ... <MESSAGE>`

Check block only if Octave was compiled with feature `HAVE_XXX`. `RUNTIME_COND` is an optional expression to evaluate to check whether some condition is met when the test is executed. If `RUNTIME_COND` is false, the test is skipped. If `<MESSAGE>` is present, the test block is interpreted as for `xtest`.

`%!xtest`
`%!xtest <MESSAGE>`

Check block, report a test failure but do not abort testing. If `<MESSAGE>` is present, then the text of the message is displayed if the test fails, like this:

```
!!!!! known bug:  MESSAGE
```

If the message is an integer, it is interpreted as a bug ID for the Octave bug tracker and reported as

```
!!!!! known bug: https://octave.org/testfailure/?BUG-ID
```

in which BUG-ID is the integer bug number. The intent is to allow clearer documentation of known problems.

If `MESSAGE` is an integer preceeded by an asterisk (e.g., `*12345`), it is interpreted as the id for a bug report that has been closed. That usually means that the issue probed in this test has been resolved. If such tests are failing, they are reported as regressions by the `test` function:

```
!!!!! regression: https://octave.org/testfailure/?BUG-ID
```

`%!error`
`%!error <MESSAGE>`
`%!warning`
`%!warning <MESSAGE>`
: Check for correct error or warning message. If `<MESSAGE>` is supplied it is interpreted as a regular expression pattern that is expected to match the error or warning message.

`%!demo`
: Demo only executes in interactive mode.

`%!#`
: Comment. Ignore everything within the block

`%!shared x,y,z`
: Declare variables for use in multiple tests.

`%!function`
: Define a function for use in multiple tests.

`%!endfunction`
: Close a function definition.

`%!assert (x, y, tol)`
`%!assert <MESSAGE> (x, y, tol)`
`%!fail (CODE, PATTERN)`
`%!fail <MESSAGE> (CODE, PATTERN)`
: Shorthand for `%!test assert (x, y, tol)` or `%!test fail (CODE, PATTERN)`. If `<MESSAGE>` is present, the test block is interpreted as for `xtest`.

When coding tests the Octave convention is that lines that begin with a block type do not have a semicolon at the end. Any code that is within a block, however, is normal Octave code and usually will have a trailing semicolon. For example,

```
## bare block instantiation
%!assert (sin (0), 0)
```

but

```
## test block with normal Octave code
%!test
%! assert (sin (0), 0);
```

You can also create test scripts for built-in functions and your own C++ functions. To do so, put a file with the bare function name (no .m extension) in a directory in the load path and it will be discovered by the `test` function. Alternatively, you can embed tests directly in your C++ code:

```
/*
%!test disp ("this is a test")
*/
```

or

```
#if 0
%!test disp ("this is a test")
#endif
```

However, in this case the raw source code will need to be on the load path and the user will have to remember to type `test ("funcname.cc")`.

`assert (cond)`
`assert (cond, errmsg)`
`assert (cond, errmsg, ...)`
`assert (cond, msg_id, errmsg, ...)`
`assert (observed, expected)`
`assert (observed, expected, tol)`

Produce an error if the specified condition is not met.

`assert` can be called in three different ways.

`assert (cond)`
`assert (cond, errmsg)`
`assert (cond, errmsg, ...)`
`assert (cond, msg_id, errmsg, ...)`

Called with a single argument *cond*, `assert` produces an error if *cond* is false (numeric zero).

Any additional arguments are passed to the `error` function for processing.

`assert (observed, expected)`

Produce an error if observed is not the same as expected.

Note that *observed* and *expected* can be scalars, vectors, matrices, strings, cell arrays, or structures.

`assert (observed, expected, tol)`

Produce an error if observed is not the same as expected but equality comparison for numeric data uses a tolerance *tol*.

If *tol* is positive then it is an absolute tolerance which will produce an error if `abs (observed - expected) > abs (tol)`.

If *tol* is negative then it is a relative tolerance which will produce an error if `abs (observed - expected) > abs (tol * expected)`.

If *expected* is zero *tol* will always be interpreted as an absolute tolerance.

If *tol* is not scalar its dimensions must agree with those of *observed* and *expected* and tests are performed on an element-by-element basis.

**See also:** [fail], page 1159, [test], page 1151, [error], page 255, [isequal], page 175.

`assert_equal (observed, expected)`
`assert_equal (observed, expected, tol)`

Produce an error if observed is not the same as expected.

`assert_equal (observed, expected)` tests for the two input arguments being equal with respect to their class, size, and value. All octave values are supported as input for *observed* and *expected* and compared for equality according to the following rules.

- Numeric inputs that support exceptional values (i.e., Inf, -Inf, NaN, NA) are treated under the additional assumption that they are equal. They are also tested for sparsity and complexity.
- Logical inputs are additionally tested for sparsity.
- Character arrays are compared with the `strcmp` function.

- Structures are additionally checked against having the same unordered field-names.
- Cell arrays are recursively tested on an element by element basis.
- Function handles are tested with the `isequal` function.
- Classdef objects are explicitly tested if and only if they have an overloaded `eq` method and unless `all (eq (`*observed*`, `*expected*`), "all")` is true, an error is produced. The handling of missing values is strictly defined by the overloaded `eq` method.

`assert_equal (`*observed*`, `*expected*`, `*tol*`)` further specifies a tolerance value, which is explicitly used for numerical inputs. Specifying *tol* as 0, is equivalent to calling `assert_equal` with only two input arguments. Any nonzero tolerance is also ignored when *observed* and *expected* are non-numeric. Logical and character arrays are not considered to be numeric.

If *tol* is positive then it is an absolute tolerance which will produce an error if `abs (`*observed* `-` *expected*`) > abs (`*tol*`)`.

If *tol* is negative then it is a relative tolerance which will produce an error if `abs (`*observed* `-` *expected*`) > abs (`*tol* `*` *expected*`)`.

If *expected* is zero *tol* will always be interpreted as an absolute tolerance.

If *tol* is not scalar its dimensions must agree with those of *observed* and *expected* and tests are performed on an element-by-element basis.

**See also:** [fail], page 1159, [test], page 1151, [error], page 255, [isequal], page 175, [eq], page 175, [strcmp], page 88, [assert], page 1157.

*status* `= fail (`*code*`)`
*status* `= fail (`*code*`, `*pattern*`)`
*status* `= fail (`*code*`, "`*warning*`")`
*status* `= fail (`*code*`, "`*warning*`", `*pattern*`)`

Return true if *code* fails with an error message matching *pattern*, otherwise produce an error.

*code* must be in the form of a string that is passed to the Octave interpreter via the `evalin` function, i.e., a (quoted) string constant or a string variable.

Note that if *code* runs successfully, rather than failing, the error printed is:

```
expected error <.> but got none
```

If called with two arguments, the return value will be true only if *code* fails with an error message containing *pattern* (case sensitive). If the code fails with a different error than the one specified in *pattern* then the message produced is:

```
expected <pattern>
but got <text of actual error>
```

The angle brackets are not part of the output.

When called with the `"warning"` option `fail` will produce an error if executing the code produces no warning.

**See also:** [assert], page 1157, [error], page 255.

## B.2 Demonstration Functions

`demo` *name*
`demo` *name n*
`demo` (`"`*name*`"`)
`demo` (`"`*name*`"`, *n*)

Run example code block *n* associated with the function *name*.

If *n* is not specified, all examples are run.

The preferred location for example code blocks is embedded within the script m-file immediately following the code that it exercises. Alternatively, the examples may be stored in a file with the same name but no extension located on Octave's load path. To separate examples from regular script code all lines are prefixed by `%!`. Each example must also be introduced by the keyword `"demo"` flush left to the prefix with no intervening spaces. The remainder of the example can contain arbitrary Octave code. For example:

```
%!demo
%! t = 0:0.01:2*pi;
%! x = sin (t);
%! plot (t, x);
%! title ("one cycle of a sine wave");
%! #-----------------------------------------------
%! # the figure window shows one cycle of a sine wave
```

Note that the code is displayed before it is executed so that a simple comment at the end suffices for labeling what is being shown. For plots, labeling can also be done with `title` or `text`. It is generally **not** necessary to use `disp` or `printf` within the demo.

Demos are run in a stand-alone function environment with no access to external variables. This means that every demo must have separate initialization code. Alternatively, all demos can be combined into a single large demo with the code

```
%! input ("Press <enter> to continue: ", "s");
```

between the sections, but this usage is discouraged. Other techniques to avoid multiple initialization blocks include using multiple plots with a new `figure` command between each plot, or using `subplot` to put multiple plots in the same window.

Finally, because `demo` evaluates within a function context it is not possible to define new functions within the code. Anonymous functions make a good substitute in most instances. If function blocks **must** be used then the code `eval (example ("function", n))` will allow Octave to see them. This has its own problems, however, as `eval` only evaluates one line or statement at a time. In this case the function declaration must be wrapped with `"if 1 <demo stuff> endif"` where `"if"` is on the same line as `"demo"`. For example:

```
%!demo if 1
%!  function y = f(x)
%!    y = x;
%!  endfunction
%!  f(3)
%! endif
```

**See also:** [rundemos], page 1161, [example], page 1161, [test], page 1151.

`example` *name*
`example` *name* *n*
`example ("`*name*`")`
`example ("`*name*`",` *n*`)`
`[`*codestr*`,` *codeidx*`] = example (...)`

Display the code for example *n* associated with the function *name*, but do not run it.

If *n* is not specified, all examples are displayed.

When called with output arguments, the examples are returned in the form of a string *codestr*, with *codeidx* indicating the ending position of the various examples.

For a complete explanation see [demo], page 1160.

**See also:** [demo], page 1160, [test], page 1151.

`oruntests ()`
`oruntests (`*directory*`)`

Execute built-in tests for all m-files in the specified *directory*.

Test blocks in any C++ source files (`*.cc`) will also be executed for use with dynamically linked oct-file functions.

If no directory is specified, operate on all directories in Octave's search path for functions.

**See also:** [rundemos], page 1161, [test], page 1151, [path], page 229.

`rundemos ()`
`rundemos (`*directory*`)`

Execute built-in demos for all m-files in the specified *directory*.

Demo blocks in any C++ source files (`*.cc`) will also be executed for use with dynamically linked oct-file functions.

If no directory is specified, operate on all directories in Octave's search path for functions.

**See also:** [demo], page 1160, [oruntests], page 1161, [path], page 229.

`speed (`*f*`,` *init*`,` *max_n*`,` *f2*`,` *tol*`)`
`[`*order*`,` *n*`,` *T_f*`,` *T_f2*`] = speed (...)`

Determine the execution time of an expression (*f*) for various input values (*n*).

The *n* are log-spaced from 1 to *max_n*. For each *n*, an initialization expression (*init*) is computed to create any data needed for the test. If a second expression (*f2*) is given then the execution times of the two expressions are compared. When called without output arguments the results are printed to stdout and displayed graphically.

*f* The code expression to evaluate.

*max_n* The maximum test length to run. The default value is 100. Alternatively, use `[min_n, max_n]` or specify the *n* exactly with `[n1, n2, ..., nk]`.

*init* Initialization expression for function argument values. Use *k* for the test number and *n* for the size of the test. This should compute values for all variables used by *f*. Note that *init* will be evaluated first for $k = 0$, so things which are constant throughout the test series can be computed once. The default value is `x = randn (n, 1)`.

*f2* An alternative expression to evaluate, so that the speed of two expressions can be directly compared. The default is `[]`.

*tol* Tolerance used to compare the results of expression *f* and expression *f2*. If *tol* is positive, the tolerance is an absolute one. If *tol* is negative, the tolerance is a relative one. The default is `eps`. If *tol* is `Inf`, then no comparison will be made.

*order* The time complexity of the expression $O(a * n^p)$. This is a structure with fields `a` and `p`.

*n* The values *n* for which the expression was calculated **AND** the execution time was greater than zero.

*T_f* The nonzero execution times recorded for the expression *f* in seconds.

*T_f2* The nonzero execution times recorded for the expression *f2* in seconds. If required, the mean time ratio is simply `mean (T_f ./ T_f2)`.

The slope of the execution time graph shows the approximate power of the asymptotic running time $O(n^p)$. This power is plotted for the region over which it is approximated (the latter half of the graph). The estimated power is not very accurate, but should be sufficient to determine the general order of an algorithm. It should indicate if, for example, the implementation is unexpectedly $O(n^2)$ rather than $O(n)$ because it extends a vector each time through the loop rather than pre-allocating storage. In the current version of Octave, the following is not the expected $O(n)$.

```
speed ("for i = 1:n, y{i} = x(i); endfor", "", [1000, 10000])
```

But it is if you preallocate the cell array `y`:

```
speed ("for i = 1:n, y{i} = x(i); endfor", ...
       "x = rand (n, 1); y = cell (size (x));", [1000, 10000])
```

An attempt is made to approximate the cost of individual operations, but it is wildly inaccurate. You can improve the stability somewhat by doing more work for each `n`. For example:

```
speed ("airy(x)", "x = rand (n, 10)", [10000, 100000])
```

When comparing two different expressions (*f*, *f2*), the slope of the line on the speedup ratio graph should be larger than 1 if the new expression is faster. Better algorithms have a shallow slope. Generally, vectorizing an algorithm will not change the slope of the execution time graph, but will shift it relative to the original. For example:

```
speed ("sum (x)", "", [10000, 100000], ...
       "v = 0; for i = 1:length (x), v += x(i); endfor")
```

The following is a more complex example. If there was an original version of `xcorr` using for loops and a second version using an FFT, then one could compare the run

speed for various lags as follows, or for a fixed lag with varying vector lengths as follows:

```
speed ("xcorr (x, n)", "x = rand (128, 1);", 100,
       "xcorr_orig (x, n)", -100*eps)
speed ("xcorr (x, 15)", "x = rand (20+n, 1);", 100,
       "xcorr_orig (x, n)", -100*eps)
```

Assuming one of the two versions is in xcorr_orig, this would compare their speed and their output values. Note that the FFT version is not exact, so one must specify an acceptable tolerance on the comparison `100*eps`. In this case, the comparison should be computed relatively, as `abs ((x - y) ./ y)` rather than absolutely as `abs (x - y)`.

Type *`example ("speed")`* to see some real examples or *`demo ("speed")`* to run them.

# Appendix C Obsolete Functions

After being marked as deprecated for two major releases, the following functions have been removed from Octave. The third column of the table shows the version of Octave in which the function was removed. Prior to removal, each function in the list was marked as deprecated for at least two major releases. All deprecated functions issue warnings explaining that they will be removed in a future version of Octave, and which function should be used instead.

Replacement functions do not always accept precisely the same arguments as the obsolete function, but should provide equivalent functionality.

| **Obsolete Function** | **Replacement** | **Version** |
|---|---|---|
| `beta_cdf` | `betacdf` in Octave Forge statistics pkg | 3.4.0 |
| `beta_inv` | `betainv` in Octave Forge statistics pkg | 3.4.0 |
| `beta_pdf` | `betapdf` in Octave Forge statistics pkg | 3.4.0 |
| `beta_rnd` | `betarnd` in Octave Forge statistics pkg | 3.4.0 |
| `binomial_cdf` | `binocdf` in Octave Forge statistics pkg | 3.4.0 |
| `binomial_inv` | `binoinv` in Octave Forge statistics pkg | 3.4.0 |
| `binomial_pdf` | `binopdf` in Octave Forge statistics pkg | 3.4.0 |
| `binomial_rnd` | `binornd` in Octave Forge statistics pkg | 3.4.0 |
| `chisquare_cdf` | `chi2cdf` in Octave Forge statistics pkg | 3.4.0 |
| `chisquare_inv` | `chi2inv` in Octave Forge statistics pkg | 3.4.0 |
| `chisquare_pdf` | `chi2pdf` in Octave Forge statistics pkg | 3.4.0 |
| `chisquare_rnd` | `chi2rnd` in Octave Forge statistics pkg | 3.4.0 |
| `clearplot` | `clf` | 3.4.0 |
| `com2str` | `num2str` | 3.4.0 |
| `exponential_cdf` | `expcdf` in Octave Forge statistics pkg | 3.4.0 |
| `exponential_inv` | `expinv` in Octave Forge statistics pkg | 3.4.0 |
| `exponential_pdf` | `exppdf` in Octave Forge statistics pkg | 3.4.0 |

| | | |
|---|---|---|
| exponential_rnd | exprnd in Octave Forge statistics pkg | 3.4.0 |
| f_cdf | fcdf in Octave Forge statistics pkg | 3.4.0 |
| f_inv | finv in Octave Forge statistics pkg | 3.4.0 |
| f_pdf | fpdf in Octave Forge statistics pkg | 3.4.0 |
| f_rnd | frnd in Octave Forge statistics pkg | 3.4.0 |
| gamma_cdf | gamcdf in Octave Forge statistics pkg | 3.4.0 |
| gamma_inv | gaminv in Octave Forge statistics pkg | 3.4.0 |
| gamma_pdf | gampdf in Octave Forge statistics pkg | 3.4.0 |
| gamma_rnd | gamrnd in Octave Forge statistics pkg | 3.4.0 |
| geometric_cdf | geocdf in Octave Forge statistics pkg | 3.4.0 |
| geometric_inv | geoinv in Octave Forge statistics pkg | 3.4.0 |
| geometric_pdf | geopdf in Octave Forge statistics pkg | 3.4.0 |
| geometric_rnd | geornd in Octave Forge statistics pkg | 3.4.0 |
| hypergeometric_cdf | hygecdf in Octave Forge statistics pkg | 3.4.0 |
| hypergeometric_inv | hygeinv in Octave Forge statistics pkg | 3.4.0 |
| hypergeometric_pdf | hygepdf in Octave Forge statistics pkg | 3.4.0 |
| hypergeometric_rnd | hygernd in Octave Forge statistics pkg | 3.4.0 |
| intersection | intersect | 3.4.0 |
| is_bool | isbool | 3.4.0 |
| is_complex | iscomplex | 3.4.0 |
| is_list | None | 3.4.0 |
| is_matrix | ismatrix | 3.4.0 |
| is_scalar | isscalar | 3.4.0 |
| is_square | issquare | 3.4.0 |
| is_stream | None | 3.4.0 |
| is_struct | isstruct | 3.4.0 |
| is_symmetric | issymmetric | 3.4.0 |
| is_vector | isvector | 3.4.0 |

| | | |
|---|---|---|
| `lognormal_cdf` | `logncdf` in Octave Forge statistics pkg | 3.4.0 |
| `lognormal_inv` | `logninv` in Octave Forge statistics pkg | 3.4.0 |
| `lognormal_pdf` | `lognpdf` in Octave Forge statistics pkg | 3.4.0 |
| `lognormal_rnd` | `lognrnd` in Octave Forge statistics pkg | 3.4.0 |
| `meshdom` | `meshgrid` | 3.4.0 |
| `normal_cdf` | `normcdf` in Octave Forge statistics pkg | 3.4.0 |
| `normal_inv` | `norminv` in Octave Forge statistics pkg | 3.4.0 |
| `normal_pdf` | `normpdf` in Octave Forge statistics pkg | 3.4.0 |
| `normal_rnd` | `normrnd` in Octave Forge statistics pkg | 3.4.0 |
| `pascal_cdf` | `nbincdf` in Octave Forge statistics pkg | 3.4.0 |
| `pascal_inv` | `nbininv` in Octave Forge statistics pkg | 3.4.0 |
| `pascal_pdf` | `nbinpdf` in Octave Forge statistics pkg | 3.4.0 |
| `pascal_rnd` | `nbinrnd` in Octave Forge statistics pkg | 3.4.0 |
| `poisson_cdf` | `poisscdf` in Octave Forge statistics pkg | 3.4.0 |
| `poisson_inv` | `poissinv` in Octave Forge statistics pkg | 3.4.0 |
| `poisson_pdf` | `poisspdf` in Octave Forge statistics pkg | 3.4.0 |
| `poisson_rnd` | `poissrnd` in Octave Forge statistics pkg | 3.4.0 |
| `polyinteg` | `polyint` | 3.4.0 |
| `struct_contains` | `isfield` | 3.4.0 |
| `struct_elements` | `fieldnames` | 3.4.0 |
| `t_cdf` | `tcdf` in Octave Forge statistics pkg | 3.4.0 |
| `t_inv` | `tinv` in Octave Forge statistics pkg | 3.4.0 |
| `t_pdf` | `tpdf` in Octave Forge statistics pkg | 3.4.0 |
| `t_rnd` | `trnd` in Octave Forge statistics pkg | 3.4.0 |

| | | |
|---|---|---|
| `uniform_cdf` | `unifcdf` in Octave Forge statistics pkg | 3.4.0 |
| `uniform_inv` | `unifinv` in Octave Forge statistics pkg | 3.4.0 |
| `uniform_pdf` | `unifpdf` in Octave Forge statistics pkg | 3.4.0 |
| `uniform_rnd` | `unifrnd` in Octave Forge statistics pkg | 3.4.0 |
| `weibull_cdf` | `wblcdf` in Octave Forge statistics pkg | 3.4.0 |
| `weibull_inv` | `wblinv` in Octave Forge statistics pkg | 3.4.0 |
| `weibull_pdf` | `wblpdf` in Octave Forge statistics pkg | 3.4.0 |
| `weibull_rnd` | `wblrnd` in Octave Forge statistics pkg | 3.4.0 |
| `wiener_rnd` | `wienrnd` in Octave Forge statistics pkg | 3.4.0 |
| `create_set` | `unique` | 3.6.0 |
| `dmult` | `diag (A) * B` | 3.6.0 |
| `iscommand` | None | 3.6.0 |
| `israwcommand` | None | 3.6.0 |
| `lchol` | `chol (..., "lower")` | 3.6.0 |
| `loadimage` | `load` or `imread` | 3.6.0 |
| `mark_as_command` | None | 3.6.0 |
| `mark_as_rawcommand` | None | 3.6.0 |
| `spatan2` | `atan2` | 3.6.0 |
| `spchol` | `chol` | 3.6.0 |
| `spchol2inv` | `chol2inv` | 3.6.0 |
| `spcholinv` | `cholinv` | 3.6.0 |
| `spcumprod` | `cumprod` | 3.6.0 |
| `spcumsum` | `cumsum` | 3.6.0 |
| `spdet` | `det` | 3.6.0 |
| `spdiag` | `sparse (diag (...))` | 3.6.0 |
| `spfind` | `find` | 3.6.0 |
| `sphcat` | `horzcat` | 3.6.0 |
| `spinv` | `inv` | 3.6.0 |
| `spkron` | `kron` | 3.6.0 |
| `splchol` | `chol (..., "lower")` | 3.6.0 |
| `split` | `char (strsplit (s, t))` | 3.6.0 |
| `splu` | `lu` | 3.6.0 |
| `spmax` | `max` | 3.6.0 |
| `spmin` | `min` | 3.6.0 |
| `spprod` | `prod` | 3.6.0 |
| `spqr` | `qr` | 3.6.0 |

| | | |
|---|---|---|
| spsum | sum | 3.6.0 |
| spsumsq | sumsq | 3.6.0 |
| spvcat | vertcat | 3.6.0 |
| str2mat | char | 3.6.0 |
| unmark_command | None | 3.6.0 |
| unmark_rawcommand | None | 3.6.0 |
| autocor | xcorr in Octave Forge signal pkg | 3.8.0 |
| autocov | xcov in Octave Forge signal pkg | 3.8.0 |
| betai | betainc | 3.8.0 |
| cellidx | ismember | 3.8.0 |
| cquad | quadcc | 3.8.0 |
| dispatch | None | 3.8.0 |
| fstat | stat | 3.8.0 |
| gammai | gammainc | 3.8.0 |
| glpkmex | glpk | 3.8.0 |
| is_duplicate_entry | unique | 3.8.0 |
| is_global | isglobal | 3.8.0 |
| krylovb | [Uret, ~, Ucols] = krylov (...) | 3.8.0 |
| perror | None | 3.8.0 |
| replot | refresh | 3.8.0 |
| saveimage | imwrite | 3.8.0 |
| setstr | char | 3.8.0 |
| strerror | None | 3.8.0 |
| values | unique | 3.8.0 |
| cut | histc | 4.0.0 |
| cor | corr | 4.0.0 |
| corrcoef | corr | 4.0.0 |
| __error_text__ | lasterr | 4.0.0 |
| error_text | lasterr | 4.0.0 |
| polyderiv | polyder | 4.0.0 |
| shell_cmd | system | 4.0.0 |
| studentize | zscore | 4.0.0 |
| sylvester_matrix | hadamard (2^k) | 4.0.0 |
| default_save_options | save_default_options | 4.2.0 |
| gen_doc_cache | doc_cache_create | 4.2.0 |
| interp1q | interp1 | 4.2.0 |
| isequalwithequalnans | isequaln | 4.2.0 |
| java_convert_matrix | java_matrix_autoconversion | 4.2.0 |
| java_debug | debug_java | 4.2.0 |
| java_invoke | javaMethod | 4.2.0 |
| java_new | javaObject | 4.2.0 |

| | | |
|---|---|---|
| java_unsigned_conversion | java_unsigned_autoconversion | 4.2.0 |
| javafields | fieldnames | 4.2.0 |
| javamethods | methods | 4.2.0 |
| re_read_readline_init_file | readline_re_read_init_file | 4.2.0 |
| read_readline_init_file | readline_read_init_file | 4.2.0 |
| saving_history | history_save | 4.2.0 |
| allow_noninteger_range_as_index | None | 4.4.0 |
| bicubic | interp2 | 4.4.0 |
| delaunay3 | delaunay | 4.4.0 |
| do_braindead_shortcircuit_evaluation | None | 4.4.0 |
| dump_prefs | None | 4.4.0 |
| find_dir_in_path | dir_in_loadpath | 4.4.0 |
| finite | isfinite | 4.4.0 |
| fmod | rem | 4.4.0 |
| fnmatch | glob or regexp | 4.4.0 |
| gmap40 | None | 4.4.0 |
| loadaudio | audioread | 4.4.0 |
| luinc | ichol or ilu | 4.4.0 |
| mouse_wheel_zoom | mousehweelzoom property | 4.4.0 |
| nfields | numfields | 4.4.0 |
| octave_tmp_file_name | tempname | 4.4.0 |
| playaudio | audioplayer | 4.4.0 |
| saveaudio | audiowrite | 4.4.0 |
| setaudio | None | 4.4.0 |
| syl | sylvester | 4.4.0 |
| usage | print_usage | 4.4.0 |
| bitmax | flintmax | 5.1.0 |
| mahalanobis | mahal in Octave Forge statistics pkg | 5.1.0 |
| md5sum | hash | 5.1.0 |
| octave_config_info | __octave_config_info__ | 5.1.0 |
| onenormest | normest1 | 5.1.0 |
| sleep | pause | 5.1.0 |
| usleep | pause | 5.1.0 |
| wavread | audioread | 5.1.0 |
| wavwrite | audiowrite | 5.1.0 |
| output_max_field_width | output_precision | 7.1.0 |
| is_keyword | iskeyword | 7.1.0 |

| | | |
|---|---|---|
| runtests | oruntests | 8.1.0 |
| disable_diagonal_ matrix | optimize_diagonal_ matrix | 9.1.0 |
| disable_permutation_ matrix | optimize_ permutation_matrix | 9.1.0 |
| disable_range | optimize_range | 9.1.0 |
| shift | circshift | 10.1.0 |
| sparse_auto_mutate | None | 10.1.0 |

# Appendix D Known Causes of Trouble

This section describes known problems that affect users of Octave. Most of these are not Octave bugs per se—if they were, we would fix them. But the result for a user may be like the result of a bug.

Some of these problems are due to bugs in other software, some are missing features that are too much work to add, and some are places where people's opinions differ as to what is best.

## D.1 Actual Bugs We Haven't Fixed Yet

- Output that comes directly from Fortran functions is not sent through the pager and may appear out of sequence with other output that is sent through the pager. One way to avoid this is to force pending output to be flushed before calling a function that will produce output from within Fortran functions. To do this, use the command

  ```
  fflush (stdout)
  ```

  Another possible workaround is to use the command

  ```
  page_screen_output (false);
  ```

  to turn the pager off.

A list of ideas for future enhancements is distributed with Octave. See the file `PROJECTS` in the top level directory in the source distribution.

## D.2 Reporting Bugs

Your bug reports play an essential role in making Octave reliable.

When you encounter a problem, the first thing to do is to see if it is already known. See Appendix D [Trouble], page 1173. If it isn't known, then you should report the problem.

Reporting a bug may help you by bringing a solution to your problem, or it may not. In any case, the principal function of a bug report is to help the entire community by making the next version of Octave work better. Bug reports are your contribution to the maintenance of Octave.

In order for a bug report to serve its purpose, you must include the information that makes it possible to fix the bug.

### D.2.1 Have You Found a Bug?

If you are not sure whether you have found a bug, here are some guidelines:

- If Octave gets a fatal signal, for any input whatever, that is a bug. Reliable interpreters never crash.
- If Octave produces incorrect results, for any input whatever, that is a bug.
- Some output may appear to be incorrect when it is in fact due to a program whose behavior is undefined, which happened by chance to give the desired results on another system. For example, the range operator may produce different results because of differences in the way floating point arithmetic is handled on various systems.
- If Octave produces an error message for valid input, that is a bug.

- If Octave does not produce an error message for invalid input, that is a bug. However, you should note that your idea of "invalid input" might be my idea of "an extension" or "support for traditional practice".
- If you are an experienced user of programs like Octave, your suggestions for improvement are welcome in any case.

### D.2.2 Where to Report Bugs

To report a bug in Octave, submit a bug report to the Octave bug tracker `https://bugs.octave.org`.

**Do not send bug reports to 'help-octave'**. Most users of Octave do not want to receive bug reports.

### D.2.3 How to Report Bugs

Submit bug reports for Octave to the Octave bug tracker `https://bugs.octave.org`.

The fundamental principle of reporting bugs usefully is this: **report all the facts**. If you are not sure whether to state a fact or leave it out, state it!

Often people omit facts because they think they know what causes the problem and they conclude that some details don't matter. Thus, you might assume that the name of the variable you use in an example does not matter. Well, probably it doesn't, but one cannot be sure. Perhaps the bug is a stray memory reference which happens to fetch from the location where that name is stored in memory; perhaps, if the name were different, the contents of that location would fool the interpreter into doing the right thing despite the bug. Play it safe and give a specific, complete example.

Keep in mind that the purpose of a bug report is to enable someone to fix the bug if it is not known. Always write your bug reports on the assumption that the bug is not known.

Sometimes people give a few sketchy facts and ask, "Does this ring a bell?" This cannot help us fix a bug. It is better to send a complete bug report to begin with.

Try to make your bug report self-contained. If we have to ask you for more information, it is best if you include all the previous information in your response, as well as the information that was missing.

To enable someone to investigate the bug, you should include all these things:

- The version of Octave. You can get this by noting the version number that is printed when Octave starts, or running it with the '`-v`' option.
- A complete input file that will reproduce the bug.

  A single statement may not be enough of an example—the bug might depend on other details that are missing from the single statement where the error finally occurs.
- The command arguments you gave Octave to execute that example and observe the bug. To guarantee you won't omit something important, list all the options.

  If we were to try to guess the arguments, we would probably guess wrong and then we would not encounter the bug.
- The type of machine you are using, and the operating system name and version number.
- The command-line arguments you gave to the `configure` command when you installed the interpreter.

- A complete list of any modifications you have made to the interpreter source.

  Be precise about these changes—show a context diff for them.
- Details of any other deviations from the standard procedure for installing Octave.
- A description of what behavior you observe that you believe is incorrect. For example, "The interpreter gets a fatal signal," or, "The output produced at line 208 is incorrect."

  Of course, if the bug is that the interpreter gets a fatal signal, then one can't miss it. But if the bug is incorrect output, we might not notice unless it is glaringly wrong.

  Even if the problem you experience is a fatal signal, you should still say so explicitly. Suppose something strange is going on, such as, your copy of the interpreter is out of sync, or you have encountered a bug in the C library on your system. Your copy might crash and the copy here would not. If you said to expect a crash, then when the interpreter here fails to crash, we would know that the bug was not happening. If you don't say to expect a crash, then we would not know whether the bug was happening. We would not be able to draw any conclusion from our observations.

  Often the observed symptom is incorrect output when your program is run. Unfortunately, this is not enough information unless the program is short and simple. It is very helpful if you can include an explanation of the expected output, and why the actual output is incorrect.
- If you wish to suggest changes to the Octave source, send them as context diffs. If you even discuss something in the Octave source, refer to it by context, not by line number, because the line numbers in the development sources probably won't match those in your sources.

Here are some things that are not necessary:

- A description of the envelope of the bug.

  Often people who encounter a bug spend a lot of time investigating which changes to the input file will make the bug go away and which changes will not affect it. Such information is usually not necessary to enable us to fix bugs in Octave, but if you can find a simpler example to report *instead* of the original one, that is a convenience. Errors in the output will be easier to spot, running under the debugger will take less time, etc. Most Octave bugs involve just one function, so the most straightforward way to simplify an example is to delete all the function definitions except the one in which the bug occurs.

  However, simplification is not vital; if you don't want to do this, report the bug anyway and send the entire test case you used.
- A patch for the bug. Patches can be helpful, but if you find a bug, you should report it, even if you cannot send a fix for the problem.

### D.2.4 Sending Patches for Octave

If you would like to write bug fixes or improvements for Octave, that is very helpful. When you send your changes, please follow these guidelines to avoid causing extra work for us in studying the patches.

If you don't follow these guidelines, your information might still be useful, but using it will take extra work. Maintaining Octave is a lot of work in the best of circumstances, and we can't keep up unless you do your best to help.

- Send an explanation with your changes of what problem they fix or what improvement they bring about. For a bug fix, just include a copy of the bug report, and explain why the change fixes the bug.
- Always include a proper bug report for the problem you think you have fixed. We need to convince ourselves that the change is right before installing it. Even if it is right, we might have trouble judging it if we don't have a way to reproduce the problem.
- Include all the comments that are appropriate to help people reading the source in the future understand why this change was needed.
- Don't mix together changes made for different reasons. Send them *individually*.

  If you make two changes for separate reasons, then we might not want to install them both. We might want to install just one.
- Use '`diff -c`' to make your diffs. Diffs without context are hard for us to install reliably. More than that, they make it hard for us to study the diffs to decide whether we want to install them. Unified diff format is better than contextless diffs, but not as easy to read as '`-c`' format.

  If you have GNU diff, use '`diff -cp`', which shows the name of the function that each change occurs in.
- Write the change log entries for your changes.

  Read the `ChangeLog` file to see what sorts of information to put in, and to learn the style that we use. The purpose of the change log is to show people where to find what was changed. So you need to be specific about what functions you changed; in large functions, it's often helpful to indicate where within the function the change was made.

  On the other hand, once you have shown people where to find the change, you need not explain its purpose. Thus, if you add a new function, all you need to say about it is that it is new. If you feel that the purpose needs explaining, it probably does—but the explanation will be much more useful if you put it in comments in the code.

  If you would like your name to appear in the header line for who made the change, send us the header line.

## D.3 How To Get Help with Octave

The mailing list `help@octave.org` exists for the discussion of matters related to using and installing Octave. If would like to join the discussion, please send a short note to `help-request@octave.org`.

**Please do not** send requests to be added or removed from the mailing list, or other administrative trivia to the list itself.

If you think you have found a bug in Octave or in the installation procedure, however, you should submit a complete bug report to the Octave bug tracker at `https://bugs.octave.org`. But before you submit a bug report, please read `https://www.octave.org/bugs.html` to learn how to submit a useful bug report.

## D.4 How to Distinguish Between Octave and Matlab

Octave and MATLAB are very similar, but handle Java slightly different. Therefore it may be necessary to detect the environment and use the appropriate functions. The following

function can be used to detect the environment. Due to the persistent variable it can be called repeatedly without a heavy performance hit.

Example:

```
%%
%% Return: true if the environment is Octave.
%%
function retval = isOctave
  persistent cacheval;  % speeds up repeated calls

  if isempty (cacheval)
    cacheval = (exist ("OCTAVE_VERSION", "builtin") > 0);
  end

  retval = cacheval;
end
```

# Appendix E Installing Octave

The procedure for installing Octave from source on a Unix-like system is described next. Building on other platforms will follow similar steps. Note that this description applies to Octave releases. Building the development sources from the Mercurial archive requires additional steps as described in the development source itself.

## E.1 Build Dependencies

Octave is a fairly large program with many build dependencies. You may be able to find pre-packaged versions of the dependencies distributed as part of your system, or you may have to build some or all of them yourself.

### E.1.1 Obtaining the Dependencies Automatically

On some systems you can obtain many of Octave's build dependencies automatically. The commands for doing this vary by system. Similarly, the names of pre-compiled packages vary by system and do not always match exactly the names listed in Section E.1.2 [Build Tools], page 1179, and Section E.1.3 [External Packages], page 1180.

You will usually need the development version of an external dependency so that you get the libraries and header files for building software, not just for running already compiled programs. These packages typically have names that end with the suffix `-dev` or `-devel`.

On systems with `apt-get` (Debian, Ubuntu, etc.), you may be able to install most of the tools and external packages using a command similar to

```
apt-get build-dep octave
```

The specific package name may be `octave3.2` or `octave3.4`. The set of required tools and external dependencies does not change frequently, so it is not important that the version match exactly, but you should use the most recent one available.

On systems with `yum` (Fedora, Red Hat, etc.), you may be able to install most of the tools and external packages using a command similar to

```
yum-builddep octave
```

The `yum-builddep` utility is part of the `yum-utils` package.

For either type of system, the package name may include a version number. The set of required tools and external dependencies does not change frequently, so it is not important that the version exactly match the version you are installing, but you should use the most recent one available.

### E.1.2 Build Tools

The following tools are required:

C++, C, and Fortran compilers
: The Octave sources are primarily written in C++, but some portions are also written in C and Fortran. The Octave sources are intended to be portable. Recent versions of the GNU compiler collection (GCC) should work (`https://gcc.gnu.org`). If you use GCC, you should avoid mixing versions. For example, be sure that you are not using the obsolete `g77` Fortran compiler with modern versions of `gcc` and `g++`.

GNU Make
: Tool for building software (https://www.gnu.org/software/make). Octave's build system requires GNU Make. Other versions of Make will not work. Fortunately, GNU Make is highly portable and easy to install.

AWK, sed, and other Unix utilities
: Basic Unix system utilities are required for building Octave. All will be available with any modern Unix system and also on Windows with either Cygwin or MinGW and MSYS.

Additionally, the following tools may be needed:

Bison
: Parser generator (https://www.gnu.org/software/bison). You will need Bison if you modify the `oct-parse.yy` source file or if you delete the files that are generated from it.

Flex
: Lexer analyzer (https://www.gnu.org/software/flex). You will need Flex if you modify the `lex.ll` source file or if you delete the files that are generated from it.

Autoconf
: Package for software configuration (https://www.gnu.org/software/autoconf). Autoconf is required if you modify Octave's `configure.ac` file or other files that it requires.

Automake
: Package for Makefile generation (https://www.gnu.org/software/automake). Automake is required if you modify Octave's `Makefile.am` files or other files that they depend on.

Libtool
: Package for building software libraries (https://www.gnu.org/software/libtool). Libtool is required by Automake.

gperf
: Perfect hash function generator (https://www.gnu.org/software/gperf). You will need gperf if you modify the `octave.gperf` file or if you delete the file that is generated from it.

Texinfo
: Package for generating online and print documentation (https://www.gnu.org/software/texinfo). You will need Texinfo to build Octave's documentation or if you modify the documentation source files or the docstring of any Octave function.

### E.1.3 External Packages

The following external packages are required:

BLAS
: Basic Linear Algebra Subroutine library. Accelerated BLAS libraries such as OpenBLAS (https://www.openblas.net/) or ATLAS (http://math-atlas.sourceforge.net) are recommended for best performance. The reference implementation (http://www.netlib.org/blas) is slow and suffers from certain bugs in corner case inputs.

LAPACK
: Linear Algebra Package (http://www.netlib.org/lapack).

PCRE
: The Perl Compatible Regular Expression library (https://www.pcre.org).

The following external package is optional but strongly recommended:

GNU Readline
: Command-line editing library (https://www.gnu.org/s/readline).

If you wish to build Octave without GNU readline installed, you must use the `--disable-readline` option when running the configure script.

The following external software packages are optional. Octave can be built without them but certain features might be missing:

ARPACK
: Library for the solution of large-scale eigenvalue problems (https://forge.scilab.org/index.php/p/arpack-ng). ARPACK is required to provide the functions `eigs` and `svds`.

cURL
: Library for transferring data with URL syntax (https://curl.haxx.se). cURL is required to provide the `urlread` and `urlwrite` functions and the `ftp` class.

FFTW3
: Library for computing discrete Fourier transforms (http://www.fftw.org). FFTW3 is used to provide better performance for functions that compute discrete Fourier transforms (`fft`, `ifft`, `fft2`, etc.)

FLTK
: Portable GUI toolkit (http://www.fltk.org). FLTK can be used to provide windows for Octave's OpenGL-based graphics functions.

fontconfig
: Library for configuring and customizing font access (https://www.freedesktop.org/wiki/Software/fontconfig). Fontconfig is used to manage fonts for Octave's OpenGL-based graphics functions.

FreeType
: Portable font engine (https://www.freetype.org). FreeType is used to perform font rendering for Octave's OpenGL-based graphics functions.

GLPK
: GNU Linear Programming Kit (https://www.gnu.org/software/glpk). GPLK is required for the function `glpk`.

gl2ps
: OpenGL to PostScript printing library (https://www.geuz.org/gl2ps/). gl2ps is required for printing when using OpenGL-based graphics toolkits (currently either FLTK or Qt).

gnuplot
: Interactive graphics program (http://www.gnuplot.info). gnuplot can be used as a graphics renderer for Octave; prior to Octave 4.0, gnuplot was the default graphics renderer.

GraphicsMagick++
: Image processing library (http://www.graphicsmagick.org). GraphicsMagick++ is used to provide the `imread` and `imwrite` functions.

HDF5
: Library for manipulating portable data files (https://www.hdfgroup.org/HDF5). HDF5 is required for Octave's `load` and `save` commands to read and write HDF data files.

Java Development Kit
: Java programming language compiler and libraries. The OpenJDK free software implementation is recommended (http://openjdk.java.net/), although other JDK implementations may work. Java is required to be able to call Java functions from within Octave.

OpenGL
: API for portable 2-D and 3-D graphics (https://www.opengl.org). An OpenGL implementation can be used to provide a renderer for Octave's graphics functions. Octave's OpenGL-based graphics functions usually outperform the gnuplot-based graphics functions because plot data can be rendered directly instead of sending data and commands to gnuplot for interpretation and rendering. Since Octave 4.0, the default graphics renderer (`"qt"`) has been OpenGL-based.

PortAudio
: PortAudio (http://www.portaudio.com/) provides a very simple API for recording and/or playing sound using a simple callback function or a blocking read/write interface. It is required for the audio processing functions `audioplayer`, `audiorecorder`, and `audiodevinfo`.

Qhull
: Computational geometry library (http://www.qhull.org). Qhull is required to provide the functions `convhull`, `convhulln`, `delaunay`, `delaunayn`, `voronoi`, and `voronoin`.

QRUPDATE
: QR factorization updating library (https://sourceforge.net/projects/qrupdate). QRUPDATE is used to provide improved performance for the functions `qrdelete`, `qrinsert`, `qrshift`, and `qrupdate`.

QScintilla
: Source code highlighter and manipulator; a Qt port of Scintilla (http://www.riverbankcomputing.co.uk/software/qscintilla). QScintilla is used for syntax highlighting and code completion in the GUI.

Qt
: GUI and utility libraries (https://www.qt.io). Qt is required for building the GUI. It is a large framework, but the only components required are the GUI, core, and network modules. Since Octave 4.0, the default graphics renderer (`"qt"`) has been Qt-based, which has been OpenGL-based.

RapidJSON
: A fast JSON parser/generator for C++ with both SAX/DOM style API (https://rapidjson.org/). RapidJSON is required to read or write from or to JSON files with the functions `jsondecode` and `jsonencode`.

SuiteSparse
: Sparse matrix factorization library (http://faculty.cse.tamu.edu/davis/suitesparse.html). SuiteSparse is required to provide sparse matrix factorizations and solution of linear equations for sparse systems.

SUNDIALS
: The SUite of Nonlinear and DIfferential/ALgebraic Equation Solvers (https://computation.llnl.gov/projects/sundials) is required for the Ordinary Differential Equations (ODE) solvers `ode15i` and `ode15s`.

zlib
: Data compression library (https://zlib.net). The zlib library is required for Octave's `load` and `save` commands to handle compressed data, including MATLAB v5 MAT files.

## E.2 Running Configure and Make

- Run the shell script `configure`. This will determine the features your system has (or doesn't have) and create a file named `Makefile` from each of the files named `Makefile.in`.

  For a complete list of configure options, run `configure --help`. Here is a summary of the configure options that are most frequently used when building Octave:

  `--help` Print a summary of the options recognized by the configure script.

  `--prefix=`*prefix*
  Install Octave in subdirectories below *prefix*. The default value of *prefix* is `/usr/local`.

  `--srcdir=`*dir*
  Look for Octave sources in the directory *dir*.

  `--disable-64`
  Disable using 64-bit integers for indexing arrays and use 32-bit integers instead. On systems with 32-bit pointers, this option is always disabled. If the configure script determines that your BLAS library uses 32-bit integers, then operations using the following libraries are limited to arrays with dimensions that are smaller than $2^{31}$ elements:

  - BLAS
  - LAPACK
  - QRUPDATE
  - SuiteSparse
  - ARPACK

  Additionally, the following libraries use `int` internally, so maximum problem sizes are always limited:

  - GLPK
  - Qhull

  See Section E.3 [Compiling Octave with 64-bit Indexing], page 1187, for more details about building Octave with more complete support for large arrays.

  `--enable-address-sanitizer-flags`
  Enable compiler options `-fsanitize=address` and `-fomit-frame-pointer` for memory access checking. This option is primarily used for debugging Octave. Building Octave with this option has a negative impact on performance and is not recommended for general use. It may also interfere with proper functioning of the GUI.

  `--disable-docs`
  Disable building all forms of the documentation (Info, PDF, HTML). The default is to build documentation, but your system will need functioning Texinfo and TEX installs for this to succeed.

`--enable-float-truncate`
: This option allows for truncation of intermediate floating point results in calculations. It is only necessary for certain platforms.

`--enable-readline`
: Use the readline library to provide for editing of the command line in terminal environments. This option is on by default.

`--enable-shared`
: Create shared libraries (this is the default). If you are planning to use the dynamic loading features, you will probably want to use this option. It will make your `.oct` files much smaller and on some systems it may be necessary to build shared libraries in order to use dynamically linked functions.

  You may also want to build a shared version of `libstdc++`, if your system doesn't already have one.

`--with-blas=<lib>`
: By default, configure looks for the best BLAS matrix libraries on your system, including optimized implementations such as the free ATLAS 3.0, as well as vendor-tuned libraries. (The use of an optimized BLAS will generally result in several-times faster matrix operations.) Use this option to specify a particular BLAS library that Octave should use.

`--with-lapack=<lib>`
: By default, configure looks for the best LAPACK matrix libraries on your system, including optimized implementations such as the free ATLAS 3.0, as well as vendor-tuned libraries. (The use of an optimized LAPACK will generally result in several-times faster matrix operations.) Use this option to specify a particular LAPACK library that Octave should use.

`--with-magick=<lib>`
: Select the Magick++ library to use for image I/O. For many distributions, possible values are `"GraphicsMagick++"` (default) or `"ImageMagick++"`.

`--with-sepchar=<char>`
: Use `<char>` as the path separation character. This option can help when running Octave on non-Unix systems.

`--without-amd`
: Don't use AMD, disable some sparse matrix functionality.

`--without-camd`
: Don't use CAMD, disable some sparse matrix functionality.

`--without-colamd`
: Don't use COLAMD, disable some sparse matrix functionality.

`--without-ccolamd`
: Don't use CCOLAMD, disable some sparse matrix functionality.

`--without-cholmod`
: Don't use CHOLMOD, disable some sparse matrix functionality.

`--without-curl`
: Don't use the cURL library, disable the ftp objects, `urlread` and `urlwrite` functions.

`--without-cxsparse`
: Don't use CXSPARSE, disable some sparse matrix functionality.

`--without-fftw3`
: Use the included FFTPACK library for computing Fast Fourier Transforms instead of the FFTW3 library.

`--without-fftw3f`
: Use the included FFTPACK library for computing Fast Fourier Transforms instead of the FFTW3 library when operating on single precision (float) values.

`--without-glpk`
: Don't use the GLPK library for linear programming.

`--without-hdf5`
: Don't use the HDF5 library, disable reading and writing of HDF5 files.

`--without-opengl`
: Don't use OpenGL, disable native graphics toolkit for plotting. You will need `gnuplot` installed in order to make plots.

`--without-qhull_r`
: Don't use (re-entrant) Qhull, disable `delaunay`, `convhull`, and related functions.

`--without-qrupdate`
: Don't use QRUPDATE, disable QR and Cholesky update functions.

`--without-umfpack`
: Don't use UMFPACK, disable some sparse matrix functionality.

`--without-z`
: Don't use the zlib library, disable data file compression and support for recent MAT file formats.

`--without-framework-carbon`
: Don't use framework Carbon headers, libraries, or specific source code even if the configure test succeeds (the default is to use Carbon framework if available). This is a platform specific configure option for Mac systems.

`--without-framework-opengl`
: Don't use framework OpenGL headers, libraries, or specific source code even if the configure test succeeds. If this option is given then OpenGL headers and libraries in standard system locations are tested (the default value is `--with-framework-opengl`). This is a platform specific configure option for Mac systems.

See the file `INSTALL` for more general information about the command line options used by configure. That file also contains instructions for compiling in a directory other than the one where the source is located.

- Run make.

  You will need a recent version of GNU Make as Octave relies on certain features not generally available in all versions of make. Modifying Octave's makefiles to work with other make programs is probably not worth your time; instead, we simply recommend installing GNU Make.

  There are currently three options for plotting in Octave: the external program gnuplot, the internal graphics engine using OpenGL coupled with either FLTK or Qt widgets. Gnuplot is a command-driven interactive function plotting program.

  To compile Octave, you will need a recent version of `g++` or other ANSI C++ compiler. In addition, you will need a Fortran 77 compiler or `f2c`. If you use `f2c`, you will need a script like `fort77` that works like a normal Fortran compiler by combining `f2c` with your C compiler in a single script.

  If you plan to modify the parser you will also need GNU `bison` and `flex`. If you modify the documentation, you will need GNU Texinfo.

  GNU Make, `gcc` (and `libstdc++`), `gnuplot`, `bison`, `flex`, and Texinfo are all available from many anonymous ftp archives. The primary site is `ftp://ftp.gnu.org`, but it is often very busy. A list of sites that mirror the software on `ftp://ftp.gnu.org` is available by anonymous ftp from `ftp://ftp.gnu.org/pub/gnu/GNUinfo/FTP`.

  Octave requires approximately 1.4 GB of disk storage to unpack and compile from source (significantly less, 400 MB, if you don't compile with debugging symbols). To compile without debugging symbols try the command

  ```
  make CFLAGS=-O CXXFLAGS=-O LDFLAGS=
  ```

  instead of just `make`.

- If you encounter errors while compiling Octave, first see Section E.4 [Installation Problems], page 1190, for a list of known problems and if there is a workaround or solution for your problem. If not, see Appendix D [Trouble], page 1173, for information about how to report bugs.

- Once you have successfully compiled Octave, run `make install`.

  This will install a copy of Octave, its libraries, and its documentation in the destination directory. As distributed, Octave is installed in the following directories. In the table below, *prefix* defaults to `/usr/local`, *version* stands for the current version number of the interpreter, and *arch* is the type of computer on which Octave is installed (for example, '`i586-unknown-gnu`').

  *prefix*`/bin`
  : Octave and other binaries that people will want to run directly.

  *prefix*`/lib/octave-`*version*
  : Libraries like liboctave.a and liboctinterp.a.

  *prefix*`/include/octave-`*version*`/octave`
  : Include files distributed with Octave.

  *prefix*`/share`
  : Architecture-independent data files.

  *prefix*`/share/man/man1`
  : Unix-style man pages describing Octave.

*prefix*/share/info
: Info files describing Octave.

*prefix*/share/octave/*version*/m
: Function files distributed with Octave. This includes the Octave version, so that multiple versions of Octave may be installed at the same time.

*prefix*/libexec/octave/*version*/exec/*arch*
: Executables to be run by Octave rather than the user.

*prefix*/lib/octave/*version*/oct/*arch*
: Object files that will be dynamically loaded.

*prefix*/share/octave/*version*/imagelib
: Image files that are distributed with Octave.

## E.3 Compiling Octave with 64-bit Indexing

Note: the following only applies to systems that have 64-bit pointers. Configuring Octave with `--enable-64` cannot magically make a 32-bit system have a 64-bit address space.

On 64-bit systems, Octave uses 64-bit integers for indexing arrays by default. If the configure script determines that your BLAS library uses 32-bit integers, then operations using the following libraries are limited to arrays with dimensions that are smaller than $2^{31}$ elements:

- BLAS
- LAPACK
- QRUPDATE
- SuiteSparse
- SUNDIALS IDA
- ARPACK

Additionally, the following libraries use `int` internally, so maximum problem sizes are always limited:

- GLPK
- Qhull

Except for GLPK and Qhull, these libraries may also be configured to use 64-bit integers, but most systems do not provide packages built this way. If you wish to experiment with large arrays, the following information may be helpful.

To determine the integer size of the BLAS library used by Octave, the following code can be executed:

```
clear all;
N = 2^31;
## The following line requires about 8 GB of RAM!
a = b = ones (N, 1, "single");
c = a' * b
```

If the BLAS library uses 32-bit integers, an error will be thrown:

```
error: integer dimension or index out of range for Fortran
INTEGER type
```

Otherwise, if the BLAS library uses 64-bit integers, the result is:

```
c = 2^31 = 2147483648
```

Note that the test case above usually requires twice the memory, if $a$ and $b$ are not assigned by `a = b = ...`. Note further, that the data type `single` has a precision of about 23 binary bits. In this particular example no rounding errors occur.

Generally, it is best to have all of these libraries in versions that support 32-bit indexing, or all of these libraries must support 64-bit indexing. Mixing libraries with 64-bit indexing with libraries with 32-bit indexing can cause unpredictable behavior including program crashes with possible loss of data.

The following instructions were tested with the development version of Octave and GCC 4.3.4 on an x86_64 Debian system and may be out of date now. Please report any problems or corrections on the Octave bug tracker.

The versions listed below are the versions used for testing. If newer versions of these packages are available, you should try to use them, although there may be some differences.

All libraries and header files will be installed in subdirectories of `$prefix64` (you must choose the location of this directory).

- BLAS and LAPACK (http://www.netlib.org/lapack)

  Reference versions for both libraries are included in the reference LAPACK 3.2.1 distribution from http://www.netlib.org/.

  - Copy the file `make.inc.example` and name it `make.inc`. The options `-fdefault-integer-8` and `-fPIC` (on 64-bit CPU) have to be added to the variable `OPTS` and `NOOPT`.
  - Once you have compiled this library make sure that you use it for compiling Suite Sparse and Octave. In the following we assume that you installed the LAPACK library as $prefix64/lib/liblapack.a.

- QRUPDATE (https://sourceforge.net/projects/qrupdate)

  In the `Makeconf` file:

  - Add `-fdefault-integer-8` to `FFLAGS`.
  - Adjust the BLAS and LAPACK variables as needed if your 64-bit aware BLAS and LAPACK libraries are in a non-standard location.
  - Set `PREFIX` to the top-level directory of your install tree.
  - Run `make solib` to make a shared library.
  - Run `make install` to install the library.

- SuiteSparse (http://faculty.cse.tamu.edu/davis/suitesparse.html)

  Pass the following options to `make` to enable 64-bit integers for BLAS library calls. On 64-bit Windows systems, use `-DLONGBLAS="long long"` instead.

  ```
  CFLAGS='-DLONGBLAS=long'
  CXXFLAGS='-DLONGBLAS=long'
  ```

  The SuiteSparse makefiles don't generate shared libraries. On some systems, you can generate them by doing something as simple as

```
top=$(pwd)
for f in *.a; do
  mkdir tmp
  cd tmp
  ar vx ../$f
  gcc -shared -o ../${f%%.a}.so *.o
  cd $top
  rm -rf tmp
done
```

Other systems may require a different solution.

- SUNDIALS IDA (https://computing.llnl.gov/projects/sundials/ida)

  When configuring with `cmake` add the flag `-DSUNDIALS_INDEX_SIZE=64`.

- ARPACK (https://forge.scilab.org/index.php/p/arpack-ng/)
  - Add `-fdefault-integer-8` to `FFLAGS` when running configure.
  - Run `make` to build the library.
  - Run `make install` to install the library.
- ATLAS instead of reference BLAS and LAPACK

  Suggestions on how to compile ATLAS would be most welcome.
- GLPK
- Qhull (http://www.qhull.org)

  Both GLPK and Qhull use `int` internally so maximum problem sizes may be limited.
- Octave

  Octave's 64-bit index support is activated with the configure option `--enable-64`.

  ```
  ./configure \
    LD_LIBRARY_PATH="$prefix64/lib" \
    CPPFLAGS="-I$prefix64/include" LDFLAGS="-L$prefix64/lib" \
    --enable-64
  ```

  You must ensure that all Fortran sources, except those in the `liboctave/external/ranlib` directory, are compiled such that INTEGERS are 8-bytes wide. If you are using gfortran, the configure script should automatically set the Makefile variable `F77_INTEGER_8_FLAG` to `-fdefault-integer-8`. If you are using another compiler, you must set this variable yourself. You should NOT set this flag in `FFLAGS`, otherwise the files in `liboctave/external/ranlib` will be miscompiled.
- Other dependencies

  Probably nothing special needs to be done for the following dependencies. If you discover that something does need to be done, please submit a bug report.
  - pcre
  - zlib
  - hdf5
  - fftw3
  - cURL
  - GraphicsMagick++

- OpenGL
- freetype
- fontconfig
- fltk

## E.4 Installation Problems

This section contains a list of problems (and some apparent problems that don't really mean anything is wrong) that may show up during installation of Octave.

- On some SCO systems, `info` fails to compile if `HAVE_TERMIOS_H` is defined in `config.h`. Simply removing the definition from `info/config.h` should allow it to compile.
- If `configure` finds `dlopen`, `dlsym`, `dlclose`, and `dlerror`, but not the header file `dlfcn.h`, you need to find the source for the header file and install it in the directory `usr/include`. This is reportedly a problem with Slackware 3.1. For Linux/GNU systems, the source for `dlfcn.h` is in the `ldso` package.
- Building `.oct` files doesn't work.

  You should probably have a shared version of `libstdc++`. A patch is needed to build shared versions of version 2.7.2 of `libstdc++` on the HP-PA architecture. You can find the patch at `ftp://ftp.cygnus.com/pub/g++/libg++-2.7.2-hppa-gcc-fix`.
- On some DEC alpha systems there may be a problem with the `libdxml` library, resulting in floating point errors and/or segmentation faults in the linear algebra routines called by Octave. If you encounter such problems, then you should modify the configure script so that `SPECIAL_MATH_LIB` is not set to `-ldxml`.
- On FreeBSD systems Octave may hang while initializing some internal constants. The fix appears to be to use

  ```
  options      GPL_MATH_EMULATE
  ```

  rather than

  ```
  options      MATH_EMULATE
  ```

  in the kernel configuration files (typically found in the directory `/sys/i386/conf`). After making this change, you'll need to rebuild the kernel, install it, and reboot.
- If you encounter errors like

  ```
  passing ‘void (*)()' as argument 2 of
    ‘octave_set_signal_handler(int, void (*)(int))'
  ```

  or

  ```
  warning: ANSI C++ prohibits conversion from ‘(int)'
           to ‘(...)'
  ```

  while compiling `sighandlers.cc`, you may need to edit some files in the `gcc` include subdirectory to add proper prototypes for functions there. For example, Ultrix 4.2 needs proper declarations for the `signal` function and the `SIG_IGN` macro in the file `signal.h`.

  On some systems the `SIG_IGN` macro is defined to be something like this:

  ```
  #define  SIG_IGN  (void (*)())1
  ```

when it should really be something like:

```
#define  SIG_IGN  (void (*)(int))1
```

to match the prototype declaration for the `signal` function. This change should also be made for the `SIG_DFL` and `SIG_ERR` symbols. It may be necessary to change the definitions in `sys/signal.h` as well.

The `gcc fixincludes` and `fixproto` scripts should probably fix these problems when `gcc` installs its modified set of header files, but I don't think that's been done yet.

**You should not change the files in `/usr/include`.** You can find the `gcc` include directory tree by running the command

```
gcc -print-libgcc-file-name
```

The directory of `gcc` include files normally begins in the same directory that contains the file `libgcc.a`.

- Some of the Fortran subroutines may fail to compile with older versions of the Sun Fortran compiler. If you get errors like

```
zgemm.f:
        zgemm:
warning: unexpected parent of complex expression subtree
zgemm.f, line 245: warning: unexpected parent of complex
  expression subtree
warning: unexpected parent of complex expression subtree
zgemm.f, line 304: warning: unexpected parent of complex
  expression subtree
warning: unexpected parent of complex expression subtree
zgemm.f, line 327: warning: unexpected parent of complex
  expression subtree
pcc_binval: missing IR_CONV in complex op
make[2]: *** [zgemm.o] Error 1
```

  when compiling the Fortran subroutines in the `liboctave/external` subdirectory, you should either upgrade your compiler or try compiling with optimization turned off.

- On NeXT systems, if you get errors like this:

```
/usr/tmp/cc007458.s:unknown:Undefined local
      symbol LBB7656
/usr/tmp/cc007458.s:unknown:Undefined local
      symbol LBE7656
```

  when compiling `Array.cc` and `Matrix.cc`, try recompiling these files without `-g`.

- Some people have reported that calls to system() and the pager do not work on SunOS systems. This is apparently due to having `G_HAVE_SYS_WAIT` defined to be 0 instead of 1 when compiling `libg++`.

- On systems where the reference BLAS library is used the following matrix-by-vector multiplication incorrectly handles NaN values of the form `NaN * 0`.

```
[NaN, 1; 0, 0] * [0; 1]
⇒
[ 1
  0 ]

correct result ⇒
[ NaN
  0   ]
```

Install a different BLAS library such as OpenBLAS or ATLAS to correct this issue.

- On NeXT systems, linking to `libsys_s.a` may fail to resolve the following functions

  ```
  _tcgetattr
  _tcsetattr
  _tcflow
  ```

  which are part of `libposix.a`. Unfortunately, linking Octave with `-posix` results in the following undefined symbols.

  ```
  .destructors_used
  .constructors_used
  _objc_msgSend
  _NXGetDefaultValue
  _NXRegisterDefaults
  .objc_class_name_NXStringTable
  .objc_class_name_NXBundle
  ```

  One kluge around this problem is to extract `termios.o` from `libposix.a`, put it in Octave's `src` directory, and add it to the list of files to link together in the makefile. Suggestions for better ways to solve this problem are welcome!

- If Octave crashes immediately with a floating point exception, it is likely that it is failing to initialize the IEEE 754 floating point values for infinity and NaN.

  If your system actually does support IEEE 754 arithmetic, you should be able to fix this problem by modifying the function `octave_ieee_init` in the file `lo-ieee.cc` to correctly initialize Octave's internal infinity and NaN variables.

  If your system does not support IEEE 754 arithmetic but Octave's configure script incorrectly determined that it does, you can work around the problem by editing the file `config.h` to not define `HAVE_ISINF`, `HAVE_FINITE`, and `HAVE_ISNAN`.

  In any case, please report this as a bug since it might be possible to modify Octave's configuration script to automatically determine the proper thing to do.

- If Octave is unable to find a header file because it is installed in a location that is not normally searched by the compiler, you can add the directory to the include search path by specifying (for example) `CPPFLAGS=-I/some/nonstandard/directory` as an argument to `configure`. Other variables that can be specified this way are `CFLAGS`, `CXXFLAGS`, `FFLAGS`, and `LDFLAGS`. Passing them as options to the configure script also records them in the `config.status` file. By default, `CPPFLAGS` and `LDFLAGS` are empty, `CFLAGS` and `CXXFLAGS` are set to `"-g -O2"` and `FFLAGS` is set to `"-O"`.

# Appendix F Grammar and Parser

This appendix will eventually contain a semi-formal description of Octave's language.

## F.1 Keywords

The function `iskeyword` can be used to check whether an identifier is reserved by Octave.

*tf* = iskeyword (*name*)
*keyword_list* = iskeyword ()

Return true if *name* is an Octave keyword.

If *name* is omitted, return a list of keywords.

**See also:** [isvarname], page 145, [exist], page 152.

### Keyword Categories

[Variable Declaration], page 1193 | [Function Definition], page 1194 | [Control Statements], page 1195 | [Iterating Structures], page 1198 | [Classdef Structures], page 1199 | [Execution Environment], page 1200

### Alphabetical Keyword Listing

The identifiers below are keywords, and may not be used as variable or function names.

[__FILE__], page 1200 | [__LINE__], page 1201 | [arguments], page 1195 | [break], page 1199 | [case], page 1196 | [catch], page 1197 | [classdef], page 1199 | [continue], page 1199 | [do], page 1198 | [else], page 1196 | [elseif], page 1196 | [end], page 1194 | [end_try_catch], page 1197 | [end_unwind_protect], page 1198 | [endarguments], page 1195 | [endclassdef], page 1200 | [endenumeration], page 1200 | [endevents], page 1200 | [endfor], page 1198 | [endfunction], page 1194 | [endif], page 1196 | [endmethods], page 1200 | [endparfor], page 1199 | [endproperties], page 1200 | [endswitch], page 1197 | [endwhile], page 1198 | [enumeration], page 1200 | [events], page 1200 | [for], page 1198 | [function], page 1194 | [global], page 1193 | [if], page 1195 | [methods], page 1200 | [otherwise], page 1197 | [parfor], page 1199 | [persistent], page 1193 | [properties], page 1200 | [return], page 1195 | [switch], page 1196 | [try], page 1197 | [until], page 1199 | [unwind_protect], page 1197 | [unwind_protect_cleanup], page 1198 | [while], page 1198

### Variable Declaration

`global` *var*

Declare variables to have global scope.

```
global x;
if (isempty (x))
  x = 1;
endif
```

**See also:** [persistent], page 1193.

`persistent` *var*

Declare variables as persistent.

A variable that has been declared persistent within a function will retain its contents in memory between subsequent calls to the same function. The difference between persistent variables and global variables is that persistent variables are local in scope to a particular function and are not visible elsewhere.

**See also:** [global], page 1193.

## Function Definition

`function` [*output1*, ...] = *function_name* (*input1*, ...)
`function` *function_name* (*input1*, ...)
`function` *output* = *function_name*

Begin a function body with name `function_name`, with *outputs* as results, and with *inputs* as parameters.

The function can later be invoked in Octave using the syntax

```
[output1, output2, ...] = function_name (input1, input2, ...)
```

**See also:** [return], page 1195.

`endfunction`

Terminate a function definition.

**See also:** [function], page 1194.

`end`

The termination of a code block *or* the last element of an array.

The keyword `end` is used to terminate blocks of code begun with `for`, `parfor`, `if`, `do`, `while`, `function`, `switch`, `try`, or `unwind_protect`.

When using classdef (object-oriented) programming, the keyword terminates blocks of code begun with `classdef`, `properties`, `methods`, `events`, or `enumeration`.

As an index of an array, the magic index `"end"` refers to the last valid entry in an indexing operation.

Examples:

```
x = [ 1 2 3; 4 5 6 ];
x(1,end)
 ⇒ 3
x(end,1)
 ⇒ 4
x(end,end)
 ⇒ 6
```

Programming Notes:

1. The `end` keyword cannot be used within `subsref`, `subsasgn`, or `substruct` for manual indexing operations.
2. For custom classes, enabling use of `end` in indexing expressions requires writing an overloaded function such as:

```
function last_index = end (obj, end_dim, ndim_obj)
  if (end_dim == ndim_obj)
    last_index = prod (size (obj)(end_dim:ndim_obj));
  else
    last_index = size (obj, end_dim);
  endif
endfunction
```

For more information, see Chapter 34 [Object Oriented Programming], page 973.

**See also:** [for], page 1198, [parfor], page 1199, [if], page 1195, [do], page 1198, [while], page 1198, [function], page 1194, [switch], page 1196, [try], page 1197, [unwind_protect], page 1197.

`return`
: Return execution control immediately from a function or script to the calling code.

  `return` is used to stop executing code and exit an m-file immediately rather than continuing until the end of the function or script is reached.

  If the function or script was invoked directly, rather than from calling code in an m-file, then Octave returns to the command line.

  **See also:** [function], page 1194.

`arguments`
: Begin a function arguments block.

  **Warning:** Function arguments validation blocks are not yet implemented in Octave. Use of the keyword `arguments` will cause the parser to emit a warning, and may throw an error depending on the contents of the code in the block.

`endarguments`
: Terminate a function arguments block.

  **Warning:** Function arguments validation blocks are not yet implemented in Octave. Use of the keyword `arguments` will cause the parser to emit a warning, and may throw an error depending on the contents of the code in the block.

  **See also:** [arguments], page 1195.

## Control Statements

`if` (*`cond`*) ... *endif*
`if` (*`cond`*) ... *else* ... *endif*
`if` (*`cond`*) ... *elseif* (*`cond`*) ... *endif*
`if` (*`cond`*) ... *elseif* (*`cond`*) ... *else* ... *endif*
: Begin an if block.

  The conditional *cond* is true if it is not empty and if *all* values are nonzero.

```
x = 1;
if (x == 1)
  disp ("one");
elseif (x == 2)
  disp ("two");
else
  disp ("not one or two");
endif
```

**See also:** [switch], page 1196.

`else`

Alternate action for an if block.

See `if` for an example.

**See also:** [if], page 1195.

`elseif (`*cond*`)`

Alternate conditional test for an if block.

The conditional *cond* is true if it is not empty and if *all* values are nonzero.

See `if` for an example.

**See also:** [if], page 1195.

`endif`

Terminate an if block.

See `if` for an example.

**See also:** [if], page 1195.

`switch` *statement*

Begin a switch block.

```
yesno = "yes";

switch (yesno)
  case {"Yes" "yes" "YES" "y" "Y"}
    value = 1;
  case {"No" "no" "NO" "n" "N"}
    value = 0;
  otherwise
    error ("invalid value");
endswitch
```

**See also:** [if], page 1195, [case], page 1196, [otherwise], page 1197.

`case` *value*
`case` {*value1*, ...}

A case statement in a switch block.

The code associated with a `case` statement is executed if one of the tests below succeeds given a `switch` argument *statement* and a *value* of type:

- number : *statement* `==` *value*

- character array or string : `strcmp (`*`statement`*`, `*`value`*`)`
- object : `eq (`*`statement`*`, `*`value`*`)`
  (programmer must write overloaded `eq` function which performs '`==`')
- cell array : One of the *values* in the cell array must match *statement* using the criteria above.

Programming Notes: Octave cases are exclusive and do not fall-through as do C-language cases. A `switch` statement must have at least one `case`.

See `switch` for an example.

**See also:** [switch], page 1196.

`otherwise`

The default statement in a `switch` block which is executed when no other `case` statements match the input.

**See also:** [switch], page 1196, [case], page 1196.

`endswitch`

Terminate a switch block.

See `switch` for an example.

**See also:** [switch], page 1196.

`try`

Begin a try-catch block.

If an error occurs within a `try` block, then the code in the `catch` block will be run. Execution will proceed after the `catch` block (though it is often recommended to use the `lasterr` function to re-throw the error after cleanup is completed).

**See also:** [catch], page 1197, [unwind_protect], page 1197.

`catch`
`catch` *`exception_var`*

Begin the cleanup part of a try-catch block.

If *exception_var* is specified then a variable of that name will defined in the scope of the `catch` block. The variable is a scalar struct with fields `identifier`, `message`, and `stack` which describes the error that was caught.

**See also:** [try], page 1197.

`end_try_catch`

Terminate a try-catch block.

**See also:** [try], page 1197, [catch], page 1197.

`unwind_protect`

Begin an unwind_protect block.

If an error occurs within the first part of an `unwind_protect` block the commands within the `unwind_protect_cleanup` block are executed before the error is thrown. If an error is not thrown, then the `unwind_protect_cleanup` block is still executed. In

other words, the `unwind_protect_cleanup` code is guaranteed to execute regardless of success or failure in the `unwind_protect` block.

**See also:** [unwind_protect_cleanup], page 1198, [try], page 1197.

`unwind_protect_cleanup`
Begin the cleanup section of an `unwind_protect` block.

**See also:** [unwind_protect], page 1197.

`end_unwind_protect`
Terminate an unwind_protect block.

**See also:** [unwind_protect], page 1197.

## Iterating Structures

`for` *i* = *range*
Begin a for loop.

```
for i = 1:10
  i
endfor
```

**See also:** [parfor], page 1199, [do], page 1198, [while], page 1198.

`endfor`
Terminate a for loop.

See `for` for an example.

**See also:** [for], page 1198.

`while` (*cond*)
Begin a while loop.

The conditional *cond* is true if it is not empty and if *all* values are nonzero.

```
i = 0;
while (i < 10)
  i++
endwhile
```

**See also:** [do], page 1198, [endwhile], page 1198, [for], page 1198, [until], page 1199.

`endwhile`
Terminate a while loop.

See `while` for an example.

**See also:** [do], page 1198, [while], page 1198.

`do`
Begin a do-until loop.

This differs from a `while` loop in that the body of the loop is executed at least once.

```
i = 0;
do
  i++
until (i == 10)
```

**See also:** [for], page 1198, [until], page 1199, [while], page 1198.

`until (`*cond*`)`
End a do-until loop.

The conditional *cond* is true if it is not empty and if *all* values are nonzero.

See `do` for an example.

**See also:** [do], page 1198.

`parfor` *i* = *range*
`parfor (`*i* = *range*, *maxproc*`)`
Begin a for loop that may execute in parallel.

A `parfor` loop has the same syntax as a `for` loop. If your Octave session has a parallel processing pool enabled, the iterations of the `parfor` loop will be executed in parallel across the pool's workers. Otherwise, `parfor` will behave exactly as `for`.

When operating in parallel mode, a `parfor` loop's iterations are not guaranteed to occur sequentially, and there are additional restrictions about the data access operations you can do inside the loop body.

**Warning:** parallel processing pools are currently unimplemented in Octave; `parfor` currently behaves exactly as a normal `for` loop.

```
parfor i = 1:10
  i
endparfor
```

**See also:** [for], page 1198, [do], page 1198, [while], page 1198.

`endparfor`
Terminate a parfor loop.

See `parfor` for an example.

**See also:** [parfor], page 1199.

`break`
Exit the innermost enclosing `do`, `while`, or `for` loop.

**See also:** [continue], page 1199, [do], page 1198, [while], page 1198, [for], page 1198, [parfor], page 1199.

`continue`
Jump to the end of the innermost enclosing `do`, `while`, or `for` loop.

**See also:** [break], page 1199, [do], page 1198, [while], page 1198, [for], page 1198, [parfor], page 1199.

## Classdef Structures

`classdef`
Begin a classdef block.

**See also:** [properties], page 1200, [methods], page 1200, [events], page 1200, [enumeration], page 1200.

`endclassdef`
: Terminate a classdef definition.

  **See also:** [classdef], page 1199.

`properties`
: Mark the beginning of a block of properties in a classdef definition.

  Note that the [function `"properties"`], page 994, is not a keyword, but rather an ordinary function that lists the properties of a classdef class or object.

  **See also:** [endproperties], page 1200.

`endproperties`
: Terminate a properties block in a classdef definition.

  **See also:** [properties], page 1200.

`methods`
: Mark the beginning of a block of methods in a classdef definition.

  Note that the [function `"methods"`], page 974, is not a keyword, but rather an ordinary function that lists the methods of a class or object.

  **See also:** [endmethods], page 1200.

`endmethods`
: Terminate a methods block in a classdef definition.

  **See also:** [methods], page 1200.

`events`
: Begin an events block in a classdef definition.

`endevents`
: Terminate an events block in a classdef definition.

  **See also:** [events], page 1200.

`enumeration`
: Begin an enumeration block in a classdef definition.

`endenumeration`
: Terminate an enumeration block in a classdef definition.

  **See also:** [enumeration], page 1200.

## Execution Environment

`__FILE__`
: When the lexer recognizes the `"__FILE__"` keyword it returns a character array containing the full name and path of the file that is being executed. `"__FILE__"` will return `"stdin"` if called from the command line.

  **See also:** [__LINE__], page 1201.

`__LINE__`
: When the lexer recognizes the `"__LINE__"` keyword it returns a numeric value containing the current input line number of the function or file being executed. `"__LINE__"` will return `1` if called from the command line.

  **See also:** [__FILE__], page 1200.

## F.2 Parser

The parser has a number of variables that affect its internal operation. These variables are generally documented in the manual alongside the code that they affect.

In addition, there are three non-specific parser customization functions. `add_input_event_hook` can be used to schedule a user function for periodic evaluation. `remove_input_event_hook` will stop a user function from being evaluated periodically.

*id* = add_input_event_hook (*fcn*)
*id* = add_input_event_hook (*fcn*, *data*)
: Add the named function or function handle *fcn* to the list of functions to call periodically when Octave is waiting for input.

  The function should have the form

  ```
  fcn (data)
  ```

  If *data* is omitted, Octave calls the function without any arguments.

  The returned identifier may be used to remove the function handle from the list of input hook functions.

  **See also:** [remove_input_event_hook], page 1201.

remove_input_event_hook (*name*)
remove_input_event_hook (*fcn_id*)
: Remove the named function or function handle with the given identifier from the list of functions to call periodically when Octave is waiting for input.

  **See also:** [add_input_event_hook], page 1201.

Finally, when the parser cannot identify an input token it calls a particular function to handle this. By default, this is the internal function `"__unimplemented__"` which makes suggestions about possible Octave substitutes for MATLAB functions.

*val* = missing_function_hook ()
*old_val* = missing_function_hook (*new_val*)
*old_val* = missing_function_hook (*new_val*, "*local*")
: Query or set the internal variable that specifies the function to call to provide extra information when an unknown identifier is referenced.

  When called from inside a function with the `"local"` option, the variable is changed locally for the function and any subroutines it calls. The original variable value is restored when exiting the function.

  **See also:** [missing_component_hook], page 1095.

# Appendix G GNU GENERAL PUBLIC LICENSE

Version 3, 29 June 2007

Copyright © 2007 Free Software Foundation, Inc. https://fsf.org/

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

## Preamble

The GNU General Public License is a free, copyleft license for software and other kinds of works.

The licenses for most software and other practical works are designed to take away your freedom to share and change the works. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change all versions of a program—to make sure it remains free software for all its users. We, the Free Software Foundation, use the GNU General Public License for most of our software; it applies also to any other work released this way by its authors. You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for them if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs, and that you know you can do these things.

To protect your rights, we need to prevent others from denying you these rights or asking you to surrender the rights. Therefore, you have certain responsibilities if you distribute copies of the software, or if you modify it: responsibilities to respect the freedom of others.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must pass on to the recipients the same freedoms that you received. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

Developers that use the GNU GPL protect your rights with two steps: (1) assert copyright on the software, and (2) offer you this License giving you legal permission to copy, distribute and/or modify it.

For the developers' and authors' protection, the GPL clearly explains that there is no warranty for this free software. For both users' and authors' sake, the GPL requires that modified versions be marked as changed, so that their problems will not be attributed erroneously to authors of previous versions.

Some devices are designed to deny users access to install or run modified versions of the software inside them, although the manufacturer can do so. This is fundamentally incompatible with the aim of protecting users' freedom to change the software. The systematic pattern of such abuse occurs in the area of products for individuals to use, which is precisely where it is most unacceptable. Therefore, we have designed this version of the GPL to prohibit the practice for those products. If such problems arise substantially in other domains, we stand ready to extend this provision to those domains in future versions of the GPL, as needed to protect the freedom of users.

Finally, every program is threatened constantly by software patents. States should not allow patents to restrict development and use of software on general-purpose computers, but in those that do, we wish to avoid the special danger that patents applied to a free program could make it effectively proprietary. To prevent this, the GPL assures that patents cannot be used to render the program non-free.

The precise terms and conditions for copying, distribution and modification follow.

## TERMS AND CONDITIONS

0. Definitions.

   "This License" refers to version 3 of the GNU General Public License.

   "Copyright" also means copyright-like laws that apply to other kinds of works, such as semiconductor masks.

   "The Program" refers to any copyrightable work licensed under this License. Each licensee is addressed as "you". "Licensees" and "recipients" may be individuals or organizations.

   To "modify" a work means to copy from or adapt all or part of the work in a fashion requiring copyright permission, other than the making of an exact copy. The resulting work is called a "modified version" of the earlier work or a work "based on" the earlier work.

   A "covered work" means either the unmodified Program or a work based on the Program.

   To "propagate" a work means to do anything with it that, without permission, would make you directly or secondarily liable for infringement under applicable copyright law, except executing it on a computer or modifying a private copy. Propagation includes copying, distribution (with or without modification), making available to the public, and in some countries other activities as well.

   To "convey" a work means any kind of propagation that enables other parties to make or receive copies. Mere interaction with a user through a computer network, with no transfer of a copy, is not conveying.

   An interactive user interface displays "Appropriate Legal Notices" to the extent that it includes a convenient and prominently visible feature that (1) displays an appropriate copyright notice, and (2) tells the user that there is no warranty for the work (except to the extent that warranties are provided), that licensees may convey the work under this License, and how to view a copy of this License. If the interface presents a list of user commands or options, such as a menu, a prominent item in the list meets this criterion.

1. Source Code.

   The "source code" for a work means the preferred form of the work for making modifications to it. "Object code" means any non-source form of a work.

   A "Standard Interface" means an interface that either is an official standard defined by a recognized standards body, or, in the case of interfaces specified for a particular programming language, one that is widely used among developers working in that language.

The "System Libraries" of an executable work include anything, other than the work as a whole, that (a) is included in the normal form of packaging a Major Component, but which is not part of that Major Component, and (b) serves only to enable use of the work with that Major Component, or to implement a Standard Interface for which an implementation is available to the public in source code form. A "Major Component", in this context, means a major essential component (kernel, window system, and so on) of the specific operating system (if any) on which the executable work runs, or a compiler used to produce the work, or an object code interpreter used to run it.

The "Corresponding Source" for a work in object code form means all the source code needed to generate, install, and (for an executable work) run the object code and to modify the work, including scripts to control those activities. However, it does not include the work's System Libraries, or general-purpose tools or generally available free programs which are used unmodified in performing those activities but which are not part of the work. For example, Corresponding Source includes interface definition files associated with source files for the work, and the source code for shared libraries and dynamically linked subprograms that the work is specifically designed to require, such as by intimate data communication or control flow between those subprograms and other parts of the work.

The Corresponding Source need not include anything that users can regenerate automatically from other parts of the Corresponding Source.

The Corresponding Source for a work in source code form is that same work.

2. Basic Permissions.

   All rights granted under this License are granted for the term of copyright on the Program, and are irrevocable provided the stated conditions are met. This License explicitly affirms your unlimited permission to run the unmodified Program. The output from running a covered work is covered by this License only if the output, given its content, constitutes a covered work. This License acknowledges your rights of fair use or other equivalent, as provided by copyright law.

   You may make, run and propagate covered works that you do not convey, without conditions so long as your license otherwise remains in force. You may convey covered works to others for the sole purpose of having them make modifications exclusively for you, or provide you with facilities for running those works, provided that you comply with the terms of this License in conveying all material for which you do not control copyright. Those thus making or running the covered works for you must do so exclusively on your behalf, under your direction and control, on terms that prohibit them from making any copies of your copyrighted material outside their relationship with you.

   Conveying under any other circumstances is permitted solely under the conditions stated below. Sublicensing is not allowed; section 10 makes it unnecessary.

3. Protecting Users' Legal Rights From Anti-Circumvention Law.

   No covered work shall be deemed part of an effective technological measure under any applicable law fulfilling obligations under article 11 of the WIPO copyright treaty adopted on 20 December 1996, or similar laws prohibiting or restricting circumvention of such measures.

When you convey a covered work, you waive any legal power to forbid circumvention of technological measures to the extent such circumvention is effected by exercising rights under this License with respect to the covered work, and you disclaim any intention to limit operation or modification of the work as a means of enforcing, against the work's users, your or third parties' legal rights to forbid circumvention of technological measures.

4. Conveying Verbatim Copies.

   You may convey verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice; keep intact all notices stating that this License and any non-permissive terms added in accord with section 7 apply to the code; keep intact all notices of the absence of any warranty; and give all recipients a copy of this License along with the Program.

   You may charge any price or no price for each copy that you convey, and you may offer support or warranty protection for a fee.

5. Conveying Modified Source Versions.

   You may convey a work based on the Program, or the modifications to produce it from the Program, in the form of source code under the terms of section 4, provided that you also meet all of these conditions:

   a. The work must carry prominent notices stating that you modified it, and giving a relevant date.

   b. The work must carry prominent notices stating that it is released under this License and any conditions added under section 7. This requirement modifies the requirement in section 4 to "keep intact all notices".

   c. You must license the entire work, as a whole, under this License to anyone who comes into possession of a copy. This License will therefore apply, along with any applicable section 7 additional terms, to the whole of the work, and all its parts, regardless of how they are packaged. This License gives no permission to license the work in any other way, but it does not invalidate such permission if you have separately received it.

   d. If the work has interactive user interfaces, each must display Appropriate Legal Notices; however, if the Program has interactive interfaces that do not display Appropriate Legal Notices, your work need not make them do so.

   A compilation of a covered work with other separate and independent works, which are not by their nature extensions of the covered work, and which are not combined with it such as to form a larger program, in or on a volume of a storage or distribution medium, is called an "aggregate" if the compilation and its resulting copyright are not used to limit the access or legal rights of the compilation's users beyond what the individual works permit. Inclusion of a covered work in an aggregate does not cause this License to apply to the other parts of the aggregate.

6. Conveying Non-Source Forms.

   You may convey a covered work in object code form under the terms of sections 4 and 5, provided that you also convey the machine-readable Corresponding Source under the terms of this License, in one of these ways:

a. Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by the Corresponding Source fixed on a durable physical medium customarily used for software interchange.

b. Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by a written offer, valid for at least three years and valid for as long as you offer spare parts or customer support for that product model, to give anyone who possesses the object code either (1) a copy of the Corresponding Source for all the software in the product that is covered by this License, on a durable physical medium customarily used for software interchange, for a price no more than your reasonable cost of physically performing this conveying of source, or (2) access to copy the Corresponding Source from a network server at no charge.

c. Convey individual copies of the object code with a copy of the written offer to provide the Corresponding Source. This alternative is allowed only occasionally and noncommercially, and only if you received the object code with such an offer, in accord with subsection 6b.

d. Convey the object code by offering access from a designated place (gratis or for a charge), and offer equivalent access to the Corresponding Source in the same way through the same place at no further charge. You need not require recipients to copy the Corresponding Source along with the object code. If the place to copy the object code is a network server, the Corresponding Source may be on a different server (operated by you or a third party) that supports equivalent copying facilities, provided you maintain clear directions next to the object code saying where to find the Corresponding Source. Regardless of what server hosts the Corresponding Source, you remain obligated to ensure that it is available for as long as needed to satisfy these requirements.

e. Convey the object code using peer-to-peer transmission, provided you inform other peers where the object code and Corresponding Source of the work are being offered to the general public at no charge under subsection 6d.

A separable portion of the object code, whose source code is excluded from the Corresponding Source as a System Library, need not be included in conveying the object code work.

A "User Product" is either (1) a "consumer product", which means any tangible personal property which is normally used for personal, family, or household purposes, or (2) anything designed or sold for incorporation into a dwelling. In determining whether a product is a consumer product, doubtful cases shall be resolved in favor of coverage. For a particular product received by a particular user, "normally used" refers to a typical or common use of that class of product, regardless of the status of the particular user or of the way in which the particular user actually uses, or expects or is expected to use, the product. A product is a consumer product regardless of whether the product has substantial commercial, industrial or non-consumer uses, unless such uses represent the only significant mode of use of the product.

"Installation Information" for a User Product means any methods, procedures, authorization keys, or other information required to install and execute modified versions of a covered work in that User Product from a modified version of its Corresponding Source.

The information must suffice to ensure that the continued functioning of the modified object code is in no case prevented or interfered with solely because modification has been made.

If you convey an object code work under this section in, or with, or specifically for use in, a User Product, and the conveying occurs as part of a transaction in which the right of possession and use of the User Product is transferred to the recipient in perpetuity or for a fixed term (regardless of how the transaction is characterized), the Corresponding Source conveyed under this section must be accompanied by the Installation Information. But this requirement does not apply if neither you nor any third party retains the ability to install modified object code on the User Product (for example, the work has been installed in ROM).

The requirement to provide Installation Information does not include a requirement to continue to provide support service, warranty, or updates for a work that has been modified or installed by the recipient, or for the User Product in which it has been modified or installed. Access to a network may be denied when the modification itself materially and adversely affects the operation of the network or violates the rules and protocols for communication across the network.

Corresponding Source conveyed, and Installation Information provided, in accord with this section must be in a format that is publicly documented (and with an implementation available to the public in source code form), and must require no special password or key for unpacking, reading or copying.

7. Additional Terms.

   "Additional permissions" are terms that supplement the terms of this License by making exceptions from one or more of its conditions. Additional permissions that are applicable to the entire Program shall be treated as though they were included in this License, to the extent that they are valid under applicable law. If additional permissions apply only to part of the Program, that part may be used separately under those permissions, but the entire Program remains governed by this License without regard to the additional permissions.

   When you convey a copy of a covered work, you may at your option remove any additional permissions from that copy, or from any part of it. (Additional permissions may be written to require their own removal in certain cases when you modify the work.) You may place additional permissions on material, added by you to a covered work, for which you have or can give appropriate copyright permission.

   Notwithstanding any other provision of this License, for material you add to a covered work, you may (if authorized by the copyright holders of that material) supplement the terms of this License with terms:

   a. Disclaiming warranty or limiting liability differently from the terms of sections 15 and 16 of this License; or
   b. Requiring preservation of specified reasonable legal notices or author attributions in that material or in the Appropriate Legal Notices displayed by works containing it; or
   c. Prohibiting misrepresentation of the origin of that material, or requiring that modified versions of such material be marked in reasonable ways as different from the original version; or

d. Limiting the use for publicity purposes of names of licensors or authors of the material; or

e. Declining to grant rights under trademark law for use of some trade names, trademarks, or service marks; or

f. Requiring indemnification of licensors and authors of that material by anyone who conveys the material (or modified versions of it) with contractual assumptions of liability to the recipient, for any liability that these contractual assumptions directly impose on those licensors and authors.

All other non-permissive additional terms are considered "further restrictions" within the meaning of section 10. If the Program as you received it, or any part of it, contains a notice stating that it is governed by this License along with a term that is a further restriction, you may remove that term. If a license document contains a further restriction but permits relicensing or conveying under this License, you may add to a covered work material governed by the terms of that license document, provided that the further restriction does not survive such relicensing or conveying.

If you add terms to a covered work in accord with this section, you must place, in the relevant source files, a statement of the additional terms that apply to those files, or a notice indicating where to find the applicable terms.

Additional terms, permissive or non-permissive, may be stated in the form of a separately written license, or stated as exceptions; the above requirements apply either way.

8. Termination.

You may not propagate or modify a covered work except as expressly provided under this License. Any attempt otherwise to propagate or modify it is void, and will automatically terminate your rights under this License (including any patent licenses granted under the third paragraph of section 11).

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, you do not qualify to receive new licenses for the same material under section 10.

9. Acceptance Not Required for Having Copies.

You are not required to accept this License in order to receive or run a copy of the Program. Ancillary propagation of a covered work occurring solely as a consequence of using peer-to-peer transmission to receive a copy likewise does not require acceptance.

However, nothing other than this License grants you permission to propagate or modify any covered work. These actions infringe copyright if you do not accept this License. Therefore, by modifying or propagating a covered work, you indicate your acceptance of this License to do so.

10. Automatic Licensing of Downstream Recipients.

    Each time you convey a covered work, the recipient automatically receives a license from the original licensors, to run, modify and propagate that work, subject to this License. You are not responsible for enforcing compliance by third parties with this License.

    An "entity transaction" is a transaction transferring control of an organization, or substantially all assets of one, or subdividing an organization, or merging organizations. If propagation of a covered work results from an entity transaction, each party to that transaction who receives a copy of the work also receives whatever licenses to the work the party's predecessor in interest had or could give under the previous paragraph, plus a right to possession of the Corresponding Source of the work from the predecessor in interest, if the predecessor has it or can get it with reasonable efforts.

    You may not impose any further restrictions on the exercise of the rights granted or affirmed under this License. For example, you may not impose a license fee, royalty, or other charge for exercise of rights granted under this License, and you may not initiate litigation (including a cross-claim or counterclaim in a lawsuit) alleging that any patent claim is infringed by making, using, selling, offering for sale, or importing the Program or any portion of it.

11. Patents.

    A "contributor" is a copyright holder who authorizes use under this License of the Program or a work on which the Program is based. The work thus licensed is called the contributor's "contributor version".

    A contributor's "essential patent claims" are all patent claims owned or controlled by the contributor, whether already acquired or hereafter acquired, that would be infringed by some manner, permitted by this License, of making, using, or selling its contributor version, but do not include claims that would be infringed only as a consequence of further modification of the contributor version. For purposes of this definition, "control" includes the right to grant patent sublicenses in a manner consistent with the requirements of this License.

    Each contributor grants you a non-exclusive, worldwide, royalty-free patent license under the contributor's essential patent claims, to make, use, sell, offer for sale, import and otherwise run, modify and propagate the contents of its contributor version.

    In the following three paragraphs, a "patent license" is any express agreement or commitment, however denominated, not to enforce a patent (such as an express permission to practice a patent or covenant not to sue for patent infringement). To "grant" such a patent license to a party means to make such an agreement or commitment not to enforce a patent against the party.

    If you convey a covered work, knowingly relying on a patent license, and the Corresponding Source of the work is not available for anyone to copy, free of charge and under the terms of this License, through a publicly available network server or other readily accessible means, then you must either (1) cause the Corresponding Source to be so

available, or (2) arrange to deprive yourself of the benefit of the patent license for this particular work, or (3) arrange, in a manner consistent with the requirements of this License, to extend the patent license to downstream recipients. "Knowingly relying" means you have actual knowledge that, but for the patent license, your conveying the covered work in a country, or your recipient's use of the covered work in a country, would infringe one or more identifiable patents in that country that you have reason to believe are valid.

If, pursuant to or in connection with a single transaction or arrangement, you convey, or propagate by procuring conveyance of, a covered work, and grant a patent license to some of the parties receiving the covered work authorizing them to use, propagate, modify or convey a specific copy of the covered work, then the patent license you grant is automatically extended to all recipients of the covered work and works based on it.

A patent license is "discriminatory" if it does not include within the scope of its coverage, prohibits the exercise of, or is conditioned on the non-exercise of one or more of the rights that are specifically granted under this License. You may not convey a covered work if you are a party to an arrangement with a third party that is in the business of distributing software, under which you make payment to the third party based on the extent of your activity of conveying the work, and under which the third party grants, to any of the parties who would receive the covered work from you, a discriminatory patent license (a) in connection with copies of the covered work conveyed by you (or copies made from those copies), or (b) primarily for and in connection with specific products or compilations that contain the covered work, unless you entered into that arrangement, or that patent license was granted, prior to 28 March 2007.

Nothing in this License shall be construed as excluding or limiting any implied license or other defenses to infringement that may otherwise be available to you under applicable patent law.

12. No Surrender of Others' Freedom.

    If conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot convey a covered work so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not convey it at all. For example, if you agree to terms that obligate you to collect a royalty for further conveying from those to whom you convey the Program, the only way you could satisfy both those terms and this License would be to refrain entirely from conveying the Program.

13. Use with the GNU Affero General Public License.

    Notwithstanding any other provision of this License, you have permission to link or combine any covered work with a work licensed under version 3 of the GNU Affero General Public License into a single combined work, and to convey the resulting work. The terms of this License will continue to apply to the part which is the covered work, but the special requirements of the GNU Affero General Public License, section 13, concerning interaction through a network will apply to the combination as such.

14. Revised Versions of this License.

The Free Software Foundation may publish revised and/or new versions of the GNU General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies that a certain numbered version of the GNU General Public License "or any later version" applies to it, you have the option of following the terms and conditions either of that numbered version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of the GNU General Public License, you may choose any version ever published by the Free Software Foundation.

If the Program specifies that a proxy can decide which future versions of the GNU General Public License can be used, that proxy's public statement of acceptance of a version permanently authorizes you to choose that version for the Program.

Later license versions may give you additional or different permissions. However, no additional obligations are imposed on any author or copyright holder as a result of your choosing to follow a later version.

15. Disclaimer of Warranty.

    THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. Limitation of Liability.

    IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MODIFIES AND/OR CONVEYS THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

17. Interpretation of Sections 15 and 16.

    If the disclaimer of warranty and limitation of liability provided above cannot be given local legal effect according to their terms, reviewing courts shall apply local law that most closely approximates an absolute waiver of all civil liability in connection with the Program, unless a warranty or assumption of liability accompanies a copy of the Program in return for a fee.

# END OF TERMS AND CONDITIONS

## How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively state the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
one line to give the program's name and a brief idea of what it does.
Copyright (C) year name of author

This program is free software: you can redistribute it and/or modify
it under the terms of the GNU General Public License as published by
the Free Software Foundation, either version 3 of the License, or (at
your option) any later version.

This program is distributed in the hope that it will be useful, but
WITHOUT ANY WARRANTY; without even the implied warranty of
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the GNU
General Public License for more details.

You should have received a copy of the GNU General Public License
along with this program.  If not, see https://www.gnu.org/licenses/.
```

Also add information on how to contact you by electronic and paper mail.

If the program does terminal interaction, make it output a short notice like this when it starts in an interactive mode:

```
program Copyright (C) year name of author
This program comes with ABSOLUTELY NO WARRANTY; for details type ‘show w’.
This is free software, and you are welcome to redistribute it
under certain conditions; type ‘show c’ for details.
```

The hypothetical commands ‘`show w`’ and ‘`show c`’ should show the appropriate parts of the General Public License. Of course, your program's commands might be different; for a GUI interface, you would use an "about box".

You should also get your employer (if you work as a programmer) or school, if any, to sign a "copyright disclaimer" for the program, if necessary. For more information on this, and how to apply and follow the GNU GPL, see https://www.gnu.org/licenses/.

The GNU General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Lesser General Public License instead of this License. But first, please read https://www.gnu.org/licenses/why-not-lgpl.html.

# Concept Index


## A

acknowledgements ... 1
addition ... 170, 984
and operator ... 176, 984
anonymous functions ... 248
`ans` ... 145
answers, incorrect ... 1173, 1175
application-defined data ... 548
apply ... 695
area series ... 554
arguments in function call ... 167
arithmetic operators ... 170, 984
array, creating a Java array ... 1143
assignment expressions ... 179
assignment operators ... 179
axes graphics object ... 450

## B

bar series ... 554
batch processing ... 37
block comments ... 39
body of a loop ... 197
boolean expressions ... 176, 984
boolean operators ... 176, 984
`break` statement ... 200
broadcast ... 691
broadcasting ... 691
BSX ... 691
bug criteria ... 1173
bug tracker ... 1174
bugs ... 1173
bugs, investigating ... 1175
bugs, known ... 1173
bugs, reporting ... 1174
built-in data types ... 41

## C

callbacks ... 546
calling Java from Octave ... 1138
calling Octave from Java ... 1138
`case` statement ... 195
`catch` ... 202
cell arrays ... 46, 131
Chained indexing ... 162
character strings ... 46, 75
Cholesky factorization ... 657
Citations ... 5
Citing Octave ... 5
classes, making available to Octave ... 1138
classpath, adding new path ... 1146
classpath, difference between static and dynamic ... 1138
classpath, displaying ... 1145
classpath, dynamic ... 1145, 1146
classpath, removing path ... 1146
classpath, setting ... 1138
classpath, static ... 1145
`classpath.txt` ... 1138
clearing the screen ... 27
code profiling ... 280
colors, graphics ... 545
comma-separated lists ... 141
command and output logs ... 35
command completion ... 28
command descriptions ... 13
command echoing ... 35
command history ... 29
command options ... 15
command-line editing ... 26
commands functions ... 250
comments ... 39
comparison expressions ... 174, 984
complex-conjugate transpose ... 170, 984
Component Indexing ... 159
containers ... 117
continuation lines ... 203
`continue` statement ... 201
contour series ... 555
contributing to Octave ... 6
contributors ... 1
conversion specifications (`printf`) ... 316
conversion specifications (`scanf`) ... 322
copy-on-write ... 704
copyright ... 1203
core dump ... 1173
COW ... 704
creating graphics objects ... 450
cs-lists ... 141
customizing `readline` ... 33
customizing the prompt ... 33

## D

DAE ... 791
data sources in object groups ... 553
data structures ... 46, 117
data types ... 41
data types, built-in ... 41
data types, user-defined ... 46
decrement operator ... 183
default arguments ... 216
default graphics properties ... 543
defining functions ... 205
Degree Symbol ... 434
deprecated functions ... 1165
description format ... 12

diagonal and permutation matrices ..... 717
diagonal matrix expressions ..... 720
dialog, displaying a dialog for selecting directories ..... 999
dialog, displaying a dialog for selecting files ..... 999
dialog, displaying a dialog for storing files ..... 1000
dialog, displaying a font selection dialog ..... 1006
dialog, displaying a help dialog ..... 1001
dialog, displaying a list dialog ..... 1002
dialog, displaying a message dialog ..... 1003
dialog, displaying a modal dialog ..... 1006
dialog, displaying a question dialog ..... 1004
dialog, displaying a warning dialog ..... 1005
dialog, displaying an error dialog ..... 1001
dialog, displaying an input dialog ..... 1002
diary of commands and output ..... 35
differential equations ..... 791
diffs, submitting ..... 1175
distribution of Octave ..... 6
division ..... 170, 984
`do-until` statement ..... 198
documentation fonts ..... 11
documentation notation ..... 11
documenting functions ..... 40
documenting Octave programs ..... 39
documenting user scripts ..... 40
Dulmage-Mendelsohn decomposition ..... 746
dynamic classpath ..... 1138, 1145
dynamic classpath, adding new path ..... 1146
dynamic naming ..... 122
dynamic-linking ..... 1097
Dynamically Linked Functions ..... 1097

## E

echoing executing commands ..... 35
editing the command line ..... 26
element-by-element evaluation ..... 176
`else` statement ..... 193
`elseif` statement ..... 193
`end` statement ..... 193
`end`, indexing ..... 161
`end:` and `:end` ..... 161
`end_try_catch` ..... 202
`end_unwind_protect` ..... 202
`endfor` statement ..... 198
`endfunction` statement ..... 205
`endif` statement ..... 193
`endswitch` statement ..... 195
`endwhile` statement ..... 197
equality operator ..... 174, 984
equality, tests for ..... 174, 984
equations, nonlinear ..... 707
erroneous messages ..... 1173
erroneous results ..... 1173, 1175
error bar series ..... 556
error ids ..... 260
error message notation ..... 12
error messages ..... 36
error messages, incorrect ..... 1173
escape sequence notation ..... 75
evaluation notation ..... 11
executable scripts ..... 37
execution speed ..... 704
exiting octave ..... 7, 20
exponentiation ..... 170, 984
expression, range ..... 56
expressions ..... 159
expressions, assignment ..... 179
expressions, boolean ..... 176, 984
expressions, comparison ..... 174, 984
expressions, logical ..... 176, 984

## F

factorial function ..... 169
fatal signal ..... 1173
field, returning value of Java object field ..... 1144
field, setting value of Java object field ..... 1144
fields, displaying available fields of a Java object ..... 1144
figure deletefcn ..... 170
figure graphics object ..... 449
finding minimums ..... 712
finish.m ..... 20
flag character (`printf`) ..... 317
flag character (`scanf`) ..... 323
`for` statement ..... 198
Frobenius norm ..... 654
function application ..... 695
function descriptions ..... 12
function file ..... 224
function handle ..... 170
`function` statement ..... 205
functions, deprecated ..... 1165
functions, obsolete ..... 1165
functions, user-defined ..... 205
funding Octave development ..... 6

## G

general p-norm ..... 654
`global` statement ..... 147
global variables ..... 147
GNUTERM ..... 562
grammar rules ..... 1193
graphics ..... 333
graphics colors ..... 545
graphics data structures ..... 448
graphics line styles ..... 545
graphics marker styles ..... 546
graphics object properties ..... 460
graphics object, axes ..... 450
graphics object, figure ..... 449
graphics object, image ..... 450
graphics object, light ..... 450

graphics object, line ... 450
graphics object, patch ... 450
graphics object, root ... 449
graphics object, surface ... 450
graphics object, text ... 450
graphics objects ... 449
graphics objects, saving ... 459
graphics properties, default ... 543
graphics toolkits ... 561
greater than operator ... 174, 984
group objects ... 558, 559, 560

## H

handle functions ... 453
handle, function handles ... 248
hash table ... 130
help, online ... 21
help, user-defined functions ... 40
help, where to find ... 1176
Hermitian operator ... 170, 984
Hessenberg decomposition ... 659
history ... 1
history of commands ... 29

## I

`if` statement ... 193
image graphics object ... 450
improving Octave ... 1174, 1175
incorrect error messages ... 1173
incorrect output ... 1173, 1175
incorrect results ... 1173, 1175
increment operator ... 183
indirect function call ... 170
infinity norm ... 654
initialization ... 19
input conversions, for `scanf` ... 323
input history ... 29
`~/.inputrc` ... 33
installation trouble ... 1173
installing Octave ... 1179
instance, creating a Java instance ... 1143
introduction ... 7
introduction to graphics structures ... 448
invalid input ... 1173

## J

Java, calling from Octave ... 1138
Java, using with Octave ... 1138
`javaclasspath.txt` ... 1138
Jupyter Notebooks ... 246

## K

Kendall's Tau ... 868
key/value store ... 130
keywords ... 1193
known causes of trouble ... 1173

## L

language definition ... 1193
less than operator ... 174, 984
light graphics object ... 450
line graphics object ... 450
line series ... 557
line styles, graphics ... 545
linear algebra ... 647
linear algebra, techniques ... 647
Linear Indexing ... 161
loading data ... 294
local minimum ... 712
logging commands and output ... 35
logical expressions ... 176, 984
Logical Indexing ... 161
logical operators ... 176, 984
loop ... 197
looping over structure elements ... 199
LP ... 813
LU decomposition ... 659
lvalue ... 179

## M

Map ... 130
map ... 695
marker styles, graphics ... 546
matching failure, in `scanf` ... 322
matrices ... 52
matrices, diagonal and permutation ... 717
matrix factorizations ... 657
matrix functions, basic ... 647
matrix multiplication ... 170, 984
matrix, functions of ... 672
matrix, permutation functions ... 722
matrix, specialized solvers ... 676
matrix, zero elements ... 723
maximum field width (`scanf`) ... 323
memory management ... 704
memory, displaying Java memory status ... 1147
memory, limitations on JVM ... 1142
messages, error ... 36
method, invoking a method of a Java object ... 1145
methods, displaying available methods of a Java object ... 1145
mex ... 1123
mex-files ... 1123
minimum field width (`printf`) ... 317
missing data ... 45
mkoctfile ... 1098

multi-line comments ..... 39
multiplication ..... 170, 984

## N

negation ..... 170, 984
NLP ..... 813
nonlinear equations ..... 707
nonlinear programming ..... 813
not operator ..... 176, 984
numeric constant ..... 44, 51
numeric value ..... 44, 51

## O

object groups ..... 549
object, creating a Java object ..... 1143
obsolete functions ..... 1165
oct ..... 1098
oct-files ..... 1098
Octave and MATLAB, how to distinguish between ..... 1176
Octave API ..... 1097
Octave command options ..... 15
Octave, calling from Java ..... 1138
`.octaverc` ..... 16, 19
`~/.octaverc` ..... 16, 19
ODE ..... 791
online help ..... 21
opengl rendering slow windows ..... 562
opengl single precision date time ..... 563
operator precedence ..... 184
operators, arithmetic ..... 170, 984
operators, assignment ..... 179
operators, boolean ..... 176, 984
operators, decrement ..... 183
operators, increment ..... 183
operators, logical ..... 176, 984
operators, relational ..... 174, 984
optimization ..... 704, 813
options, Octave command ..... 15
`--braindead` ..... 17
`--built-in-docstrings-file` *filename* ..... 15
`--doc-cache-file` *filename* ..... 15
`--echo-commands` ..... 15
`--eval` *code* ..... 15
`--exec-path` *path* ..... 15
`--gui` ..... 15
`--help` ..... 15
`--image-path` *path* ..... 16
`--info-file` *filename* ..... 16
`--info-program` *program* ..... 16
`--init-trace` ..... 16
`--interactive` ..... 16
`--line-editing` ..... 16
`--no-gui` ..... 16
`--no-history` ..... 16
`--no-init-all` ..... 16
`--no-init-path` ..... 16
`--no-init-site` ..... 16
`--no-init-user` ..... 16
`--no-line-editing` ..... 16
`--no-window-system` ..... 17
`--norc` ..... 16
`--path` *path* ..... 17
`--persist` ..... 17
`--quiet` ..... 17
`--silent` ..... 17
`--texi-macros-file` *filename* ..... 17
`--traditional` ..... 17
`--version` ..... 18
`-f` ..... 16
`-h` ..... 15
`-H` ..... 16
`-i` ..... 16
`-p` *path* ..... 17
`-q` ..... 17
`-v` ..... 18
`-W` ..... 17
`-x` ..... 15
or operator ..... 176, 984
oregonator ..... 794
`otherwise` statement ..... 195
output conversions, for `printf` ..... 318

## P

parser ..... 1201
patch graphics object ..... 450
patches, submitting ..... 1175
path, adding to classpath ..... 1146
path, removing from classpath ..... 1146
permutation matrix functions ..... 722
`persistent` statement ..... 148
persistent variables ..... 148
personal startup file ..... 19
PKG_ADD ..... 1091
PKG_DEL ..... 1091
plotting ..... 333
plotting, high-level ..... 333
plotting, multiple plot windows ..... 425
plotting, multiple plots per figure ..... 423
plotting, object manipulation ..... 426
plotting, saving and printing plots ..... 435
plotting, three-dimensional ..... 382
plotting, two-dimensional functions ..... 378
plotting, window manipulation ..... 427
precision (`printf`) ..... 318
printing notation ..... 12
printing plots ..... 435
profiler ..... 280
program, self contained ..... 37
Progress Bar ..... 1007
project startup file ..... 19
prompt customization ..... 33
pseudoinverse ..... 655, 720

## Q

QP ... 813
QR factorization ... 661
quadratic programming ... 813
quitting octave ... 7, 20
quiver group ... 558
quotient ... 170, 984

## R

range expressions ... 56
`readline` customization ... 33
recycling ... 691
relational operators ... 174, 984
reporting bugs ... 1173, 1174
results, incorrect ... 1173, 1175
root graphics object ... 449

## S

saving data ... 294
saving graphics objects ... 459
saving plots ... 435
Schur decomposition ... 666
script files ... 205
scripts ... 37
select JVM version ... 1142
self contained programs ... 37
series objects ... 554, 555, 556, 557, 559
short-circuit evaluation ... 177
side effect ... 179
SIMD ... 691
singular value decomposition ... 669
site exiting file ... 20
site startup file ... 16, 19
Spearman's Rho ... 868
speedups ... 704
stair group ... 559
startup ... 19
startup files ... 19
`startup.m` ... 19
statements ... 193
static classpath ... 1138, 1145
stem series ... 559
strings ... 46, 75
structural rank ... 753
structure elements, looping over ... 199
structures ... 46, 117
submitting diffs ... 1175
submitting patches ... 1175
subtraction ... 170, 984
suggestions ... 1174
surface graphics object ... 450
surface group ... 560
`switch` statement ... 195

## T

test functions ... 1151
tests for equality ... 174, 984
text graphics object ... 450
toolkit customization ... 562
toolkits, graphics ... 561
transform groups ... 561
transpose ... 170, 984
transpose, complex-conjugate ... 170, 984
troubleshooting ... 1173
`try` statement ... 202

## U

uitable properties ... 529
unary minus ... 170, 984
undefined behavior ... 1173
undefined function value ... 1173
`unwind_protect` statement ... 202
`unwind_protect_cleanup` ... 202
use of comments ... 39
user-defined data types ... 46
user-defined functions ... 205
user-defined variables ... 145
using Octave with Java ... 1138

## V

validating arguments ... 216
`varargin` ... 213
`varargout` ... 212
variable-length argument lists ... 213
variable-length return lists ... 212
variables, global ... 147
variables, persistent ... 148
variables, user-defined ... 145
vectorization ... 689
vectorize ... 689
version startup file ... 19

## W

warning ids ... 265
warranty ... 1203
`while` statement ... 197
wrong answers ... 1173, 1175

# Function Index


_

__FILE__ ... 1200
__LINE__ ... 1201
_Exit ... 1063

## A

abs ... 605
accumarray ... 700
accumdim ... 702
acos ... 607
acosd ... 610
acosh ... 608
acot ... 608
acotd ... 610
acoth ... 609
acsc ... 608
acscd ... 610
acsch ... 609
add_input_event_hook ... 1201
addlistener ... 550
addpath ... 227
addpref ... 1021
addproperty ... 549
addtodate ... 1036
airy ... 628
all ... 565
allchild ... 458
amd ... 743
ancestor ... 458
and ... 176
angle ... 605
annotation ... 421
any ... 565
arch_fit ... 930
arch_rnd ... 931
arch_test ... 931
area ... 363, 364
arg ... 605
arguments ... 1195
argv ... 18
arma_rnd ... 931
arrayfun ... 695
ascii ... 1053
asctime ... 1027
asec ... 607
asecd ... 610
asech ... 608
asin ... 607
asind ... 610
asinh ... 608
assert ... 1158
assert_equal ... 1158
assignin ... 191
atan ... 607
atan2 ... 609
atan2d ... 610
atand ... 610
atanh ... 608
atexit ... 20
audiodevinfo ... 964
audioformats ... 964
audioinfo ... 963
audioplayer ... 965
audioread ... 963
audiorecorder ... 967
audiowrite ... 964
auto_repeat_debug_command ... 279
autoload ... 235
autoreg_matrix ... 932
autumn ... 953
available_graphics_toolkits ... 561
axes ... 450
axis ... 367

## B

balance ... 647
bandwidth ... 648
bar ... 339
barh ... 340
bartlett ... 932
base2dec ... 106
base64_decode ... 1058
base64_encode ... 1058
beep ... 258
beep_on_error ... 258
besselh ... 631
besseli ... 630
besselj ... 629
besselk ... 630
bessely ... 629
beta ... 631
betainc ... 632
betaincinv ... 632
betaln ... 633
bicg ... 676
bicgstab ... 678
bin2dec ... 103
binary ... 1053
bincoeff ... 633
bitand ... 64
bitcmp ... 65
bitget ... 64
bitor ... 65
bitpack ... 43
bitset ... 63
bitshift ... 65
bitunpack ... 44

bitxor ... 65
blackman ... 932
blanks ... 75
blkdiag ... 580
blkmm ... 675
bone ... 954
bounds ... 832
box ... 420
break ... 1199
brighten ... 958
bsxfun ... 692
built_in_docstrings_file ... 24
builtin ... 235
bunzip2 ... 1048
byte_size ... 1104
bzip2 ... 1050

## C

calendar ... 1036
camlight ... 397
camlookat ... 401
camorbit ... 402
campos ... 401
camroll ... 402
camtarget ... 403
camup ... 404
camva ... 404
camzoom ... 405
canonicalize_file_name ... 1047
cart2pol ... 641
cart2sph ... 642
case ... 1196
cast ... 42
cat ... 571
catch ... 1197
cbrt ... 604
ccolamd ... 743
cd ... 1052, 1069
ceil ... 616
cell ... 133
cell2mat ... 141
cell2struct ... 141
celldisp ... 132
cellfun ... 697
cellindexmat ... 139
cellslices ... 137
cellstr ... 140
center ... 857
cgs ... 680
char ... 83
chdir ... 1069
chol ... 657
chol2inv ... 657
choldelete ... 658
cholinsert ... 658
cholinv ... 657
cholshift ... 659
cholupdate ... 658
circshift ... 574
citation ... 6
cla ... 430
clabel ... 419
class ... 41
classdef ... 1199
clc ... 27
clear ... 154
clearAllMemoizedCaches ... 704
clearvars ... 155
clf ... 430
clim ... 368
clock ... 1029
close ... 431, 1051
closereq ... 431
cmdline_options ... 18
cmpermute ... 959
cmunique ... 959
colamd ... 744
colloc ... 780
colon ... 982
colorbar ... 420
colorcube ... 954
colormap ... 951
colperm ... 745
colstyle ... 546
columns ... 47
comet ... 366
comet3 ... 366
command_line_path ... 230
commandhistory ... 1023
commandwindow ... 1023
common_size ... 567
commutation_matrix ... 633
compan ... 878
compare_versions ... 1076
compass ... 362
completion_append_char ... 28
completion_matches ... 28
complex ... 52
computer ... 1072
cond ... 648
condeig ... 649
condest ... 751
confirm_recursive_rmdir ... 1040
conj ... 605
containers.Map ... 130
continue ... 1199
contour ... 348
contour3 ... 350
contourc ... 349
contourf ... 349
contrast ... 958
conv ... 879
conv2 ... 880
convhull ... 916
convhulln ... 917

convn ... 880
cool ... 954
copper ... 954
copyfile ... 1038
copyobj ... 459
corr ... 867
corrcoef ... 867
corrcov ... 867
cos ... 607
cosd ... 609
cosh ... 608
cosint ... 634
cospi ... 611
cot ... 607
cotd ... 610
coth ... 608
cov ... 865
cplxpair ... 605
cputime ... 1030
crash_dumps_octave_core ... 309
cross ... 623
csc ... 607
cscd ... 610
csch ... 608
cstrcat ... 84
csvread ... 302
csvwrite ... 302
csymamd ... 745
ctime ... 1026
ctranspose ... 172
cubehelix ... 954
cummax ... 619
cummin ... 620
cumprod ... 614
cumsum ... 613
cumtrapz ... 780
curl ... 624
cylinder ... 413

## D

daspect ... 409
daspk ... 794
daspk_options ... 795
dasrt ... 800
dasrt_options ... 802
dassl ... 798
dassl_options ... 799
date ... 1030
datenum ... 1032
datestr ... 1034
datetick ... 1037
datevec ... 1036
dawson ... 634
dbclear ... 276, 277
dbcont ... 274
dbdown ... 280
dblist ... 278
dblquad ... 781
dbnext ... 279
dbquit ... 274
dbstack ... 279
dbstatus ... 276
dbstep ... 279
dbstop ... 274
dbtype ... 278
dbup ... 280
dbwhere ... 278
deal ... 212
deblank ... 79
debug_java ... 1148
debug_on_error ... 273
debug_on_interrupt ... 273
debug_on_warning ... 273
dec2base ... 105
dec2bin ... 103
dec2hex ... 104
decic ... 808
deconv ... 880
deg2rad ... 606
del2 ... 624
delaunay ... 905
delaunayn ... 906
delete ... 430, 431, 1053
dellistener ... 551
demo ... 1160
det ... 649
detrend ... 932
diag ... 579
dialog ... 1006
diary ... 14, 35
diff ... 567
diffpara ... 932
diffuse ... 395
dims ... 1104
dir ... 1052, 1070
dir_encoding ... 231
dir_in_loadpath ... 230
discrete_cdf ... 869
discrete_inv ... 869
discrete_pdf ... 869
discrete_rnd ... 870
disp ... 287
display ... 988
dither ... 949
divergence ... 624
dlmread ... 302
dlmwrite ... 301
dmperm ... 746
do ... 1198
do_string_escapes ... 80
doc ... 22
doc_cache_create ... 25
doc_cache_file ... 24
dos ... 1060
dot ... 623

double ... 52
drawnow ... 427
dsearchn ... 912
dup2 ... 1064
duplication_matrix ... 634
durbinlevinson ... 933

## E

e ... 643
echo ... 36
edit ... 225
edit_history ... 30
EDITOR ... 33
eig ... 649
eigs ... 755
elem ... 1103
ellipj ... 634
ellipke ... 635
ellipsoid ... 414
else ... 1196
elseif ... 1196
empirical_cdf ... 869
empirical_inv ... 869
empirical_pdf ... 869
empirical_rnd ... 870
end ... 1194
end_try_catch ... 1197
end_unwind_protect ... 1198
endarguments ... 1195
endclassdef ... 1200
endenumeration ... 1200
endevents ... 1200
endfor ... 1198
endfunction ... 1194
endgrent ... 1072
endif ... 1196
endmethods ... 1200
endparfor ... 1199
endproperties ... 1200
endpwent ... 1071
endswitch ... 1197
endsWith ... 91
endwhile ... 1198
enumeration ... 1200
eomday ... 1037
eps ... 645
eq ... 175
erase ... 94
erf ... 635
erfc ... 635
erfcinv ... 636
erfcx ... 635
erfi ... 636
erfinv ... 636
errno ... 261
errno_list ... 261
error ... 255
errorbar ... 351
errordlg ... 1001
etime ... 1030
etree ... 737
etreeplot ... 738
eval ... 187
evalc ... 187
evalin ... 191
events ... 1200
example ... 1161
exec ... 1063
EXEC_PATH ... 1062
exist ... 153
exit ... 20
exp ... 603
expint ... 636
expm ... 672
expm1 ... 603
eye ... 580
ezcontour ... 380
ezcontourf ... 380
ezmesh ... 410
ezmeshc ... 411
ezplot ... 379
ezplot3 ... 409
ezpolar ... 381
ezsurf ... 411
ezsurfc ... 412

## F

factor ... 625
factorial ... 625
fail ... 1159
false ... 67
fclear ... 330
fclose ... 313
fcntl ... 1066
fdisp ... 300
feather ... 362, 363
feof ... 329
ferror ... 329
feval ... 189
fflush ... 293
fft ... 925
fft2 ... 925
fftconv ... 928
fftfilt ... 928
fftn ... 926
fftshift ... 933
fftw ... 927
fgetl ... 314
fgets ... 314
fieldnames ... 125
figure ... 425
file_in_loadpath ... 229
file_in_path ... 1045
fileattrib ... 1043

filebrowser ... 1023
fileparts ... 1046
fileread ... 300
filesep ... 1046
fill ... 364
fill3 ... 365
filter ... 928
filter2 ... 929
find ... 568
findall ... 543
findfigs ... 458
findobj ... 542
findstr ... 91
fix ... 616
fixed_point_format ... 55
flag ... 955
flintmax ... 61
flip ... 570
fliplr ... 569
flipud ... 570
floor ... 616
fminbnd ... 712
fminsearch ... 714
fminunc ... 713
foo ... 13
fopen ... 312
for ... 1198
fork ... 1063
format ... 288
fplot ... 378
fprintf ... 315
fputs ... 314
fractdiff ... 933
frame2im ... 951
fread ... 325
freport ... 330
freqz ... 929
freqz_plot ... 930
frewind ... 331
fscanf ... 321
fseek ... 330
fskipl ... 315
fsolve ... 707
ftell ... 330
ftp ... 1051
full ... 732
fullfile ... 1046
func2str ... 249
function ... 1194
functions ... 248
funm ... 673
fwrite ... 327
fzero ... 710

## G

gallery ... 593, 594, 595, 596, 597, 598
gamma ... 637
gammainc ... 637
gammaincinv ... 637
gammaln ... 639
gca ... 455
gcbf ... 548
gcbo ... 547
gcd ... 625
gcf ... 455
gco ... 456
ge ... 175
genpath ... 228
get ... 457, 967, 969
get_first_help_sentence ... 25
get_help_text ... 25
get_help_text_from_file ... 25
get_home_directory ... 1069
getappdata ... 548
getaudiodata ... 969
getegid ... 1068
getenv ... 1068
geteuid ... 1067
getfield ... 127
getframe ... 950
getgid ... 1068
getgrent ... 1072
getgrgid ... 1072
getgrnam ... 1072
gethostname ... 1051
getpgrp ... 1067
getpid ... 1067
getpixelposition ... 1018
getplayer ... 969
getppid ... 1067
getpref ... 1021
getpwent ... 1071
getpwnam ... 1071
getpwuid ... 1071
getrusage ... 1080
getuid ... 1068
ginput ... 446
givens ... 650
glob ... 1044
global ... 1193
glpk ... 813
gls ... 824
gmres ... 682
gmtime ... 1026
gnuplot_binary ... 562
gplot ... 738
grabcode ... 242
gradient ... 622
graphics_toolkit ... 561
gray ... 955
gray2ind ... 948
grid ... 420

griddata ..... 918
griddata3 ..... 919
griddatan ..... 919
groot ..... 454
gsvd ..... 651
gt ..... 175
gtext ..... 446
guidata ..... 1017
guihandles ..... 1018
gunzip ..... 1048
gzip ..... 1048

## H

hadamard ..... 598
hamming ..... 933
hankel ..... 598
hanning ..... 933
hash ..... 1081
have_window_system ..... 1018
hdl2struct ..... 459
height ..... 47
help ..... 21
helpdlg ..... 1001
hess ..... 659
hex2dec ..... 104
hex2num ..... 106
hggroup ..... 549
hgload ..... 445
hgsave ..... 444
hgtransform ..... 561
hidden ..... 385
hilb ..... 599
hist ..... 341
histc ..... 860
history ..... 29
history_control ..... 31
history_file ..... 32
history_save ..... 31
history_size ..... 32
history_timestamp_format_string ..... 32
hold ..... 429
home ..... 27
horzcat ..... 572
hot ..... 955
housh ..... 671
hsv ..... 955
hsv2rgb ..... 960
humps ..... 716
hurst ..... 934
hypot ..... 622

## I

i ..... 644
ichol ..... 763
idivide ..... 63
if ..... 1195
ifelse ..... 178
ifft ..... 925
ifft2 ..... 926
ifftn ..... 926
ifftshift ..... 933
ignore_function_time_stamp ..... 227
ilu ..... 764, 765
im2double ..... 947
im2frame ..... 951
imag ..... 606
image ..... 946
IMAGE_PATH ..... 941
imagesc ..... 946, 947
imfinfo ..... 942
imformats ..... 944
importdata ..... 308
imread ..... 939
imshow ..... 945
imwrite ..... 940
ind2gray ..... 948
ind2rgb ..... 949
ind2sub ..... 163
index ..... 92
inf ..... 644
Inf ..... 644
inferiorto ..... 986
info ..... 23
info_file ..... 23
info_program ..... 23
inpolygon ..... 915
input ..... 293
inputdlg ..... 1002
inputname ..... 208
inputParser ..... 222
inputParser.CaseSensitive ..... 223
inputParser.FunctionName ..... 223
inputParser.KeepUnmatched ..... 223
inputParser.Parameters ..... 222
inputParser.PartialMatching ..... 223
inputParser.Results ..... 223
inputParser.StructExpand ..... 223
inputParser.Unmatched ..... 223
inputParser.UsingDefaults ..... 223
int16 ..... 60
int2str ..... 101
int32 ..... 60
int64 ..... 60
int8 ..... 59
integral ..... 778
integral2 ..... 785
integral3 ..... 786
interp1 ..... 895
interp2 ..... 899

interp3 ..... 900
interpft ..... 897
interpn ..... 901
intersect ..... 873
intmax ..... 60
intmin ..... 61
inv ..... 651
inverse ..... 651
invhilb ..... 599
ipermute ..... 573
iqr ..... 833
is_absolute_filename ..... 1047
is_dq_string ..... 76
is_function_handle ..... 248
is_leap_year ..... 1030
is_rooted_relative_filename ..... 1047
is_same_file ..... 1047
is_sq_string ..... 76
is_valid_file_id ..... 313
isa ..... 41
isalnum ..... 113
isalpha ..... 113
isappdata ..... 549
isargout ..... 215
isascii ..... 115
isaxes ..... 454
isbanded ..... 71
isbool ..... 69
iscell ..... 132
iscellstr ..... 140
ischar ..... 77
iscntrl ..... 114
iscolormap ..... 953
iscolumn ..... 70
iscomplex ..... 69
isdebugmode ..... 279
isdefinite ..... 71
isdeployed ..... 1074
isdiag ..... 71
isdigit ..... 114
isdir ..... 1044
isempty ..... 49
isenv ..... 1068
isequal ..... 175
isequaln ..... 175
isfield ..... 126
isfigure ..... 454
isfile ..... 1044
isfinite ..... 567
isfloat ..... 69
isfolder ..... 1044
isglobal ..... 148
isgraph ..... 114
isgraphics ..... 454
isguirunning ..... 1018
ishandle ..... 454
ishermitian ..... 71
ishghandle ..... 454
ishold ..... 429
isieee ..... 1074
isindex ..... 164
isinf ..... 567
isinteger ..... 59
isjava ..... 1143
iskeyword ..... 1193
isletter ..... 113
islogical ..... 69
islower ..... 113
ismac ..... 1074
ismatrix ..... 69
ismember ..... 875
ismembertol ..... 875
ismethod ..... 975
isna ..... 46
isnan ..... 567
isnull ..... 49
isnumeric ..... 69
isobject ..... 974
isocaps ..... 391
isocolors ..... 392
isonormals ..... 390
isosurface ..... 388
ispc ..... 1074
isplaying ..... 967
ispref ..... 1022
isprime ..... 72
isprint ..... 114
isprop ..... 449
ispunct ..... 114
isreal ..... 69
isrecording ..... 969
isrow ..... 70
isscalar ..... 70
issorted ..... 577
isspace ..... 114
issparse ..... 734
issquare ..... 70
isstring ..... 77
isstrprop ..... 115
isstruct ..... 125
isstudent ..... 1074
issymmetric ..... 70
istril ..... 72
istriu ..... 72
isuniform ..... 72
isunix ..... 1074
isupper ..... 114
isvarname ..... 145
isvector ..... 70
isxdigit ..... 114
I ..... 644

## J

j ... 644
java_get ... 1144
java_matrix_autoconversion ... 1148
java_set ... 1144
java_unsigned_autoconversion ... 1148
javaaddpath ... 1146
javaArray ... 1143
javachk ... 1146
javaclasspath ... 1145
javamem ... 1147
javaMethod ... 1145
javaObject ... 1143
javarmpath ... 1146
jet ... 955
jsondecode ... 111
jsonencode ... 109
jupyter_notebook ... 246
J ... 644

## K

kbhit ... 294
kendall ... 868
keyboard ... 277
kill ... 1067
kron ... 674
krylov ... 671
kurtosis ... 840

## L

lasterr ... 260
lasterror ... 259
lastwarn ... 265
lcm ... 626
ldivide ... 172
le ... 175
legend ... 415
legendre ... 638
length ... 48
lgamma ... 639
license ... 1077
light ... 453
lightangle ... 398
lighting ... 396
lin2mu ... 970
line ... 450, 451
lines ... 956
link ... 1039
linkaxes ... 552
linkprop ... 552
linsolve ... 652
linspace ... 584
list_in_columns ... 287
list_primes ... 627
listdlg ... 1002
listfonts ... 1018
load ... 298
loaded_graphics_toolkits ... 561
loadobj ... 978
localfunctions ... 232
localtime ... 1026
log ... 603
log10 ... 603
log1p ... 603
log2 ... 603
logical ... 66
loglog ... 338
loglogerr ... 354
logm ... 672
logspace ... 585
lookfor ... 22
lookup ... 569
lower ... 79
ls ... 1069
ls_command ... 1070
lscov ... 825
lsode ... 791
lsode_options ... 792
lsqnonneg ... 824
lstat ... 1041
lt ... 175
lu ... 659
luupdate ... 660

## M

mad ... 834
magic ... 599
make_absolute_filename ... 1047
makeinfo_program ... 23
mape ... 863
mat2cell ... 135
mat2str ... 100
material ... 396
matlab.lang.makeUniqueStrings ... 146
matlab.lang.makeValidName ... 145
matlab.net.base64decode ... 1058
matlab.net.base64encode ... 1058
matlabroot ... 1074
matrix_type ... 652
max ... 617
max_recursion_depth ... 169
max_stack_depth ... 169
mean ... 829
meansq ... 835
median ... 830
memoize ... 703
memory ... 1077
menu ... 293
merge ... 178
mesh ... 383
meshc ... 384
meshgrid ... 398
meshz ... 384

methods ... 974, 975, 1200
mex ... 1123
mexext ... 1123
mfile_encoding ... 230
mfilename ... 227
mget ... 1052
mgorth ... 655
min ... 618
minus ... 172
mislocked ... 237
missing_component_hook ... 1095
missing_function_hook ... 1201
mkdir ... 1039, 1040, 1053
mkfifo ... 1040
mkoctfile ... 1098
mkpp ... 891
mkstemp ... 328
mktime ... 1027
mldivide ... 172
mlock ... 237
mod ... 627
mode ... 831
moment ... 841
more ... 291
movefile ... 1038
movegui ... 1019
movfun ... 846
movie ... 950
movmad ... 850
movmax ... 851
movmean ... 852
movmedian ... 853
movmin ... 853
movprod ... 854
movslice ... 848
movstd ... 855
movsum ... 855
movvar ... 856
mpoles ... 879
mpower ... 172
mput ... 1052
mrdivide ... 173
msgbox ... 1003, 1004
mtimes ... 173
mu2lin ... 970
munlock ... 237
mustBeFinite ... 220
mustBeGreaterThan ... 220
mustBeGreaterThanOrEqual ... 220
mustBeInteger ... 220
mustBeLessThan ... 220
mustBeLessThanOrEqual ... 221
mustBeMember ... 221
mustBeNegative ... 221
mustBeNonempty ... 221
mustBeNonNan ... 221
mustBeNonnegative ... 221
mustBeNonpositive ... 221
mustBeNonsparse ... 221
mustBeNonzero ... 222
mustBeNumeric ... 222
mustBeNumericOrLogical ... 222
mustBePositive ... 222
mustBeReal ... 222

## N

namedargs2cell ... 130
namelengthmax ... 146
nan ... 645
NaN ... 645
nargin ... 207
narginchk ... 217
nargout ... 211
nargoutchk ... 217
native_float_format ... 300
native2unicode ... 99
NA ... 45
nchoosek ... 861
ndgrid ... 399
ndims ... 46, 1104
ne ... 175
newline ... 81
newplot ... 428
news ... 22
nextpow2 ... 604
nnz ... 734
nonzeros ... 734
norm ... 654
normalize ... 859
normest ... 750
normest1 ... 750
not ... 177
now ... 1025
nproc ... 1073
nth_element ... 577
nthargout ... 210
nthroot ... 604
null ... 654
num2cell ... 134
num2hex ... 106
num2str ... 100
numArgumentsFromSubscript ... 981
numel ... 47, 1103
numfields ... 125
nzmax ... 734

## O

ocean ... 956
octave_core_file_limit ... 310
octave_core_file_name ... 310
octave_core_file_options ... 310
OCTAVE_EXEC_HOME ... 1074
OCTAVE_HOME ... 1074
OCTAVE_VERSION ... 1075
ode15i ... 807
ode15s ... 806
ode23 ... 804
ode23s ... 805
ode45 ... 803
odeget ... 811
odeplot ... 811
odeset ... 809
ols ... 823
onCleanup ... 262
ones ... 581
open ... 1060
openfig ... 445
openvar ... 1019
operator ... 1103
optimget ... 827
optimize_diagonal_matrix ... 717
optimize_permutation_matrix ... 717
optimize_range ... 57
optimize_subsasgn_calls ... 980
optimset ... 825
or ... 177
ordeig ... 669
orderfields ... 127
ordqz ... 668
ordschur ... 667
orient ... 444
orth ... 655
oruntests ... 1161
ostreamtube ... 360
ostrsplit ... 87
otherwise ... 1197
output_precision ... 54

## P

P_tmpdir ... 329
pack ... 156
padecoef ... 890
page_output_immediately ... 292
page_screen_output ... 292
pagectranspose ... 172
PAGER ... 292
PAGER_FLAGS ... 292
pagetranspose ... 174
pan ... 426
pareto ... 347
parfor ... 1199
parseparams ... 214
pascal ... 600
patch ... 451
path ... 229
pathdef ... 229
pathsep ... 229
pause ... 966, 968, 1032
pbaspect ... 409
pcg ... 758
pchip ... 934
pclose ... 1061
pcolor ... 363
pcr ... 761
peaks ... 447
periodogram ... 934
perl ... 1060
perms ... 861
permute ... 572
persistent ... 1193
pi ... 643
pie ... 356
pie3 ... 356
pink ... 956
pinv ... 655
pipe ... 1063
pkg ... 1083, 1084
planerot ... 651
play ... 966, 969
playblocking ... 966
plot ... 334
plot3 ... 399
plotmatrix ... 346
plotyy ... 336, 337
plus ... 173
pol2cart ... 641
polar ... 355
poly ... 893
polyaffine ... 883
polyarea ... 914
polyder ... 882
polyeig ... 878
polyfit ... 883
polygcd ... 880
polyint ... 883
polyout ... 893
polyreduce ... 893
polyval ... 877
polyvalm ... 877
popen ... 1061
popen2 ... 1061
postpad ... 579
pow2 ... 604
power ... 173
powerset ... 876
ppder ... 892
ppint ... 892
ppjumps ... 892
ppval ... 892
pqpnonneg ... 820
prctile ... 844

prefdir ..... 1022
preferences ..... 1022
prepad ..... 579
primes ..... 627
print ..... 436, 437
print_empty_dimensions ..... 56
print_struct_array_contents ..... 119
print_usage ..... 258
printd ..... 343
printf ..... 315
prism ..... 956
prod ..... 612
profexplore ..... 282
profexport ..... 282
profile ..... 281
profshow ..... 282
program_invocation_name ..... 18
program_name ..... 18
properties ..... 994, 1200
PS1 ..... 34
PS2 ..... 34
PS4 ..... 35
psi ..... 639
publish ..... 239, 240
putenv ..... 1068
puts ..... 314
pwd ..... 1070
python ..... 1061

## Q

qmr ..... 685
qp ..... 819
qr ..... 661
qrdelete ..... 664
qrinsert ..... 663
qrshift ..... 664
qrupdate ..... 663
quad ..... 772
quad_options ..... 773
quad2d ..... 783
quadcc ..... 776, 777
quadgk ..... 775
quadl ..... 774
quadv ..... 773
quantile ..... 842
questdlg ..... 1004, 1005
quit ..... 20
quiver ..... 357
quiver3 ..... 357
qz ..... 664
qzhess ..... 665

## R

rad2deg ..... 606
rainbow ..... 956
rand ..... 585
rande ..... 588
randg ..... 590
randi ..... 587
randn ..... 587, 588
randp ..... 589
randperm ..... 593
range ..... 833
rank ..... 655
ranks ..... 862
rat ..... 640
rats ..... 640
rcond ..... 656
rdivide ..... 173
readdir ..... 1044
readline_re_read_init_file ..... 33
readline_read_init_file ..... 33
readlink ..... 1039
real ..... 606
reallog ..... 603
realmax ..... 646
realmin ..... 646
realpow ..... 604
realsqrt ..... 604
record ..... 968, 971
recordblocking ..... 968
rectangle ..... 381
rectint ..... 915
recycle ..... 1048
reducepatch ..... 394
reducevolume ..... 393
refresh ..... 427
refreshdata ..... 553
regexp ..... 95
regexpi ..... 98
regexprep ..... 98
regexptranslate ..... 98
register_graphics_toolkit ..... 561
rehash ..... 229
rem ..... 626
remove_input_event_hook ..... 1201
rename ..... 1038, 1053
repelem ..... 582
repelems ..... 582
repmat ..... 582
rescale ..... 308
reset ..... 544
reshape ..... 573
residue ..... 881
resize ..... 573, 1104
restoredefaultpath ..... 230
resume ..... 966, 968
rethrow ..... 261
return ..... 209, 1195
rgb2gray ..... 961

rgb2hsv ..... 960
rgb2ind ..... 948
rgbplot ..... 953
ribbon ..... 406
rindex ..... 92
rmappdata ..... 548
rmdir ..... 1040, 1053
rmfield ..... 127
rmpath ..... 228
rmpref ..... 1022
rms ..... 836
rmse ..... 864
rng ..... 591
roots ..... 878
rose ..... 347, 348
rosser ..... 600
rot90 ..... 570
rotate ..... 426
rotate3d ..... 426
rotdim ..... 571
rotx ..... 921
roty ..... 922
rotz ..... 923
round ..... 617
roundb ..... 617
rows ..... 47
rref ..... 656
rsf2csf ..... 667
rticklabels ..... 375
rticks ..... 373
run ..... 189
run_count ..... 862
run_history ..... 30
rundemos ..... 1161
runlength ..... 863
rwdata ..... 1104

## S

S_ISBLK ..... 1042
S_ISCHR ..... 1042
S_ISDIR ..... 1042
S_ISFIFO ..... 1042
S_ISLNK ..... 1043
S_ISREG ..... 1043
S_ISSOCK ..... 1043
save ..... 295
save_default_options ..... 297
save_header_format_string ..... 298
save_precision ..... 297
saveas ..... 443
savefig ..... 445
saveobj ..... 978
savepath ..... 228
scanf ..... 321
scatter ..... 345, 346
scatter3 ..... 407
schur ..... 666
sec ..... 607
secd ..... 609
sech ..... 608
SEEK_CUR ..... 331
SEEK_END ..... 331
SEEK_SET ..... 331
semilogx ..... 337
semilogxerr ..... 353
semilogy ..... 338
semilogyerr ..... 354
set ..... 457, 967, 970
setappdata ..... 548
setdiff ..... 874
setenv ..... 1068
setfield ..... 126
setgrent ..... 1072
setpref ..... 1021
setpwent ..... 1071
setxor ..... 874
shading ..... 407
shg ..... 430
shiftdim ..... 575
shrinkfaces ..... 394
sighup_dumps_octave_core ..... 309
sign ..... 628
signbit ..... 628
sigquit_dumps_octave_core ..... 309
sigterm_dumps_octave_core ..... 310
SIG ..... 1067
silent_functions ..... 208
sin ..... 607
sinc ..... 930
sind ..... 609
sinetone ..... 935
sinewave ..... 935
single ..... 58
sinh ..... 608
sinint ..... 639
sinpi ..... 610
size ..... 48
size_equal ..... 49
sizemax ..... 54
sizeof ..... 49
skewness ..... 839
slice ..... 406
smooth3 ..... 392
sombrero ..... 447
sort ..... 575
sortrows ..... 576
sound ..... 971
soundsc ..... 971
source ..... 239
spalloc ..... 732
sparse ..... 732
spaugment ..... 754
spconvert ..... 734
spdiags ..... 729
spearman ..... 868

spectral_adf ..... 935
spectral_xdf ..... 936
specular ..... 395
speed ..... 1161
spencer ..... 936
speye ..... 730
spfun ..... 696
sph2cart ..... 642
sphere ..... 413
spinmap ..... 958
spline ..... 899
splinefit ..... 884
split_long_rows ..... 54
spones ..... 730
spparms ..... 752
sprand ..... 730
sprandn ..... 730
sprandsym ..... 731
sprank ..... 753
spring ..... 957
sprintf ..... 316
spstats ..... 735
spy ..... 737
sqp ..... 821
sqrt ..... 604
sqrtm ..... 672
squeeze ..... 49
sscanf ..... 322
stairs ..... 343
startsWith ..... 90
stat ..... 1041
statistics ..... 845
std ..... 836, 837
stderr ..... 311
stdin ..... 311
stdout ..... 311
stem ..... 344
stem3 ..... 345
stemleaf ..... 342
stft ..... 936
stop ..... 966, 969
str2double ..... 102
str2func ..... 249
str2num ..... 103
strcat ..... 84
strchr ..... 92
strcmp ..... 88
strcmpi ..... 89
stream2 ..... 361
stream3 ..... 361
streamline ..... 360
streamribbon ..... 358
streamtube ..... 359
strfind ..... 93
strftime ..... 1027
string_fill_char ..... 77
strjoin ..... 88
strjust ..... 78
strmatch ..... 93
strncmp ..... 89
strncmpi ..... 89
strptime ..... 1029
strread ..... 107
strrep ..... 94
strsplit ..... 86
strtok ..... 85
strtrim ..... 79
strtrunc ..... 80
struct ..... 124
struct_levels_to_print ..... 119
struct2cell ..... 129
struct2hdl ..... 459
structfun ..... 699
strvcat ..... 83
sub2ind ..... 163
subplot ..... 423, 424
subsasgn ..... 980
subsindex ..... 982
subspace ..... 669
subsref ..... 978
substr ..... 85
substruct ..... 128
sum ..... 611
summer ..... 957
sumsq ..... 615
superiorto ..... 986
suppress_verbose_help_message ..... 25
surf ..... 385
surface ..... 452
surfc ..... 386
surfl ..... 387
surfnorm ..... 388
svd ..... 669
svd_driver ..... 670
svds ..... 757
swapbytes ..... 43
switch ..... 1196
sylvester ..... 675
symamd ..... 747
symbfact ..... 753
symlink ..... 1039
symrcm ..... 748
symvar ..... 249
synthesis ..... 936
system ..... 1059

# T

tan ..... 607
tand ..... 609
tanh ..... 608
tar ..... 1049
tempdir ..... 328
tempname ..... 328
tensorprod ..... 675
terminal_size ..... 288

test ... 1151
tetramesh ... 909
texi_macros_file ... 24
text ... 417
textread ... 303
textscan ... 304
tfqmr ... 686
thetaticklabels ... 376
thetaticks ... 373
tic ... 1031
tilde_expand ... 1046
time ... 1025
times ... 173
title ... 414
tmpfile ... 328
toc ... 1031
toeplitz ... 600
tolower ... 79
toupper ... 79
trace ... 656
transpose ... 174
trapz ... 779
treelayout ... 738
treeplot ... 738
trexc ... 667
tril ... 577
trimesh ... 908
triplequad ... 782
triplot ... 907
trisurf ... 908
triu ... 578
true ... 66
try ... 1197
tsearch ... 911
tsearchn ... 911
turbo ... 957
type ... 156
typecast ... 42
typeinfo ... 41

## U

uibuttongroup ... 1009
uicontextmenu ... 1015
uicontrol ... 1009
uifigure ... 1008
uigetdir ... 999
uigetfile ... 999
uimenu ... 1014
uint16 ... 60
uint32 ... 60
uint64 ... 60
uint8 ... 60
uipanel ... 1008
uipushtool ... 1016
uiputfile ... 1000
uiresume ... 1020
uisetfont ... 1006
uitable ... 1011
uitoggletool ... 1017
uitoolbar ... 1015
uiwait ... 1020
umask ... 1041
uminus ... 174
uname ... 1073
undo_string_escapes ... 80
unicode_idx ... 92
unicode2native ... 99
union ... 873
unique ... 871
uniquetol ... 872
unix ... 1059
unlink ... 1039
unmkpp ... 891
unpack ... 1050
unsetenv ... 1068
untabify ... 80
untar ... 1049
until ... 1199
unwind_protect ... 1197
unwind_protect_cleanup ... 1198
unwrap ... 930
unzip ... 1049
uplus ... 174
upper ... 79
urlread ... 1054
urlwrite ... 1055
usejava ... 1147
user_config_dir ... 1075
user_data_dir ... 1075

## V

validateattributes ... 218
validatestring ... 217
vander ... 600
var ... 838
vec ... 579
vech ... 579
vecnorm ... 656
ver ... 1075
verLessThan ... 1076
version ... 1075
vertcat ... 572
view ... 400
viridis ... 957
voronoi ... 912
voronoin ... 913

## W

waitbar ... 1007
waitfor ... 1020
waitforbuttonpress ... 446
waitpid ... 1064
warndlg ... 1005
warning ... 263
warranty ... 23
waterfall ... 408
WCONTINUE ... 1064
WCOREDUMP ... 1064
web ... 1054
weboptions ... 1056
webread ... 1056
webwrite ... 1056
weekday ... 1037
WEXITSTATUS ... 1065
what ... 156
which ... 156
while ... 1198
white ... 957
whitebg ... 958
who ... 150
whos ... 151
whos_line_format ... 152
width ... 47
WIFCONTINUED ... 1065
WIFEXITED ... 1065
WIFSIGNALED ... 1065
WIFSTOPPED ... 1065
wilkinson ... 601
winqueryreg ... 1080
winter ... 957
WNOHANG ... 1066
workspace ... 1023
WSTOPSIG ... 1066
WTERMSIG ... 1066
WUNTRACED ... 1066

## X

xlabel ... 418
xlim ... 369
xor ... 566
xtickangle ... 377
xticklabels ... 374
xticks ... 372

## Y

yes_or_no ... 294
ylabel ... 418
ylim ... 370
ytickangle ... 377
yticklabels ... 374
yticks ... 372
yulewalker ... 937

## Z

zeros ... 581
zip ... 1049
zlabel ... 419
zlim ... 371
zoom ... 426
zscore ... 858
ztickangle ... 377
zticklabels ... 375
zticks ... 373

# Operator Index


## !

! ..... 176, 177, 985
!= ..... 174, 175, 985

## "

" ..... 46, 75

## #

# comment marker ..... 39
#! self-contained script ..... 37
#{ block comment marker ..... 39

## %

% comment marker ..... 39
%{ block comment marker ..... 39

## &

& ..... 176, 985
&& ..... 177

## '

' ..... 46, 75, 171, 172, 985

## (

( ..... 159

## )

) ..... 159

## *

* ..... 170, 173, 985
*= ..... 183

## +

+ ..... 170, 171, 173, 174, 985
++ ..... 184
+= ..... 183

## ,

, ..... 52

## -

- ..... 170, 171, 172, 174, 985
-- ..... 184
-= ..... 183

## .

. structure field access ..... 117
.' ..... 171, 174, 985
.* ..... 171, 173, 985
... continuation marker ..... 203
./ ..... 171, 173, 985
.^ ..... 171, 173, 985
.\ ..... 171, 172, 985

## /

/ ..... 171, 173, 985
/= ..... 183

## :

: ..... 159, 985
:, indexing expressions ..... 160
:, range expressions ..... 56

## ;

; ..... 52

## <

< ..... 174, 175, 985
<= ..... 174, 175, 985

## =

= ..... 179
== ..... 174, 175, 985

## >

> ..... 174, 175, 985
>= ..... 174, 175, 985

## [

[ ..... 52

## ]

] ..... 52

## ^

^ ..... 171, 172, 985

## @

@ class methods ..... 973
@ function handle ..... 170

## \

\ ..... 171, 172, 985
\ continuation marker ..... 203

## {

{ ..... 131

## }

} ..... 131

## |

| ..... 176, 177, 985
|| ..... 177

## ~

~ ..... 176, 177
~= ..... 174, 175, 985

# Graphics Properties Index


Axes Properties ... 470
Figure Properties ... 462
Image Properties ... 492
Legend Properties ... 483
Light Properties ... 511
Line Properties ... 485
Patch Properties ... 495
Root Properties ... 460
Scatter Properties ... 501
Surface Properties ... 505
Text Properties ... 489
Uibuttongroup Properties ... 516
Uicontextmenu Properties ... 519
Uicontrol Properties ... 525
Uimenu Properties ... 513
Uipanel Properties ... 521
Uipushtool Properties ... 536
Uitoggletool Properties ... 539
Uitoolbar Properties ... 534


## A


axes alim ... 478
axes alimmode ... 478
axes alphamap ... 478
axes alphascale ... 478
axes ambientlightcolor ... 478
axes beingdeleted ... 478
axes box ... 471
axes boxstyle ... 471
axes busyaction ... 476
axes buttondownfcn ... 479
axes cameraposition ... 477
axes camerapositionmode ... 477
axes cameratarget ... 477
axes cameratargetmode ... 477
axes cameraupvector ... 477
axes cameraupvectormode ... 477
axes cameraviewangle ... 477
axes cameraviewanglemode ... 477
axes children ... 481
axes clim ... 478
axes climmode ... 478
axes clipping ... 479
axes clippingstyle ... 479
axes color ... 471
axes colormap ... 478
axes colororder ... 470
axes colororderindex ... 470
axes colorscale ... 478
axes contextmenu ... 479
axes createfcn ... 478
axes currentpoint ... 479
axes dataaspectratio ... 471
axes dataaspectratiomode ... 471
axes deletefcn ... 478
axes fontangle ... 481
axes fontname ... 481
axes fontsize ... 481
axes fontsizemode ... 481
axes fontsmoothing ... 481
axes fontunits ... 481
axes fontweight ... 481
axes gridalpha ... 475
axes gridalphamode ... 475
axes gridcolor ... 475
axes gridcolormode ... 475
axes gridlinestyle ... 475
axes handlevisibility ... 481
axes hittest ... 479
axes innerposition ... 475
axes interactions ... 476
axes interruptible ... 476
axes labelfontsizemultiplier ... 482
axes layer ... 471
axes layout ... 471
axes legend ... 482
axes linestyleorder ... 470
axes linestyleorderindex ... 471
axes linewidth ... 471
axes minorgridalpha ... 475
axes minorgridalphamode ... 475
axes minorgridcolor ... 475
axes minorgridcolormode ... 475
axes minorgridlinestyle ... 476
axes mousewheelzoom ... 479
axes nextplot ... 480
axes nextseriesindex ... 471
axes outerposition ... 480
axes parent ... 481
axes pickableparts ... 479
axes plotboxaspectratio ... 480
axes plotboxaspectratiomode ... 480
axes position ... 480
axes positionconstraint ... 480
axes projection ... 477
axes selected ... 479
axes selectionhighlight ... 480
axes sortmethod ... 481
axes tag ... 480
axes tickdir ... 471
axes tickdirmode ... 471
axes ticklabelinterpreter ... 482
axes ticklength ... 472
axes tightinset ... 482
axes title ... 480
axes titlefontsizemultiplier ... 482

axes titlefontweight ... 482
axes toolbar ... 472
axes type ... 480
axes units ... 481
axes userdata ... 480
axes view ... 477
axes visible ... 479
axes xaxis ... 472
axes xaxislocation ... 472
axes xcolor ... 472
axes xcolormode ... 472
axes xdir ... 472
axes xgrid ... 476
axes xlabel ... 482
axes xlim ... 472
axes xlimitmethod ... 472
axes xlimmode ... 472
axes xminorgrid ... 476
axes xminortick ... 472
axes xminortickvalues ... 472
axes xminortickvaluesmode ... 473
axes xscale ... 476
axes xtick ... 473
axes xticklabel ... 482
axes xticklabelmode ... 482
axes xticklabelrotation ... 482
axes xtickmode ... 473
axes yaxis ... 473
axes yaxislocation ... 473
axes ycolor ... 473
axes ycolormode ... 473
axes ydir ... 473
axes ygrid ... 476
axes ylabel ... 482
axes ylim ... 473
axes ylimitmethod ... 473
axes ylimmode ... 473
axes yminorgrid ... 476
axes yminortick ... 476
axes yminortickvalues ... 473
axes yminortickvaluesmode ... 474
axes yscale ... 476
axes ytick ... 474
axes yticklabel ... 482
axes yticklabelmode ... 482
axes yticklabelrotation ... 482
axes ytickmode ... 474
axes zaxis ... 474
axes zcolor ... 474
axes zcolormode ... 474
axes zdir ... 474
axes zgrid ... 476
axes zlabel ... 482
axes zlim ... 474
axes zlimitmethod ... 474
axes zlimmode ... 474
axes zminorgrid ... 476
axes zminortick ... 474
axes zminortickvalues ... 474
axes zminortickvaluesmode ... 475
axes zscale ... 476
axes ztick ... 475
axes zticklabel ... 482
axes zticklabelmode ... 483
axes zticklabelrotation ... 483
axes ztickmode ... 475

## F

figure alphamap ... 463
figure beingdeleted ... 463
figure busyaction ... 463
figure buttondownfcn ... 465
figure children ... 469
figure clipping ... 470
figure closerequestfcn ... 463
figure color ... 463
figure colormap ... 463
figure contextmenu ... 465
figure createfcn ... 463
figure currentaxes ... 468
figure currentcharacter ... 468
figure currentobject ... 468
figure currentpoint ... 465
figure deletefcn ... 464
figure dockcontrols ... 468
figure filename ... 469
figure graphicssmoothing ... 463
figure handlevisibility ... 469
figure hittest ... 466
figure innerposition ... 468
figure integerhandle ... 468
figure interruptible ... 463
figure inverthardcopy ... 469
figure keypressfcn ... 464
figure keyreleasefcn ... 464
figure menubar ... 466
figure name ... 463
figure nextplot ... 468
figure number ... 468
figure numbertitle ... 463
figure outerposition ... 468
figure paperorientation ... 469
figure paperposition ... 469
figure paperpositionmode ... 469
figure papersize ... 469
figure papertype ... 469
figure paperunits ... 470
figure parent ... 469
figure pickableparts ... 470
figure pointer ... 466
figure pointershapecdata ... 466
figure pointershapehotspot ... 466
figure position ... 468
figure renderer ... 470
figure renderermode ... 470

figure resize ... 466
figure resizefcn ... 466
figure selected ... 466
figure selectionhighlight ... 466
figure selectiontype ... 466
figure sizechangedfcn ... 467
figure tag ... 468
figure toolbar ... 467
figure type ... 468
figure units ... 468
figure userdata ... 468
figure visible ... 464
figure windowbuttondownfcn ... 467
figure windowbuttonmotionfcn ... 467
figure windowbuttonupfcn ... 467
figure windowkeypressfcn ... 465
figure windowkeyreleasefcn ... 465
figure windowscrollwheelfcn ... 467
figure windowstate ... 464
figure windowstyle ... 464

## I

image alphadata ... 493
image alphadatamapping ... 493
image beingdeleted ... 493
image busyaction ... 492
image buttondownfcn ... 494
image cdata ... 493
image cdatamapping ... 493
image children ... 495
image clipping ... 493
image contextmenu ... 494
image createfcn ... 493
image deletefcn ... 493
image handlevisibility ... 495
image hittest ... 494
image interruptible ... 493
image parent ... 495
image pickableparts ... 494
image selected ... 494
image selectionhighlight ... 494
image tag ... 495
image type ... 495
image userdata ... 495
image visible ... 493
image xdata ... 494
image ydata ... 494

## L

legend autoupdate ... 483
legend box ... 484
legend color ... 484
legend edgecolor ... 484
legend fontangle ... 485
legend fontname ... 485
legend fontsize ... 485
legend fontunits ... 485
legend fontweight ... 485
legend itemhitfcn ... 483
legend location ... 484
legend numcolumns ... 484
legend orientation ... 484
legend position ... 484
legend string ... 485
legend textcolor ... 485
legend title ... 484
legend units ... 485
light beingdeleted ... 511
light busyaction ... 511
light buttondownfcn ... 512
light children ... 513
light clipping ... 512
light color ... 512
light contextmenu ... 512
light createfcn ... 511
light deletefcn ... 511
light handlevisibility ... 513
light hittest ... 512
light interruptible ... 511
light parent ... 513
light pickableparts ... 512
light position ... 512
light selected ... 512
light selectionhighlight ... 512
light style ... 512
light tag ... 513
light type ... 513
light userdata ... 513
light visible ... 512
line beingdeleted ... 486
line busyaction ... 486
line buttondownfcn ... 488
line children ... 488
line clipping ... 487
line color ... 487
line contextmenu ... 488
line createfcn ... 486
line deletefcn ... 487
line displayname ... 487
line handlevisibility ... 488
line hittest ... 488
line interruptible ... 486
line linejoin ... 487
line linestyle ... 487
line linewidth ... 487
line marker ... 487

line markeredgecolor ... 487
line markerfacecolor ... 487
line markersize ... 487
line parent ... 488
line pickableparts ... 488
line selected ... 488
line selectionhighlight ... 488
line tag ... 488
line type ... 488
line userdata ... 488
line visible ... 487
line xdata ... 486
line xdatasource ... 486
line ydata ... 486
line ydatasource ... 486
line zdata ... 486
line zdatasource ... 486

## P

patch alphadatamapping ... 496
patch ambientstrength ... 498
patch backfacelighting ... 498
patch beingdeleted ... 497
patch busyaction ... 495
patch buttondownfcn ... 499
patch cdata ... 496
patch cdatamapping ... 496
patch children ... 500
patch clipping ... 498
patch contextmenu ... 499
patch createfcn ... 497
patch deletefcn ... 498
patch diffusestrength ... 498
patch displayname ... 498
patch edgealpha ... 500
patch edgecolor ... 500
patch edgelighting ... 500
patch facealpha ... 496
patch facecolor ... 496
patch facelighting ... 496
patch facenormals ... 498
patch facenormalsmode ... 498
patch faces ... 497
patch facevertexalphadata ... 496
patch facevertexcdata ... 497
patch handlevisibility ... 500
patch hittest ... 499
patch interruptible ... 495
patch linestyle ... 500
patch linewidth ... 500
patch marker ... 499
patch markeredgecolor ... 499
patch markerfacecolor ... 499
patch markersize ... 499
patch parent ... 500
patch pickableparts ... 499
patch selected ... 500
patch selectionhighlight ... 500
patch specularcolorreflectance ... 498
patch specularexponent ... 498
patch specularstrength ... 498
patch tag ... 500
patch type ... 500
patch userdata ... 500
patch vertexnormals ... 499
patch vertexnormalsmode ... 499
patch vertices ... 497
patch visible ... 498
patch xdata ... 497
patch ydata ... 497
patch zdata ... 497

## R

root beingdeleted ... 462
root busyaction ... 462
root buttondownfcn ... 462
root callbackobject ... 460
root children ... 461
root clipping ... 462
root commandwindowsize ... 460
root contextmenu ... 460
root createfcn ... 462
root currentfigure ... 460
root deletefcn ... 462
root fixedwidthfontname ... 460
root handlevisibility ... 461
root hittest ... 462
root interruptible ... 462
root monitorpositions ... 461
root parent ... 461
root pickableparts ... 462
root pointerlocation ... 461
root pointerwindow ... 461
root screendepth ... 461
root screenpixelsperinch ... 461
root screensize ... 461
root selected ... 462
root selectionhighlight ... 462
root showhiddenhandles ... 461
root tag ... 460
root type ... 460
root units ... 461
root userdata ... 461
root visible ... 462

## S

scatter annotation ... 503
scatter beingdeleted ... 503
scatter busyaction ... 501
scatter buttondownfcn ... 504
scatter cdata ... 501
scatter cdatamode ... 501
scatter cdatasource ... 501
scatter children ... 505
scatter clipping ... 503
scatter contextmenu ... 504
scatter createfcn ... 503
scatter datatiptemplate ... 504
scatter deletefcn ... 503
scatter displayname ... 503
scatter handlevisibility ... 505
scatter hittest ... 504
scatter interruptible ... 501
scatter latitudedata ... 502
scatter latitudedatasource ... 502
scatter linewidth ... 503
scatter longitudedata ... 502
scatter longitudedatasource ... 502
scatter marker ... 503
scatter markeredgealpha ... 503
scatter markeredgecolor ... 504
scatter markerfacealpha ... 504
scatter markerfacecolor ... 504
scatter parent ... 505
scatter pickableparts ... 504
scatter rdata ... 502
scatter rdatasource ... 502
scatter selected ... 505
scatter selectionhighlight ... 505
scatter seriesindex ... 501
scatter sizedata ... 504
scatter sizedatasource ... 504
scatter tag ... 505
scatter thetadata ... 502
scatter thetadatasource ... 502
scatter type ... 505
scatter userdata ... 505
scatter visible ... 503
scatter xdata ... 502
scatter xdatasource ... 502
scatter ydata ... 502
scatter ydatasource ... 502
scatter zdata ... 502
scatter zdatasource ... 502
surface alphadata ... 506
surface alphadatamapping ... 506
surface ambientstrength ... 508
surface backfacelighting ... 508
surface beingdeleted ... 507
surface busyaction ... 505
surface buttondownfcn ... 509
surface cdata ... 506
surface cdatamapping ... 506
surface cdatasource ... 506
surface children ... 510
surface clipping ... 507
surface contextmenu ... 509
surface createfcn ... 507
surface deletefcn ... 507
surface diffusestrength ... 508
surface displayname ... 508
surface edgealpha ... 510
surface edgecolor ... 510
surface edgelighting ... 510
surface facealpha ... 507
surface facecolor ... 507
surface facelighting ... 508
surface facenormals ... 508
surface facenormalsmode ... 508
surface handlevisibility ... 510
surface hittest ... 509
surface interruptible ... 506
surface linestyle ... 510
surface linewidth ... 510
surface marker ... 509
surface markeredgecolor ... 509
surface markerfacecolor ... 509
surface markersize ... 509
surface meshstyle ... 510
surface parent ... 511
surface pickableparts ... 509
surface selected ... 509
surface selectionhighlight ... 510
surface specularcolorreflectance ... 508
surface specularexponent ... 508
surface specularstrength ... 508
surface tag ... 510
surface type ... 510
surface userdata ... 510
surface vertexnormals ... 508
surface vertexnormalsmode ... 509
surface visible ... 507
surface xdata ... 506
surface xdatasource ... 506
surface ydata ... 506
surface ydatasource ... 506
surface zdata ... 507
surface zdatasource ... 507

## T

text backgroundcolor ... 492
text beingdeleted ... 489
text busyaction ... 489
text buttondownfcn ... 490
text children ... 491
text clipping ... 489
text color ... 491
text contextmenu ... 490
text createfcn ... 489
text deletefcn ... 489

text edgecolor ... 492
text editing ... 491
text extent ... 490
text fontangle ... 491
text fontname ... 491
text fontsize ... 491
text fontsmoothing ... 492
text fontunits ... 492
text fontweight ... 492
text handlevisibility ... 491
text hittest ... 490
text horizontalalignment ... 490
text interpreter ... 492
text interruptible ... 489
text linestyle ... 492
text linewidth ... 492
text margin ... 492
text parent ... 491
text pickableparts ... 490
text position ... 491
text rotation ... 491
text selected ... 490
text selectionhighlight ... 490
text string ... 492
text tag ... 490
text type ... 490
text units ... 491
text userdata ... 490
text verticalalignment ... 491
text visible ... 489

## U

uibuttongroup backgroundcolor ... 516
uibuttongroup beingdeleted ... 517
uibuttongroup bordertype ... 516
uibuttongroup borderwidth ... 516
uibuttongroup busyaction ... 516
uibuttongroup buttondownfcn ... 517
uibuttongroup children ... 518
uibuttongroup clipping ... 517
uibuttongroup contextmenu ... 517
uibuttongroup createfcn ... 517
uibuttongroup deletefcn ... 517
uibuttongroup fontangle ... 519
uibuttongroup fontname ... 519
uibuttongroup fontsize ... 519
uibuttongroup fontunits ... 519
uibuttongroup fontweight ... 519
uibuttongroup foregroundcolor ... 516
uibuttongroup handlevisibility ... 518
uibuttongroup highlightcolor ... 516
uibuttongroup hittest ... 518
uibuttongroup interruptible ... 517
uibuttongroup parent ... 519
uibuttongroup pickableparts ... 518
uibuttongroup position ... 518
uibuttongroup resizefcn ... 517
uibuttongroup selected ... 518
uibuttongroup selectedobject ... 516
uibuttongroup selectionchangedfcn ... 517
uibuttongroup selectionhighlight ... 518
uibuttongroup shadowcolor ... 516
uibuttongroup sizechangedfcn ... 516
uibuttongroup tag ... 518
uibuttongroup title ... 519
uibuttongroup titleposition ... 519
uibuttongroup type ... 518
uibuttongroup units ... 518
uibuttongroup userdata ... 518
uibuttongroup visible ... 517
uicontextmenu beingdeleted ... 520
uicontextmenu busyaction ... 519
uicontextmenu buttondownfcn ... 520
uicontextmenu callback ... 520
uicontextmenu children ... 521
uicontextmenu clipping ... 520
uicontextmenu contextmenu ... 520
uicontextmenu createfcn ... 520
uicontextmenu deletefcn ... 520
uicontextmenu handlevisibility ... 521
uicontextmenu hittest ... 520
uicontextmenu interruptible ... 520
uicontextmenu parent ... 521
uicontextmenu pickableparts ... 521
uicontextmenu position ... 521
uicontextmenu selected ... 521
uicontextmenu selectionhighlight ... 521
uicontextmenu tag ... 521
uicontextmenu type ... 521
uicontextmenu userdata ... 521
uicontextmenu visible ... 520
uicontrol backgroundcolor ... 525
uicontrol beingdeleted ... 527
uicontrol busyaction ... 526
uicontrol buttondownfcn ... 528
uicontrol callback ... 526
uicontrol cdata ... 525
uicontrol children ... 529
uicontrol clipping ... 527
uicontrol contextmenu ... 528
uicontrol createfcn ... 527
uicontrol deletefcn ... 527
uicontrol enable ... 526
uicontrol extent ... 525
uicontrol fontangle ... 529
uicontrol fontname ... 529
uicontrol fontsize ... 529
uicontrol fontunits ... 529
uicontrol fontweight ... 529
uicontrol foregroundcolor ... 525
uicontrol handlevisibility ... 529
uicontrol hittest ... 528
uicontrol horizontalalignment ... 529
uicontrol interruptible ... 526
uicontrol keypressfcn ... 526

uicontrol listboxtop ... 526
uicontrol max ... 526
uicontrol min ... 527
uicontrol parent ... 529
uicontrol pickableparts ... 528
uicontrol position ... 528
uicontrol selected ... 528
uicontrol selectionhighlight ... 528
uicontrol sliderstep ... 527
uicontrol string ... 529
uicontrol style ... 525
uicontrol tag ... 528
uicontrol tooltipstring ... 528
uicontrol type ... 528
uicontrol units ... 528
uicontrol userdata ... 528
uicontrol value ... 527
uicontrol verticalalignment ... 529
uicontrol visible ... 527
uimenu accelerator ... 514
uimenu beingdeleted ... 514
uimenu busyaction ... 513
uimenu buttondownfcn ... 515
uimenu checked ... 514
uimenu children ... 515
uimenu clipping ... 514
uimenu contextmenu ... 515
uimenu createfcn ... 514
uimenu deletefcn ... 514
uimenu enable ... 514
uimenu foregroundcolor ... 513
uimenu handlevisibility ... 515
uimenu hittest ... 515
uimenu interruptible ... 513
uimenu menuselectedfcn ... 514
uimenu parent ... 516
uimenu pickableparts ... 515
uimenu position ... 515
uimenu selected ... 515
uimenu selectionhighlight ... 515
uimenu separator ... 513
uimenu tag ... 515
uimenu text ... 514
uimenu type ... 515
uimenu userdata ... 515
uimenu visible ... 514
uipanel backgroundcolor ... 522
uipanel beingdeleted ... 523
uipanel bordertype ... 522
uipanel borderwidth ... 522
uipanel busyaction ... 522
uipanel buttondownfcn ... 523
uipanel children ... 524
uipanel clipping ... 523
uipanel contextmenu ... 523
uipanel createfcn ... 523
uipanel deletefcn ... 523
uipanel fontangle ... 524
uipanel fontname ... 524
uipanel fontsize ... 524
uipanel fontunits ... 525
uipanel fontweight ... 525
uipanel foregroundcolor ... 522
uipanel handlevisibility ... 524
uipanel highlightcolor ... 522
uipanel hittest ... 523
uipanel interruptible ... 522
uipanel parent ... 524
uipanel pickableparts ... 523
uipanel position ... 524
uipanel resizefcn ... 522
uipanel selected ... 523
uipanel selectionhighlight ... 524
uipanel shadowcolor ... 522
uipanel sizechangedfcn ... 522
uipanel tag ... 524
uipanel title ... 525
uipanel titleposition ... 525
uipanel type ... 524
uipanel units ... 524
uipanel userdata ... 524
uipanel visible ... 523
uipushtool __named_icon__ ... 537
uipushtool beingdeleted ... 537
uipushtool busyaction ... 537
uipushtool buttondownfcn ... 538
uipushtool cdata ... 537
uipushtool children ... 539
uipushtool clickedcallback ... 537
uipushtool clipping ... 538
uipushtool contextmenu ... 538
uipushtool createfcn ... 537
uipushtool deletefcn ... 537
uipushtool enable ... 539
uipushtool handlevisibility ... 539
uipushtool hittest ... 538
uipushtool interruptible ... 537
uipushtool parent ... 539
uipushtool pickableparts ... 538
uipushtool selected ... 538
uipushtool selectionhighlight ... 538
uipushtool separator ... 537
uipushtool tag ... 539
uipushtool tooltipstring ... 538
uipushtool type ... 539
uipushtool userdata ... 539
uipushtool visible ... 538
uitable backgroundcolor ... 530
uitable beingdeleted ... 531
uitable busyaction ... 530
uitable buttondownfcn ... 532
uitable celleditcallback ... 530
uitable cellselectioncallback ... 530
uitable children ... 533
uitable clipping ... 532
uitable columneditable ... 534

uitable columnformat ... 533
uitable columnname ... 533
uitable columnwidth ... 533
uitable contextmenu ... 532
uitable createfcn ... 531
uitable data ... 533
uitable deletefcn ... 531
uitable enable ... 534
uitable extent ... 533
uitable fontangle ... 534
uitable fontname ... 534
uitable fontsize ... 534
uitable fontunits ... 534
uitable fontweight ... 534
uitable foregroundcolor ... 530
uitable handlevisibility ... 533
uitable hittest ... 532
uitable interruptible ... 530
uitable keypressfcn ... 530
uitable keyreleasefcn ... 531
uitable parent ... 533
uitable pickableparts ... 532
uitable position ... 533
uitable rearrangeablecolumns ... 534
uitable rowname ... 533
uitable rowstriping ... 530
uitable selected ... 532
uitable selectionhighlight ... 532
uitable tag ... 532
uitable tooltipstring ... 532
uitable type ... 532
uitable units ... 533
uitable userdata ... 532
uitable visible ... 532
uitoggletool __named_icon__ ... 539
uitoggletool beingdeleted ... 540
uitoggletool busyaction ... 539
uitoggletool buttondownfcn ... 540
uitoggletool cdata ... 539
uitoggletool children ... 541
uitoggletool clickedcallback ... 540
uitoggletool clipping ... 540
uitoggletool contextmenu ... 540
uitoggletool createfcn ... 540
uitoggletool deletefcn ... 540
uitoggletool enable ... 542
uitoggletool handlevisibility ... 541
uitoggletool hittest ... 541
uitoggletool interruptible ... 540
uitoggletool offcallback ... 540
uitoggletool oncallback ... 540
uitoggletool parent ... 541
uitoggletool pickableparts ... 541
uitoggletool selected ... 541
uitoggletool selectionhighlight ... 541
uitoggletool separator ... 539
uitoggletool state ... 542
uitoggletool tag ... 541
uitoggletool tooltipstring ... 541
uitoggletool type ... 541
uitoggletool userdata ... 541
uitoggletool visible ... 540
uitoolbar beingdeleted ... 535
uitoolbar busyaction ... 535
uitoolbar buttondownfcn ... 535
uitoolbar children ... 536
uitoolbar clipping ... 535
uitoolbar contextmenu ... 535
uitoolbar createfcn ... 535
uitoolbar deletefcn ... 535
uitoolbar handlevisibility ... 536
uitoolbar hittest ... 535
uitoolbar interruptible ... 535
uitoolbar parent ... 536
uitoolbar pickableparts ... 536
uitoolbar selected ... 536
uitoolbar selectionhighlight ... 536
uitoolbar tag ... 536
uitoolbar type ... 536
uitoolbar userdata ... 536
uitoolbar visible ... 535