

Table of Contents

Dive Into Python

Chapter 1. Installing Python

- 1.1 Which is right for you?
- 1.2 Python on Windows
- 1.3 Python on Mac OS X
- 1.4 Python on Mac OS 9
- 1.5 Python on RedHat Linux
- 1.6 Python on Debian GNU/Linux
- 1.7 Python Installation from Source
- 1.8 Interactive Shell
- 1.9. Summary

Chapter 2. Your Program

- 2.1 Diving
- 2.2 Enclosing Functions
- 2.3 Documenting Functions
- 2.4.1 An Object
- 2.5 Indenting
- 2.6 Modules

Chapter 3. Data Types

- 3.1 Diving
- 3.2 Lists
- 3.3 Tuples
- 3.4 Variables
- 3.5 String Formatting
- 3.6 Mapping
- 3.7 List and String Splitting
- 3.8. Summary

Chapter 4. Object Inspection

- 4.1 Diving
- 4.2 Using * and Named Arguments
- 4.3 print, dir, and Other Built-In Functions
- 4.4 Object References With getattr
- 4.5 Filtering
- 4.6 The Nature of and and or
- 4.7 Built-In Functions
- 4.8 Putting Together
- 4.9. Summary

Chapter 5. Object Orientation

- 5.1 Diving
- 5.2 Importing from module import
- 5.3 Class Defining
- 5.4 Class Instantiating
- 5.5 Explicitly Wrapper Class
- 5.6 Class Methods
- 5.7 Special Class Methods
- 5.8 Class Attributes

Table of Contents

Chapter 5: Orientation

- 5.9. Private
- 5.10. Summary

Chapter 6: File Objects

- 6.1. Opening
- 6.2. File Objects
- 6.3. File Objects
- 6.4. Modes
- 6.5. Directory
- 6.6. Putting together
- 6.7. Summary

Chapter 7: Regular

- 7.1. Diving
- 7.2. City Street Addresses
- 7.3. City Roman Numerals
- 7.4. Using Syntax
- 7.5. Regular Expressions
- 7.6. Parsing Phone Numbers
- 7.7. Summary

Chapter 8: HTML

- 8.1. Diving
- 8.2. Producing
- 8.3. From HTML documents
- 8.4. HTML Processor.py
- 8.5. Locals and globals
- 8.6. Dictionary-based
- 8.7. Quoting
- 8.8. Producing
- 8.9. Putting together
- 8.10. Summary

Chapter 9: XML

- 9.1. Diving
- 9.2. Packages
- 9.3. Parsing
- 9.4. Unicode
- 9.5. Searching
- 9.6. Generating
- 9.7. Segue

Chapter 10: Scripts

- 10.1. Abstracting
- 10.2. Standard and error
- 10.3. Clackings
- 10.4. Children of a node
- 10.5. Creating handlers by node type
- 10.6. Handling arguments
- 10.7. All together

Table of Contents

Chapter 10. Scripts

10.8. Summary

Chapter 11. HTTP

11.1. Diving
11.2. How to fetch data over HTTP
11.3. HTTP features
11.4. HTTP debugging services
11.5. User-Agent
11.6. Non-Modifiable and ETag
11.7. Handling
11.8. Handling data
11.9. Putting it all together
11.10. Summary

Chapter 12. SOAP

12.1. Diving
12.2. SOAP Libraries
12.3. First with SOAP
12.4. SOAP Web Services
12.5. WSDL Introducing
12.6. SOAP Web Services with WSDL
12.7. Searching
12.8. SOAP Web Services
12.9. Summary

Chapter 13. Unit

13.1. Introduction
13.2. Diving
13.3. Introduction
13.4. Testing
13.5. Testing
13.6. Testing

Chapter 14. First-First

14.1. man.py, stage 1
14.2. man.py, stage 2
14.3. man.py, stage 3
14.4. man.py, stage 4
14.5. man.py, stage 5

Chapter 15. Refactoring

15.1. Handling
15.2. Handling requirements
15.3. Refactoring
15.4. Postscript
15.5. Summary

Chapter 16. Functional

16.1. Diving
16.2. Finding

Table of Contents

Chapter 16. Functional

- 16.1. Lists Revisited
- 16.2. Mapping
- 16.3. Data-centric
- 16.4. Dynamically
- 16.5. All Together
- 16.8. Summary

Chapter 17. Dynamic

- 17.1. Diving
- 17.2. `l17.1al.py`, stage 1
- 17.3. `l17.1al.py`, stage 2
- 17.4. `l17.1al.py`, stage 3
- 17.5. `l17.1al.py`, stage 4
- 17.6. `l17.1al.py`, stage 5
- 17.7. `l17.1al.py`, stage 6
- 17.8. Summary

Chapter 18. Performance

- 18.1. Diving
- 18.2. `l18.1al.py` Module
- 18.3. `l18.1al.py` Expressions
- 18.4. `l18.1al.py` Optimizations
- 18.5. `l18.1al.py` List Operations
- 18.6. `l18.1al.py` Manipulation
- 18.7. Summary

Appendix A. Further

Appendix B. Review

Appendix C. Tips

Appendix D. List

Appendix E. Revision

Appendix F. About

Appendix G. GNU

- G.0. Preamble
- G.1. Applicability
- G.2. Copying
- G.3. Copying
- G.4. Modifications
- G.5. Combining
- G.6. Collections
- G.7. Aggregating works
- G.8. Translation
- G.9. Termination
- G.10. Future of this license

Table of Contents

Appendix G. Creative Commons Attribution License

How to use this License for your documents

Appendix H. Python

History of Python

Conditions for accessing or otherwise using Python

Dive Into Python

20 May 2004

Copyright © 2000, 2001, 2002, 2003, 2004 Mark Pilgrim (mailto:mark@diveintopython.org)

This book lives at <http://diveintopython.org/>. If you're reading it somewhere else, you may not have the latest version.

Permission is granted to copy, distribute, and/or modify this document under the terms of the GNU Free Documentation License, Version 1.1.

The example programs in this book are free software; you can redistribute and/or modify them under the terms of the Python license as published by the Free Software Foundation.

Chapter 1. Installing Python

Welcome to Python. Let's dive in. In this chapter, you'll install the version of Python that's right for you.

1.1. Which Python is right for you?

The first thing you need to do with Python is install it. Or do you?

If you're using an account on a hosted server, your ISP may have already installed Python. Most popular Linux distributions come with Python.

Windows does not come with any version of Python, but don't despair! There are several ways to point-and-click your way to Python on Windows.

As you can see already, Python runs on a great many operating systems. The full list includes Windows, Mac OS, Mac OS X, and all varieties of Unix.

What's more, Python programs written on one platform can, with a little care, run on *any* supported platform. For instance, I regularly develop Python programs on Linux that run on Windows.

So back to the question that started this section, "Which Python is right for you?" The answer is whichever one runs on the computer you already have.

1.2. Python on Windows

On Windows, you have a couple choices for installing Python.

ActiveState makes a Windows installer for Python called ActivePython, which includes a complete version of Python, an IDE with a Python interpreter, and a number of useful tools.

ActivePython is freely downloadable, although it is not open source. It is the IDE I used to learn Python, and I recommend you try it unless you have a strong preference for another IDE.

The second option is the "official" Python installer, distributed by the people who develop Python itself. It is freely downloadable and open source.

Procedure 1.1. Option 1: Installing ActivePython

Here is the procedure for installing ActivePython:

1. Download ActivePython from <http://www.activestate.com/Products/ActivePython/>.
2. If you are using Windows 95, Windows 98, or Windows ME, you will also need to download and install Windows Installer 2.0 (<http://www.microsoft.com/windows/win98/installer/>).
3. Double-click the installer, .
4. Step through the installer program.
5. If space is tight, you can do a custom installation and deselect the documentation, but I don't recommend this unless you absolutely have no other choice.
6. After the installation is complete, close the installer and choose Start->Programs->ActiveState ActivePython 2.2->PythonWin IDE.



>

Procedure 1.2. Option 2: Installing Python from Python.org (<http://www.python.org/>)

1. Download the latest Python Windows installer by going to <http://www.python.org/ftp/python/> and selecting the highest version number.
2. Double-click the installer,  The name will depend on the version of Python available when you read this.
3. Step through the installer program.
4. If disk space is tight, you can deselect the HTMLHelp file, the utility scripts  and/or the test suite .
5. If you do not have administrative rights on your machine, you can select Advanced Options, then choose Non-Admin Install. This just works.
6. After the installation is complete, close the installer and select Start->Programs->Python 2.3->IDLE (Python GUI). You'll see some





 >

1.3. Python on Mac OS X

On Mac OS X, you have two choices for installing Python: install it, or don't install it. You probably want to install it.

Mac OS X 10.2 and later comes with a command-line version of Python preinstalled. If you are comfortable with the command line, you can

Rather than using the preinstalled version, you'll probably want to install the latest version, which also comes with a graphical interactive shell.

Procedure 1.3. Running the Preinstalled Version of Python on Mac OS X

To use the preinstalled version of Python, follow these steps:

1. Open the  folder.
2. Open the  folder.

Double-click  to open a terminal window and get to a command line.3.

4.
Type  at the command prompt.

Try it out:



Procedure 1.4. Installing the Latest Version of Python on Mac OS X

Follow these steps to download and install the latest version of Python:

1.
Download the  disk image from <http://homepages.cwi.nl/~jack/macpython/download.html>.
2.
If your browser has not already done so, double-click  to mount the disk image on your desktop.
3.
Double-click the installer, .
4.
The installer will prompt you for your administrative username and password.
5.
Step through the installer program.
6.
After installation is complete, close the installer and open the  folder.
7.
Open the  folder
8.
Double-click  to launch Python.

The MacPython IDE should display a splash screen, then take you to the interactive shell. If the interactive shell does not appear, select Window



Note that once you install the latest version, the pre-installed version is still present. If you are running scripts from the command line, you need to specify the version.

Example 1.1. Two versions of Python



1.4. Python on Mac OS 9

Mac OS 9 does not come with any version of Python, but installation is very simple, and there is only one choice.

Follow these steps to install Python on Mac OS 9:

1. Download the  file from <http://homepages.cwi.nl/~jack/macpython/download.html>.
2. If your browser does not decompress the file automatically, double-click  to decompress the file with Stuffit Expander.
3. Double-click the installer, .
4. Step through the installer program.
5. After installation is complete, close the installer and open the  folder.
6. Open the  folder.
7. Double-click  to launch Python.

The MacPython IDE should display a splash screen, and then take you to the interactive shell. If the interactive shell does not appear, select View > Shell.





1.7. Python Installation from Source

If you prefer to build from source, you can download the Python source code from <http://www.python.org/ftp/python/>. Select the highest version available.

Example 1.4. Installing from source





1.8. The Interactive Shell

Now that you have Python installed, what's this interactive shell thing you're running?

It's like this: Python leads a double life. It's an interpreter for scripts that you can run from the command line or run like applications, by dou

Launch the Python interactive shell in whatever way works on your platform, and let's dive in with the steps shown here:

Example 1.5. First Steps in the Interactive Shell



- 1 The Python interactive shell can evaluate arbitrary Python expressions, including any basic arithmetic expression.
- 2 The interactive shell can execute arbitrary Python statements, including the **print** statement.
- 3 You can also assign values to variables, and the values will be remembered as long as the shell is open (but not any longer than that).

1.9. Summary

You should now have a version of Python installed that works for you.

Depending on your platform, you may have more than one version of Python installed. If so, you need to be aware of your paths. If simply ty

Congratulations, and welcome to Python.

Chapter 2. Your First Python Program

You know how other books go on and on about programming fundamentals and finally work up to building a complete, working program? L

2.1. Diving in

Here is a complete, working Python program.

It probably makes absolutely no sense to you. Don't worry about that, because you're going to dissect it line by line. But read through it first

Example 2.1.

If you have not already done so, you can download this and other examples (<http://diveintopython.org/download/diveintopython-examples-5>).

```
  
  
  
  
  
  
  
  
  
  
  
  

```

Now run this program and see what happens.

In the ActivePython IDE on Windows, you can run the Python program you're editing by choosing File->Run... (**Ctrl-R**). Output is displaye

In the Python IDE on Mac OS, you can run a Python program with Python->Run window... (**Cmd-R**), but there is an important option you m

On UNIX-compatible systems (including Mac OS X), you can run a Python program from the command line:

The output of  will look like this:

```

```

2.2. Declaring Functions

Python has functions like most other languages, but it does not have separate header files like C++ or  sections like Pascal. When you need a

```

```

Note that the keyword `def` starts the function declaration, followed by the function name, followed by the arguments in parentheses. Multiple arguments are separated by commas.

Also note that the function doesn't define a return datatype. Python functions do not specify the datatype of their return value; they don't even have a return type.

In Visual Basic, functions (that return a value) start with `Function` and subroutines (that do not return a value) start with `Sub`. There are no subroutines in Python.

The argument, `name`, doesn't specify a datatype. In Python, variables are never explicitly typed. Python figures out what type a variable is and keeps track of it.

In Java, C++, and other statically-typed languages, you must specify the datatype of the function return value and each function argument. In Python, you don't.

2.2.1. How Python's Datatypes Compare to Other Programming Languages

An erudite reader sent me this explanation of how Python compares to other programming languages:

statically typed language

A language in which types are fixed at compile time. Most statically typed languages enforce this by requiring you to declare the type of each variable.

dynamically typed language

A language in which types are discovered at execution time; the opposite of statically typed. VBScript and Python are dynamically typed.

strongly typed language

A language in which types are always enforced. Java and Python are strongly typed. If you have an integer, you can't treat it as a string.

weakly typed language

A language in which types may be ignored; the opposite of strongly typed. VBScript is weakly typed. In VBScript, you can treat an integer as a string.

So Python is both *dynamically typed* (because it doesn't use explicit datatype declarations) and *strongly typed* (because once a variable has a type, it can't be treated as a different type).

2.3. Documenting Functions

You can document a Python function by giving it a docstring.

Example 2.2. Defining the `len` Function's Docstring

```
def len(s):  
    """  
    Return the length of the string s.  
    """  
    return len(s)
```

Triple quotes signify a multi-line string. Everything between the start and end quotes is part of a single string, including carriage returns and other special characters.

Triple quotes are also an easy way to define a string with both single and double quotes, like `"""'\"'"""` in Perl.

Everything between the triple quotes is the function's docstring, which documents what the function does. A docstring, if it exists, must be the first thing defined in the function's body.

Many Python IDEs use the docstring to provide context-sensitive documentation, so that when you type a function name, its docstring appears as a tooltip. This is a handy feature.

Further Reading on Documenting Functions

PEP 257 (<http://www.python.org/peps/pep-0257.html>) defines **conventions**.

Python Style Guide (<http://www.python.org/doc/essays/styleguide.html>) discusses how to write a good **code**.

Python Tutorial (<http://www.python.org/doc/current/tut/tut.html>) discusses conventions for spacing in **code** (<http://www.python.org/doc/current/tut/tut.html>).

2.4. Everything Is an Object

In case you missed it, I just said that Python functions have attributes, and that those attributes are available at runtime.

A function, like everything else in Python, is an object.

Open your favorite Python IDE and follow along:

Example 2.3. Accessing the Function's Attributes

```
1 import sys
2
3 def foo():
4     return 'foo'
5
6 print(foo.__doc__)
```

- 1 The first line imports the **sys** module -- a chunk of code that you can use interactively, or from a larger Python program. (You can also import a module from a file.)
- 2 When you want to use functions defined in imported modules, you need to include the module name. So you can't just say `foo`; it must be `sys.foo`.
- 3 Instead of calling the function as you would expect to, you asked for one of the function's attributes, `foo.__doc__`.

In Python is like in Perl. Once you **import** a Python module, you access its functions with `module.function`; once you **require** a Perl module, you access its functions with `module->function`.

2.4.1. The Import Search Path

Before you go any further, I want to briefly mention the library search path. Python looks in several places when you try to import a module.

Example 2.4. Import Search Path

```
1 import sys
2
3 sys.path.append('/usr/local/lib/python2.7/site-packages')
4
5 import sys
6
7 print(sys.path)
```

- 1 Importing the **sys** module makes all of its functions and attributes available.
- 2 `sys.path` is a list of directory names that constitute the current search path. (Yours will look different, depending on your operating system, where you installed Python.)
- 3 Actually, I lied; the truth is more complicated than that, because not all modules are stored as **files**. Some, like the **sys** module, are "built-in" modules.
- 4 You can add a new directory to Python's search path at runtime by appending the directory name to `sys.path` and then Python will look in that directory for modules.

2.4.2. What's an Object?

Everything in Python is an object, and almost everything has attributes and methods. All functions have a built-in attribute `__doc__` which returns the docstring. Still, this begs the question. What is an object? Different programming languages define "object" in different ways. In some, it means that *all*...

This is so important that I'm going to repeat it in case you missed it the first few times: *everything in Python is an object*. Strings are objects.

Further Reading on Objects

Python Reference Manual (<http://www.python.org/doc/current/ref/>) explains exactly what it means to say that everything in Python is an object. *eff-bot* (<http://www.effbot.org/guides/>) summarizes Python objects (<http://www.effbot.org/guides/python-objects.htm>).

2.5. Indenting Code

Python functions have no explicit `{` and no curly braces to mark where the function code starts and stops. The only delimiter is a colon (`:`) at the end of the function signature.

Example 2.5. Indenting the Function

```
def fib(n):  
    """Return the n-th Fibonacci number"""  
    if n < 0:  
        raise ValueError("n must be non-negative")  
    if n < 2:  
        return n  
    return fib(n-1) + fib(n-2)
```

Code blocks are defined by their indentation. By "code block", I mean functions, statements, loops, loops, and so forth. Indenting starts a block.

Example 2.6, "if Statements" shows an example of code indentation with statements.

Example 2.6. Statements

```
def fib(n):  
    """Return the n-th Fibonacci number"""  
    if n < 0:  
        raise ValueError("n must be non-negative")  
    if n < 2:  
        return n  
    return fib(n-1) + fib(n-2)
```

- 1 This is a function named `fib` that takes one argument, `n`. All the code within the function is indented.
- 2 Printing to the screen is very easy in Python, just use `print`. Statements can take any data type, including strings, integers, and other native Python types.
- 3 `if` statements are a type of code block. If the `if` expression evaluates to true, the indented block is executed, otherwise it falls to the `else` block.
- 4 Of course `if` and `else` blocks can contain multiple lines, as long as they are all indented the same amount. This `if` block has two lines of code indented.

After some initial protests and several snide analogies to Fortran, you will make peace with this and start seeing its benefits. One major benefit is that it makes code blocks self-delimiting.

Python uses carriage returns to separate statements and a colon and indentation to separate code blocks. C++ and Java use semicolons to separate statements.

Further Reading on Code Indentation

Python Reference Manual (<http://www.python.org/doc/current/ref/>) discusses cross-platform indentation issues and shows various indentation styles. *Python Style Guide* (<http://www.python.org/doc/essays/styleguide.html>) discusses good indentation style.

2.6. Testing Modules

Python modules are objects and have several useful attributes. You can use this to easily test your modules as you write them. Here's an example:

```
def test():
```

Some quick observations before you get to the good stuff. First, parentheses are not required around the `if` expression. Second, the `if` statement is

Like C, Python uses `==` for comparison and `=` for assignment. Unlike C, Python does not support in-line assignment, so there's no chance of accidentally

So why is this particular `if` statement a trick? Modules are objects, and all modules have a built-in attribute `__name__`. A module's `__name__` depends on how you

```
    pass
```

```
if __name__ == '__main__':
```

```
    test()
```

Knowing this, you can design a test suite for your module within the module itself by putting it in this `if` statement. When you run the module

On MacPython, there is an additional step to make the trick work. Pop up the module's options menu by clicking the black triangle in the upper

Further Reading on Importing Modules

Python Reference Manual (<http://www.python.org/doc/current/ref/>) discusses the low-level details of importing modules (<http://www.python.org/doc/current/ref/>)

Chapter 3. Native Datatypes

You'll get back to your first Python program in just a minute. But first, a short digression is in order, because you need to know about dictionaries.

3.1. Introducing Dictionaries

One of Python's built-in datatypes is the dictionary, which defines one-to-one relationships between keys and values.

A dictionary in Python is like a hash in Perl. In Perl, variables that store hashes always start with a % character. In Python, variables can be named anything.

A dictionary in Python is like an instance of the Map class in Java.

A dictionary in Python is like an instance of the Dictionary object in Visual Basic.

3.1.1. Defining Dictionaries

Example 3.1. Defining a Dictionary

```
1 d = {'cat': 'feline', 'dog': 'canine'}
2 print(d)
3 print(d['cat'])
4 print(d.get('cat'))
5 print(d.get('bird'))
```

- 1 First, you create a new dictionary with two elements and assign it to the variable d. Each element is a key-value pair, and the whole set is a dictionary.
- 2 d is a key, and its associated value, referenced by feline.
- 3 cat is a key, and its associated value, referenced by feline.
- 4 You can get values by key, but you can't get keys by value. So d.get('feline') raises an exception, because feline is not a key.

3.1.2. Modifying Dictionaries

Example 3.2. Modifying a Dictionary

```
1 d = {'cat': 'feline', 'dog': 'canine'}
2 d['bird'] = 'avian'
3 print(d)
4 d['cat'] = 'feline'
5 print(d)
```

- ❶ You can not have duplicate keys in a dictionary. Assigning a value to an existing key will wipe out the old value.
- ❷ You can add new key-value pairs at any time. This syntax is identical to modifying existing values. (Yes, this will annoy you someday)

Note that the new element (key `1` value `3`) appears to be in the middle. In fact, it was just a coincidence that the elements appeared to be in order.

Dictionaries have no concept of order among elements. It is incorrect to say that the elements are "out of order"; they are simply unordered.

When working with dictionaries, you need to be aware that dictionary keys are case-sensitive.

Example 3.3. Dictionary Keys Are Case-Sensitive

```

dict = {}
dict['a'] = 1
dict['A'] = 2
dict['a'] = 3
dict['A'] = 4
dict['a'] = 5
dict['A'] = 6
dict['a'] = 7
dict['A'] = 8
dict['a'] = 9
dict['A'] = 10

```

- ❶ Assigning a value to an existing dictionary key simply replaces the old value with a new one.
- ❷ This is not assigning a value to an existing dictionary key, because strings in Python are case-sensitive, so `'a'` is not the same as `'A'`. This creates a new key-value pair.

Example 3.4. Mixing Datatypes in a Dictionary

```

dict = {}
dict[1] = 'a'
dict[2] = 'b'
dict[3] = 'c'
dict[4] = 'd'
dict[5] = 'e'
dict[6] = 'f'
dict[7] = 'g'
dict[8] = 'h'
dict[9] = 'i'
dict[10] = 'j'

```

- ❶ Dictionaries aren't just for strings. Dictionary values can be any datatype, including strings, integers, objects, or even other dictionaries.
- ❷ Dictionary keys are more restricted, but they can be strings, integers, and a few other types. You can also mix and match key datatypes.

3.1.3. Deleting Items From Dictionaries

Example 3.5. Deleting Items from a Dictionary

```

dict = {}
dict[1] = 'a'
dict[2] = 'b'
dict[3] = 'c'
dict[4] = 'd'
dict[5] = 'e'
dict[6] = 'f'
dict[7] = 'g'
dict[8] = 'h'
dict[9] = 'i'
dict[10] = 'j'

```



- ❶ Lets you delete individual items from a dictionary by key.
- ❷ Deletes all items from a dictionary. Note that the set of empty curly braces signifies a dictionary without any items.

Further Reading on Dictionaries

How to Think Like a Computer Scientist (<http://www.ibiblio.org/obp/thinkCSpy/>) teaches about dictionaries and shows how to use them.
 Python Knowledge Base (<http://www.faqs.com/knowledge-base/index.phtml/fid/199/>) has a lot of example code using dictionaries.
 Python Cookbook (<http://www.activestate.com/ASPN/Python/Cookbook/>) discusses how to sort the values of a dictionary by key (http://www.python.org/doc/current/lib/).
Python Library Reference (<http://www.python.org/doc/current/lib/>) summarizes all the dictionary methods (<http://www.python.org/doc/current/lib/>).

3.2. Introducing Lists

Lists are Python's workhorse datatype. If your only experience with lists is arrays in Visual Basic or (God forbid) the datastore in Powerbuilder...

A list in Python is like an array in Perl. In Perl, variables that store arrays always start with the @ character; in Python, variables can be named anything.

A list in Python is much more than an array in Java (although it can be used as one if that's really all you want out of life). A better analogy would be...

3.2.1. Defining Lists

Example 3.6. Defining a List



- ❶ First, you define a list of five elements. Note that they retain their original order. This is not an accident. A list is an ordered set of elements.
- ❷ A list can be used like a zero-based array. The first element of any non-empty list is always 0.
- ❸ The last element of this five-element list is 4 because lists are always zero-based.

Example 3.7. Negative List Indices



- ❶ A negative index accesses elements from the end of the list counting backwards. The last element of any non-empty list is always -1.
- ❷ If the negative index is confusing to you, think of it this way: -1. So in this list, -1.

Example 3.8. Slicing a List

```
l = ['a', 'b', 'c', 'd', 'e']
l[0:3]
```

- 1 You can get a subset of a list, called a "slice", by specifying two indices. The return value is a new list containing all the elements of the original list between the two indices.
- 2 Slicing works if one or both of the slice indices is negative. If it helps, you can think of it this way: reading the list from left to right, the first element is at index 0, the second at index 1, and so on. Negative indices count from the end of the list: the last element is at index -1, the second to last at index -2, and so on.
- 3 Lists are zero-based, so `l[0:3]` returns the first three elements of the list, starting at index 0 up to but not including index 3.

Example 3.9. Slicing Shorthand

```
l = ['a', 'b', 'c', 'd', 'e']
l[:3]
l[2:]
l[0:5]
```

- 1 If the left slice index is 0, you can leave it out, and 0 is implied. So `l[:3]` is the same as `l[0:3]` from Example 3.8, "Slicing a List".
- 2 Similarly, if the right slice index is the length of the list, you can leave it out. So `l[2:]` is the same as `l[2:5]` because this list has five elements.
- 3 Note the symmetry here. In this five-element list, `l[:3]` returns the first 3 elements, and `l[2:]` returns the last two elements. In fact, `l[:n]` will always return the first `n` elements of the list, and `l[-n:]` will always return the last `n` elements of the list.
- 4 If both slice indices are left out, all elements of the list are included. But this is not the same as the original list; it is a new list that happens to contain the same elements.

3.2.2. Adding Elements to Lists

Example 3.10. Adding Elements to a List

```
l = ['a', 'b', 'c']
l.append('d')
l.insert(1, 'x')
l1 = ['a', 'b']
l2 = ['c', 'd']
l.extend(l2)
```

- 1 `append()` adds a single element to the end of the list.
- 2 `insert()` inserts a single element into a list. The numeric argument is the index of the first element that gets bumped out of position. Note that `insert()` does not return anything.
- 3 `extend()` concatenates lists. Note that you do not call `extend()` with multiple arguments; you call it with one argument, a list. In this case, that list has two elements.

Example 3.11. The Difference between `del` and `pop`



- ❶ Lists have two methods, `append()` and `extend()` that look like they do the same thing, but are in fact completely different. `append()` takes a single argument, which can be any data type, and simply adds it to the end of the list. Here, you're calling the `append()` method on the list `l`, passing in the list `[d, e, f]`.
- ❷ Here you started with a list of three elements (`a`, `b`, and `c`), and you extended the list with a list of another three elements (`d`, `e`, and `f`), so you now have a list of six elements.
- ❸ On the other hand, `extend()` takes one argument, which can be any data type, and simply adds it to the end of the list. Here, you're calling the `extend()` method on the list `l`, passing in the list `[d, e, f]`.
- ❹ Now the original list, which started as a list of three elements, contains four elements. Why four? Because the last element that you just added was a list, and lists are mutable. So the last element of the original list is now a list containing three elements.

3.2.3. Searching Lists

Example 3.12. Searching a List



- ❶ `index()` finds the first occurrence of a value in the list and returns the index.
- ❷ `index()` finds the *first* occurrence of a value in the list. In this case, `2` occurs twice in the list, in `l[1]` and `l[2]`, but `index()` will return only the first index, `2`.
- ❸ If the value is not found in the list, Python raises an exception. This is notably different from most languages, which will return some "not found" value.
- ❹ To test whether a value is in the list, use `in`, which returns `True` if the value is found or `False` if it is not.

Before version 2.2.1, Python had no separate boolean datatype. To compensate for this, Python accepted almost anything in a boolean context: `0` is false; all other numbers are true. An empty string (`''`) is false, all other strings are true.

An empty list (`[]`) is false; all other lists are true.

An empty tuple (`()`) is false; all other tuples are true.

An empty dictionary (`{}`) is false; all other dictionaries are true.

These rules still apply in Python 2.2.1 and beyond, but now you can also use an actual boolean, which has a value of `True` or `False`. Note the capitalization.

3.2.4. Deleting List Elements

Example 3.13. Removing Elements from a List

```
1 # Remove the first occurrence of a value from a list
2 my_list = ['a', 'b', 'c', 'd', 'e', 'f', 'g', 'h', 'i', 'j', 'k', 'l', 'm', 'n', 'o', 'p', 'q', 'r', 's', 't', 'u', 'v', 'w', 'x', 'y', 'z']
3 my_list.remove('a')
4 print(my_list)
5 # Remove only the first occurrence of a value. In this case, 'a' appeared twice in the list, but only the first occurrence is removed.
6 my_list = ['a', 'b', 'c', 'd', 'e', 'f', 'g', 'h', 'i', 'a', 'j', 'k', 'l', 'm', 'n', 'o', 'p', 'q', 'r', 's', 't', 'u', 'v', 'w', 'x', 'y', 'z']
7 my_list.remove('a')
8 print(my_list)
9 # If the value is not found in the list, Python raises an exception. This mirrors the behavior of the pop() method.
10 my_list = ['a', 'b', 'c', 'd', 'e', 'f', 'g', 'h', 'i', 'j', 'k', 'l', 'm', 'n', 'o', 'p', 'q', 'r', 's', 't', 'u', 'v', 'w', 'x', 'y', 'z']
11 my_list.remove('x')
12 # pop() is an interesting beast. It does two things: it removes the last element of the list, and it returns the value that it removed. Note that this
13 # returns the value, not the index.
14 my_list = ['a', 'b', 'c', 'd', 'e', 'f', 'g', 'h', 'i', 'j', 'k', 'l', 'm', 'n', 'o', 'p', 'q', 'r', 's', 't', 'u', 'v', 'w', 'x', 'y', 'z']
15 my_list.pop()
16 print(my_list)
```

- 1 removes the first occurrence of a value from a list.
- 2 removes *only* the first occurrence of a value. In this case, 'a' appeared twice in the list, but removed only the first occurrence.
- 3 If the value is not found in the list, Python raises an exception. This mirrors the behavior of the `pop()` method.
- 4 `pop()` is an interesting beast. It does two things: it removes the last element of the list, and it returns the value that it removed. Note that this returns the value, not the index.

3.2.5. Using List Operators

Example 3.14. List Operators

```
1 # Lists can also be concatenated with the + operator. This has the same result as list concatenation.
2 my_list = ['a', 'b', 'c', 'd', 'e', 'f', 'g', 'h', 'i', 'j', 'k', 'l', 'm', 'n', 'o', 'p', 'q', 'r', 's', 't', 'u', 'v', 'w', 'x', 'y', 'z']
3 my_list2 = ['a', 'b', 'c', 'd', 'e', 'f', 'g', 'h', 'i', 'j', 'k', 'l', 'm', 'n', 'o', 'p', 'q', 'r', 's', 't', 'u', 'v', 'w', 'x', 'y', 'z']
4 my_list + my_list2
5 # Python supports the * operator. This is equivalent to list multiplication.
6 my_list * 2
7 # The * operator works on lists as a repeater. This is equivalent to list concatenation.
8 my_list * 3
```

- 1 Lists can also be concatenated with the `+` operator. This has the same result as list concatenation. But the `+` operator returns a new (concatenated) list as a value.
- 2 Python supports the `*` operator. This is equivalent to list multiplication. The `*` operator works for lists, strings, and integers, and it can be overloaded to work for other objects.
- 3 The `*` operator works on lists as a repeater. This is equivalent to list concatenation, which concatenates the three lists into one.

Further Reading on Lists

How to Think Like a Computer Scientist (<http://www.ibiblio.org/obp/thinkCSPy/>) teaches about lists and makes an important point about...

It makes your code safer if you "write-protect" data that does not need to be changed. Using a tuple instead of a list is like having an immutable list. Remember that I said that dictionary keys can be integers, strings, and "a few other types"? Tuples are one of those types. Tuples can be used in string formatting, as you'll see shortly.

Tuples can be converted into lists, and vice-versa. The built-in `list()` function takes a list and returns a tuple with the same elements, and the `tuple()` function takes a tuple and returns a list with the same elements.

Further Reading on Tuples

How to Think Like a Computer Scientist (<http://www.ibiblio.org/obp/thinkCSPy/>) teaches about tuples and shows how to concatenate them. Python Knowledge Base (<http://www.faqs.com/knowledge-base/index.phtml/fid/199/>) shows how to sort a tuple (<http://www.faqs.com/knowledge-base/index.phtml/fid/199/>). *Python Tutorial* (<http://www.python.org/doc/current/tut/tut.html>) shows how to define a tuple with one element (<http://www.python.org/doc/current/tut/tut.html>).

3.4. Declaring variables

Now that you know something about dictionaries, tuples, and lists (oh my!), let's get back to the sample program from Chapter 2.

Python has local and global variables like most other languages, but it has no explicit variable declarations. Variables spring into existence by being assigned a value.

Example 3.17. Defining the Variable

```
def main():
    """
    """
    """
    """
    """
    """
}
```

Notice the indentation. An `if` statement is a code block and needs to be indented just like a function.

Also notice that the variable assignment is one command split over several lines, with a backslash ("`\`") serving as a line-continuation marker.

When a command is split among several lines with the line-continuation marker ("`\`"), the continued lines can be indented in any manner; Python doesn't care.

Strictly speaking, expressions in parentheses, straight brackets, or curly braces (like defining a dictionary) can be split into multiple lines with the line-continuation marker.

Third, you never declared the variable `main` you just assigned a value to it. This is like VBScript without the `Option Explicit` option. Luckily, unlike VBScript, Python doesn't complain.

3.4.1. Referencing Variables

Example 3.18. Referencing an Unbound Variable

```
x
y
z
x
1
```

You will thank Python for this one day.

3.4.2. Assigning Multiple Values at Once

One of the cooler programming shortcuts in Python is using sequences to assign multiple values at once.

Example 3.19. Assigning multiple values at once

```
#  
x  
y  
z  
a  
b  
c
```

1 `vals` is a tuple of three elements, and `xyz` is a tuple of three variables. Assigning one to the other assigns each of the values of `vals` to each of the variables in `xyz`.

This has all sorts of uses. I often want to assign names to a range of values. In C, you would use `enum` and manually list each constant and its associated value.

Example 3.20. Assigning Consecutive Values

```
#  
0  
1  
2  
3  
4  
5  
6
```

1 The built-in `range` function returns a list of integers. In its simplest form, it takes an upper limit and returns a zero-based list counting up to that limit.

2 `days` and `months` are the variables you're defining. (This example came from the `calendar` module, a fun little module that prints calendars, like the one in the next example.)

3 Now each variable has its value: `days` is 0, `months` is 1, and so forth.

You can also use multi-variable assignment to build functions that return multiple values, simply by returning a tuple of all the values. The `calendar` module does this.

Further Reading on Variables

Python Reference Manual (<http://www.python.org/doc/current/ref/>) shows examples of when you can skip the line continuation character.

How to Think Like a Computer Scientist (<http://www.ibiblio.org/obp/thinkCSpy/>) shows how to use multi-variable assignment to swap variables.

3.5. Formatting Strings

Python supports formatting values into strings. Although this can include very complicated expressions, the most basic usage is to insert values into strings.

String formatting in Python uses the same syntax as the `printf` function in C.

Example 3.21. Introducing String Formatting

```
#  
#
```



1 The whole expression evaluates to a string. The first %s replaced by the value of k the second %s replaced by the value of v All other cha

Note that %s a tuple. I told you they were good for something.

You might be thinking that this is a lot of work just to do simple string concatenation, and you would be right, except that string formatting i

Example 3.22. String Formatting vs. Concatenating



- 1 +s the string concatenation operator.
- 2 In this trivial case, string formatting accomplishes the same result as concatenation.
- 3 %s a tuple with one element. Yes, the syntax is a little strange, but there's a good reason for it: it's unambiguously a tuple. In fact, you
- 4 String formatting works with integers by specifying %d instead of %s
- 5 Trying to concatenate a string with a non-string raises an exception. Unlike string formatting, string concatenation works only when ev

As with %n in C, string formatting in Python is like a Swiss Army knife. There are options galore, and modifier strings to specially format many

Example 3.23. Formatting Numbers



- 1 The %f string formatting option treats the value as a decimal, and prints it to six decimal places.
- 2 The ".2" modifier of the %f option truncates the value to two decimal places.
- 3 You can even combine modifiers. Adding the + modifier displays a plus or minus sign before the value. Note that the ".2" modifier is st

Further Reading on String Formatting

Python Library Reference (<http://www.python.org/doc/current/lib/>) summarizes all the string formatting format characters ([http://www](http://www.python.org/doc/current/lib/))
Effective AWK Programming ([http://www-gnats.gnu.org:8080/cgi-bin/info2www?\(gawk\)Top](http://www-gnats.gnu.org:8080/cgi-bin/info2www?(gawk)Top)) discusses all the format characters (h

3.6. Mapping Lists

One of the most powerful features of Python is the list comprehension, which provides a compact way of mapping a list into another list by a

Example 3.24. Introducing List Comprehensions

```
1 # Create a list of numbers from 1 to 10
2 numbers = list(range(1, 11))
3
4 # Create a list of squares of numbers from 1 to 10
5 squares = []
6 for number in numbers:
7     squares.append(number**2)
8
9 # Create a list of squares of numbers from 1 to 10 using list comprehension
10 squares = [number**2 for number in numbers]
```

- 1 To make sense of this, look at it from right to left. It is the list you're mapping. Python loops through one element at a time, temporarily.
- 2 Note that list comprehensions do not change the original list.
- 3 It is safe to assign the result of a list comprehension to the variable that you're mapping. Python constructs the new list in memory, and

Here are the list comprehensions in the function that you declared in Chapter 2:

```
1 # Create a list of squares of numbers from 1 to 10 using list comprehension
```

First, notice that you're calling the function of the dictionary. This function returns a list of tuples of all the data in the dictionary.

Example 3.25. The keys and values Functions

```
1 # Create a dictionary
2 dictionary = {'a': 1, 'b': 2, 'c': 3}
3
4 # Create a list of keys
5 keys = dictionary.keys()
6
7 # Create a list of values
8 values = dictionary.values()
9
10 # Create a list of tuples of the form (key, value)
11 tuples = dictionary.items()
```

- 1 The keys method of a dictionary returns a list of all the keys. The list is not in the order in which the dictionary was defined (remember that dictionaries are unordered).
- 2 The values method returns a list of all the values. The list is in the same order as the list returned by keys for all values of n.
- 3 The items method returns a list of tuples of the form (key, value). The list contains all the data in the dictionary.

Now let's see what map does. It takes a list, function, and maps it to a new list by applying string formatting to each element. The new list will have the same

Example 3.26. List Comprehensions in a Step by Step

```
1 # Create a list of numbers from 1 to 10
2 numbers = list(range(1, 11))
3
4 # Create a list of squares of numbers from 1 to 10 using list comprehension
5 squares = [number**2 for number in numbers]
```

- 1 Note that you're using two variables to iterate through the list. This is another use of multi-variable assignment. The first element of
- 2 Here you're doing the same thing, but ignoring the value of k so this list comprehension ends up being equivalent to
- 3 Combining the previous two examples with some simple string formatting, you get a list of strings that include both the key and value

Further Reading on List Comprehensions

Python Tutorial (<http://www.python.org/doc/current/tut/tut.html>) discusses another way to map lists using the built-in `map` function (<http://www.python.org/doc/current/tut/tut.html>) shows how to do nested list comprehensions (<http://www.python.org/doc/current/tut/tut.html>)

3.7. Joining Lists and Splitting Strings

You have a list of key-value pairs in the form `[(k, v)]` and you want to join them into a single string. To join any list of strings into a single string, use `"".join(l)`.

Here is an example of joining a list from the `dict.items()` function:

```
"".join("%s: %s" % (k, v) for k, v in dict.items())
```

One interesting note before you continue. I keep repeating that functions are objects, strings are objects... everything is an object. You might

The `join` method joins the elements of the list into a single string, with each element separated by a semi-colon. The delimiter doesn't need to be

`join` works only on lists of strings; it does not do any type coercion. Joining a list that has one or more non-string elements will raise an exception

Example 3.27. Output of `"".join()`

```
dict = {'a': 1, 'b': 2}
"".join("%s: %s" % (k, v) for k, v in dict.items())
'a: 1 b: 2'
```

This string is then returned from the `join` function and printed by the calling block, which gives you the output that you marveled at when you sta

You're probably wondering if there's an analogous method to split a string into a list. And of course there is, and it's called `split`

Example 3.28. Splitting a String

```
dict = {'a': 1, 'b': 2}
"".join("%s: %s" % (k, v) for k, v in dict.items()).split(" ")
['a:', '1', 'b:', '2']
```

- 1 `split` reverses `join` by splitting a string into a multi-element list. Note that the delimiter (" ") is stripped out completely; it does not appear in any
- 2 `split` takes an optional second argument, which is the number of times to split. ("Ooooooh, optional arguments..." You'll learn how to do th

`split` is a useful technique when you want to search a string for a substring and then work with everything before the substring (which ends up in

Further Reading on String Methods

Chapter 4. The Power Of Introspection

This chapter covers one of Python's strengths: introspection. As you know, everything in Python is an object, and introspection is code looking at itself.

4.1. Diving In

Here is a complete, working Python program. You should understand a good deal about it just by looking at it. The numbered lines illustrate

Example 4.1. `inspect.py`

If you have not already done so, you can download this and other examples (<http://diveintopython.org/download/diveintopython-examples-5>).

```
1 import inspect
2
3 def example(a, b, c=1):
4     """Example function with a docstring"""
5     pass
6
7 if __name__ == '__main__':
8     print('Example function: %s' % inspect.getdoc(example))
9     print('Parameters: %s' % inspect.getargspec(example))
```

- 1 This module has one function, `example`. According to its function declaration, it takes three parameters: `a`, `b`, and `c`. The last two are actually optional.
- 2 The function has a multi-line docstring that succinctly describes the function's purpose. Note that no return value is mentioned; this function will return `None`.
- 3 Code within the function is indented.
- 4 The `if __name__ == '__main__':` block allows this program do something useful when run by itself, without interfering with its use as a module for other programs.
- 5 Statements use `is` for comparison, and parentheses are not required.

The `example` function is designed to be used by you, the programmer, while working in the Python IDE. It takes any object that has functions or methods.

Example 4.2. Sample Usage of `inspect`

```
1 import inspect
2
3 def example(a, b, c=1):
4     """Example function with a docstring"""
5     pass
6
7 if __name__ == '__main__':
8     print('Example function: %s' % inspect.getdoc(example))
9     print('Parameters: %s' % inspect.getargspec(example))
```

By default the output is formatted to be easy to read. Multi-line strings are collapsed into a single long line, but this option can be changed by specifying `inspect.formatargspec(verify_kwarg=False)`.

Example 4.3. Advanced Usage of `len`

```
def len_advanced(s, *, sep=' ', ignore_whitespace=False):
    """Return the number of words in s.

    A word is defined as a sequence of characters separated by spaces.
    If ignore_whitespace is True, whitespace characters are ignored and
    words separated by them are considered to be a single word.

    """
    if ignore_whitespace:
        words = s.split(sep)
    else:
        words = s.split(sep, maxsplit=1)
    return len(words)
```

4.2. Using Optional and Named Arguments

Python allows function arguments to have default values; if the function is called without the argument, the argument gets its default value. For example, the `len` function has two optional arguments:

Here is an example of `len` function with two optional arguments:

```
def len_advanced(s, *, sep=' ', ignore_whitespace=False):
    """Return the number of words in s.

    A word is defined as a sequence of characters separated by spaces.
    If ignore_whitespace is True, whitespace characters are ignored and
    words separated by them are considered to be a single word.

    """
    if ignore_whitespace:
        words = s.split(sep)
    else:
        words = s.split(sep, maxsplit=1)
    return len(words)
```

Example 4.4. Valid Calls of `len`

- 1 With only one argument, `len` gets its default value of `0` and `ignore_whitespace` gets its default value of `False`.
- 2 With two arguments, `len` gets its default value of `0`.
- 3 Here you are naming the `ignore_whitespace` argument explicitly and specifying its value. `len` still gets its default value of `0`.
- 4 Even required arguments (like `s` which has no default value) can be named, and named arguments can appear in any order.

This looks totally whacked until you realize that arguments are simply a dictionary. The "normal" method of calling functions without arguments is just a shorthand for `len(s, *, ignore_whitespace=False)`.

The only thing you need to do to call a function is specify a value (somehow) for each required argument; the manner and order in which you specify them is up to you.

Further Reading on Optional Arguments

Python Tutorial (<http://www.python.org/doc/current/tut/tut.html>) discusses exactly when and how default arguments are evaluated (and how to use them).

4.3. Using `isinstance` and Other Built-In Functions

Python has a small set of extremely useful built-in functions. All other functions are partitioned off into modules. This was actually a conscious design decision.

4.3.1. The `isinstance` Function

The `isinstance` function returns the datatype of any arbitrary object. The possible types are listed in the `types` module. This is useful for helper functions that need to check the type of an object.

Example 4.5. Introducing `type()`

```
1 # type() returns the datatype of its argument
2 # type() can take a variable and return its datatype
3 # type() also works on modules
4 # You can use the constants in the sys module to compare types of objects. This is what the isinstance() function does, as you'll see shortly.
```

- 1 `type()` takes anything -- and I mean anything -- and returns its datatype. Integers, strings, lists, dictionaries, tuples, functions, classes, modules, etc.
- 2 `type()` can take a variable and return its datatype.
- 3 `type()` also works on modules.
- 4 You can use the constants in the `sys` module to compare types of objects. This is what the `isinstance()` function does, as you'll see shortly.

4.3.2. The `str()` Function

The `str()` function coerces data into a string. Every datatype can be coerced into a string.

Example 4.6. Introducing `str()`

```
1 # str() converts any object to a string
2 # str() works on any object of any type. Here it works on a list which you've constructed in bits and pieces.
3 # str() also works on modules. Note that the string representation of the module includes the pathname of the module on disk, so yours will be different.
4 # A subtle but important behavior of str() is that it works on None, the Python null value. It returns the string 'None'. You'll use this to your advantage in later chapters.
```

- 1 For simple datatypes like integers, you would expect `str()` to work, because almost every language has a function to convert an integer to a string.
- 2 However, `str()` works on any object of any type. Here it works on a list which you've constructed in bits and pieces.
- 3 `str()` also works on modules. Note that the string representation of the module includes the pathname of the module on disk, so yours will be different.
- 4 A subtle but important behavior of `str()` is that it works on `None`, the Python null value. It returns the string `'None'`. You'll use this to your advantage in later chapters.

At the heart of the `dir()` function is the powerful `__dir__` function. `__dir__` returns a list of the attributes and methods of any object: modules, functions, strings, lists, etc.

Example 4.7. Introducing `dir()`

```
1 # dir() returns a list of the attributes and methods of any object: modules, functions, strings, lists, etc.
2 # dir() works on any object of any type. Here it works on a list which you've constructed in bits and pieces.
```



- ❶ is a list, so dir returns a list of all the methods of a list. Note that the returned list contains the names of the methods as strings, not the methods themselves.
- ❷ is a dictionary, so dir returns a list of the names of dictionary methods. At least one of these, dir, should look familiar.
- ❸ This is where it really gets interesting. dir is a module, so dir returns a list of all kinds of stuff defined in the module, including built-in attributes.

Finally, the callable function takes any object and returns True if the object can be called, or False otherwise. Callable objects include functions, class methods, and some built-in objects.

Example 4.8. Introducing dir



- ❶ The functions in the dir module are deprecated (although many people still use the dir function), but the module contains a lot of useful constants.
- ❷ dir is a function that joins a list of strings.
- ❸ dir is not callable; it is a string. (A string does have callable methods, but the string itself is not callable.)
- ❹ dir is callable; it's a function that takes two arguments.
- ❺ Any callable object may have a __dir__ attribute. By using the dir function on each of an object's attributes, you can determine which attributes you care about.

4.3.3. Built-In Functions

dir and all the rest of Python's built-in functions are grouped into a special module called `__builtin__`. (That's two underscores before and after.) If it helps,

The advantage of thinking like this is that you can access all the built-in functions and attributes as a group by getting information about the `__builtin__` module.

Example 4.9. Built-in Attributes and Functions





Python comes with excellent reference manuals, which you should peruse thoroughly to learn all the modules Python has to offer. But unlike

Further Reading on Built-In Functions

Python Library Reference (<http://www.python.org/doc/current/lib/>) documents all the built-in functions (<http://www.python.org/doc/>)

4.4. Getting Object References With `getattr()`

You already know that Python functions are objects. What you don't know is that you can get a reference to a function without knowing its name

Example 4.10. Introducing `getattr()`



- ❶ This gets a reference to the `pop` method of the list. Note that this is not calling the `pop` method; that would be `l.pop()`. This is the method itself.
- ❷ This also returns a reference to the `pop` method, but this time, the method name is specified as a string argument to the `getattr()` function. `getattr(l, 'pop')` is an inc
- ❸ In case it hasn't sunk in just how incredibly useful this is, try this: the return value of `getattr(l, 'pop')` is the method, which you can then call just as if
- ❹ `getattr()` also works on dictionaries.
- ❺ In theory, `getattr()` would work on tuples, except that tuples have no methods, so `getattr(t, 'pop')` will raise an exception no matter what attribute name you give

4.4.1. `getattr()` with Modules

`getattr()` isn't just for built-in datatypes. It also works on modules.

Example 4.11. The `getattr()` Function in `__main__`





- 1 This returns a reference to the `function` in the `module`, which you studied in Chapter 2, *Your First Python Program*. (The hex address `0x10000000` is the address of the function object.)
- 2 Using `getattr` you can get the same reference to the same function. In general, `getattr(module, name)` is equivalent to `getattr(module, name)` if `module` is a module, then `name` can be anything defined in the module.
- 3 And this is what you actually use in the `function`. `getattr` is passed into the function as an argument; `name` is a string which is the name of a method.
- 4 In this case, `name` is the name of a function, which you can prove by getting its `__name__`.
- 5 Since `getattr` is a function, it is callable.

4.4.2. `getattr` as a Dispatcher

A common usage pattern of `getattr` is as a dispatcher. For example, if you had a program that could output data in a variety of different formats, you could use `getattr` to select the appropriate function.

For example, let's imagine a program that prints site statistics in HTML, XML, and plain text formats. The choice of output format could be based on a command-line argument or a configuration file.

Example 4.12. Creating a Dispatcher with `getattr`



- 1 The `output` function takes one required argument, `format`, and one optional argument, `obj`. If `obj` is not specified, it defaults to `None` and you will end up calling `output(format)`.
- 2 You concatenate the `format` argument with "output_" to produce a function name, and then go get that function from the `module`. This allows you to dynamically select the function to call.
- 3 Now you can simply call the output function in the same way as any other function. The `output` variable is a reference to the appropriate function.

Did you see the bug in the previous example? This is a very loose coupling of strings and functions, and there is no error checking. What happens if `format` is not a valid format?

Luckily, `getattr` takes an optional third argument, a default value.

Example 4.13. `getattr` Default Values



❶ This function call is guaranteed to work, because you added a third argument to the call to `filter`. The third argument is a default value that

As you can see, `filter` is quite powerful. It is the heart of introspection, and you'll see even more powerful examples of it in later chapters.

4.5. Filtering Lists

As you know, Python has powerful capabilities for mapping lists into other lists, via list comprehensions (Section 3.6, “Mapping Lists”). This

Here is the list filtering syntax:

```
[expression for item in list if filter]
```

This is an extension of the list comprehensions that you know and love. The first two thirds are the same; the last part, starting with the `if` is the

Example 4.14. Introducing List Filtering

```
1 # Create a list of numbers
2 numbers = [1, 2, 3, 4, 5, 6, 7, 8, 9, 10]
3 # Filter out the even numbers
4 evens = filter(lambda x: x % 2 == 0, numbers)
5 # Print the filtered list
6 print(evens)
```

❶ The mapping expression here is simple (it just returns the value of each element), so concentrate on the filter expression. As Python lo

❷ Here, you are filtering out a specific value, `0`. Note that this filters all occurrences of `0` since each time it comes up, the filter expression

❸ `count` is a list method that returns the number of times a value occurs in a list. You might think that this filter would eliminate duplicates from

Let's get back to this line from `filter`:

```
lambda x: x % 2 == 0
```

This looks complicated, and it is complicated, but the basic structure is the same. The whole filter expression returns a list, which is assigned

The filter expression looks scary, but it's not. You already know about `lambda` and `in`. As you saw in the previous section, the expression `x % 2 == 0` returns a f

So this expression takes an object (named `x`). Then it gets a list of the names of the object's attributes, methods, functions, and a few other thin

Further Reading on Filtering Lists

Python Tutorial (<http://www.python.org/doc/current/tut/tut.html>) discusses another way to filter lists using the built-in `filter` function (<http://www.python.org/doc/current/tut/tut.html>)

4.6. The Peculiar Nature of `and` and `or`

In Python, `and` and `or` perform boolean logic as you would expect, but they do not return boolean values; instead, they return one of the actual values

Example 4.15. Introducing `and`

```
1 # Create a list of numbers
2 numbers = [1, 2, 3, 4, 5, 6, 7, 8, 9, 10]
3 # Use and to get the last element
4 last = numbers[-1]
5 # Print the last element
6 print(last)
```


The trick, `is` will not work like the C expression `is` when `is` is false in a boolean context.

The real trick behind the trick, then, is to make sure that the value of `is` never false. One common way of doing this is to turn `is` into `is` and `is`

Example 4.19. Using the Trick Safely

```
is
is
is
,
```

1 Since `is` is a non-empty list, it is never false. Even if `is` or 'or' or some other false value, the list `is` is true because it has one element.

By now, this trick may seem like more trouble than it's worth. You could, after all, accomplish the same thing with an `is` statement, so why go

Further Reading on the Trick

Python Cookbook (<http://www.activestate.com/ASPN/Python/Cookbook/>) discusses alternatives to the trick (<http://www.activestate.com/ASPN/Python/Cookbook/recipe/118248.html>)

4.7. Using Functions

Python supports an interesting syntax that lets you define one-line mini-functions on the fly. Borrowed from Lisp, these so-called `is` functions c

Example 4.20. Introducing Functions

```
is
is
.
is
6
is
is
6
is
6
```

1 This is a function that accomplishes the same thing as the normal function above it. Note the abbreviated syntax here: there are no par

2 You can use a function without even assigning it to a variable. This may not be the most useful thing in the world, but it just goes to s

To generalize, a function is a function that takes any number of arguments (including optional arguments) and returns the value of a single e

Functions are a matter of style. Using them is never required; anywhere you could use them, you could define a separate normal function and

4.7.1. Real-World Functions

Here are the functions in `is`

```
is
```

Notice that this uses the simple form of the trick, which is okay, because a function is always true in a boolean context. (That doesn't mean

Also notice that you're using the `strip` function with no arguments. You've already seen it used with one or two arguments, but without any arguments

Example 4.21. `strip` With No Arguments

```
1 s = "  \n  \t  "
2 print(s)
3 print(s.strip())
4 print(s.strip().strip())
```

- 1 This is a multiline string, defined by escape characters instead of triple quotes. `\n` is a carriage return, and `\t` is a tab character.
- 2 `strip` without any arguments splits on whitespace. So three spaces, a carriage return, and a tab character are all the same.
- 3 You can normalize whitespace by splitting a string with `split` and then rejoining it with `join` using a single space as a delimiter. This is what the

So what is the `strip` function actually doing with these `strip` functions, `split`, and `join` tricks?

```
strip is now a function, but which function it is depends on the value of the variable. If strip is true, strip will collapse whitespace; otherwise, strip will return i
```

To do this in a less robust language, like Visual Basic, you would probably create a function that took a string and a `strip` argument and used an `if` statement

Further Reading on `strip` Functions

- Python Knowledge Base (<http://www.faqs.com/knowledge-base/index.phtml/fid/199/>) discusses using `strip` to call functions indirectly (`strip`)
- Python Tutorial* (<http://www.python.org/doc/current/tut/tut.html>) shows how to access outside variables from inside a `strip` function (`strip`)
- The Whole Python FAQ* (<http://www.python.org/doc/FAQ.html>) has examples of obfuscated one-liners using `strip` (`strip`) <http://www.python.org/doc/FAQ.html>

4.8. Putting It All Together

The last line of code, the only one you haven't deconstructed yet, is the one that does all the work. But by now the work is easy, because even

This is the meat of `strip`

```
strip = [method for method in dir(object) if method.startswith('strip')]
strip = [method for method in strip if method != 'strip']
strip = [method for method in strip if method != 'strip']
```

Note that this is one command, split over multiple lines, but it doesn't use the line continuation character (`\`). Remember when I said that some

Now, let's take it from the end and work backwards. The

```
strip
```

shows that this is a list comprehension. As you know, `strip` is a list of all the methods you care about in `strip`. So you're looping through that list with

Example 4.22. Getting a `strip` Dynamically

```
strip = [method for method in dir(object) if method.startswith('strip')]
strip = [method for method in strip if method != 'strip']
```



- ❶ In the function, obj is the object you're getting help on, passed in as an argument.
- ❷ As you're looping through methods, method is the name of the current method.
- ❸ Using the getattr function, you're getting a reference to the method in the module.
- ❹ Now, printing the actual object of the method is easy.

The next piece of the puzzle is the use of str around the getattr. As you may recall, str is a built-in function that coerces data into a string. But a str always

Example 4.23. Why Use str on a Method



- ❶ You can easily define a function that has no args so its attribute is None. Confusingly, if you evaluate the attribute directly, the Python IDE prints None.
- ❷ You can verify that the value of the attribute is actually None by comparing it directly.
- ❸ The function takes the null value and returns a string representation of it, 'None'.

In SQL, you must use IS NULL instead of = to compare a null value. In Python, you can use either == or is, but is is faster.

Now that you are guaranteed to have a string, you can pass the string to ljust which you have already defined as a function that either does or does not

Stepping back even further, you see that you're using string formatting again to concatenate the return value of ljust with the return value of getattr.

Example 4.24. Introducing ljust



- ❶ ljust pads the string with spaces to the given length. This is what the function uses to make two columns of output and line up all the method names in the
- ❷ If the given length is smaller than the length of the string, ljust will simply return the string unchanged. It never truncates the string.

You're almost finished. Given the padded method name from the ljust method and the (possibly collapsed) object from the call to getattr, you concatenate the

Example 4.25. Printing a List

```
l = ['a', 'b', 'c']  
print(l)
```

❶ This is also a useful debugging trick when you're working with lists. And in Python, you're always working with lists.

That's the last piece of the puzzle. You should now understand this code.

```
l = ['a', 'b', 'c']  
print(l)
```

4.9. Summary

The program and its output should now make perfect sense.

```
l = ['a', 'b', 'c']  
print(l)
```

Here is the output of

```
['a', 'b', 'c']
```

Before diving into the next chapter, make sure you're comfortable doing all of these things:

- Defining and calling functions with optional and named arguments.

Using `str()` to coerce any arbitrary value into a string representation.

Using `getattr()` to get references to functions and other attributes dynamically.

Extending the list comprehension syntax to do list filtering.

Recognizing the `isinstance()` trick and using it safely.

Defining `lambda` functions.

Assigning functions to variables and calling the function by referencing the variable. I can't emphasize this enough, because this mo



5.2. Importing Modules Using `from`

Python has two ways of importing modules. Both are useful, and you should know when to use each. One way, `import`, you've already seen in Section 5.1.

Here is the basic `from` syntax:

```
from module import *
```

This is similar to the `import` syntax that you know and love, but with an important difference: the attributes and methods of the imported module are available directly in the current namespace.

In Python is like in Perl; in Python is like in Perl.

In Python is like In Java; In Python is like In Java.

Example 5.2. `from` vs. `import`



- ❶ The `module` contains no methods; it just has attributes for each Python object type. Note that the attribute, `module` must be qualified by the `module`
- ❷ `module` by itself has not been defined in this namespace; it exists only in the context of `module`
- ❸ This syntax imports the attribute `module` from the `module` directly into the local namespace.
- ❹ Now `module` can be accessed directly, without reference to `module`

When should you use `from`

- If you will be accessing attributes and methods often and don't want to type the module name over and over, use `from`
- If you want to selectively import some attributes and methods but not others, use `import`
- If the module contains attributes or functions with the same name as ones in your module, you must use `import` to avoid name conflicts.

Other than that, it's just a matter of style, and you will see Python code written both ways.

Use `import` sparingly, because it makes it difficult to determine where a particular function or attribute came from, and that makes debugging and re-

Further Reading on Module Importing Techniques

- `eff-bot` (<http://www.effbot.org/guides/>) has more to say on `from` vs. `import` (<http://www.effbot.org/guides/import-confusion.htm>).
- *Python Tutorial* (<http://www.python.org/doc/current/tut/tut.html>) discusses advanced import techniques, including `from` (<http://www.pytho>

5.3. Defining Classes

Python is fully object-oriented: you can define your own classes, inherit from your own or built-in classes, and instantiate the classes you've

Defining a class in Python is simple. As with functions, there is no separate interface definition. Just define the class and start coding. A Python

Example 5.3. The Simplest Python Class



- ❶ The name of this class is `Class` and it doesn't inherit from any other class. Class names are usually capitalized, `Class` but this is only a convention
- ❷ This class doesn't define any methods or attributes, but syntactically, there needs to be something in the definition, so you use `pass` This is
- ❸ You probably guessed this, but everything in a class is indented, just like the code within a function, `if` statement, `try` block, and so forth. The

The statement in Python is like an empty set of braces ({} in Java or C).

Of course, realistically, most classes will be inherited from other classes, and they will define their own class methods and attributes. But as you

Example 5.4. Defining the Class

In Python, the ancestor of a class is simply listed in parentheses immediately after the class name. So the class is inherited from the

In Python, the ancestor of a class is simply listed in parentheses immediately after the class name. There is no special keyword like in Java.

Python supports multiple inheritance. In the parentheses following the class name, you can list as many ancestor classes as you like, separated

5.3.1. Initializing and Coding Classes

This example shows the initialization of the class using the method.

Example 5.5. Initializing the Class

Classes can (and should) have too, just like modules and functions. is called immediately after an instance of the class is created. It would be tempting but incorrect to call this the constructor of the class. The first argument of every class method, including this is always a reference to the current instance of the class. By convention, this argument is named self. Methods can take any number of arguments, and just like functions, the arguments can be defined with default values, making them optional.

By convention, the first argument of any Python class method (the reference to the current instance) is called self. This argument fills the role of

Example 5.6. Coding the Class

Some pseudo-object-oriented languages like Powerbuilder have a concept of "extending" constructors and other events, where the ancestor class's constructor is called first. I told you that this class acts like a dictionary, and here is the first sign of it. You're assigning the argument self the value of this object. Note that the method never returns a value.

5.3.2. Knowing When to Use `self` and `cls`

When defining your class methods, you *must* explicitly list `self` as the first argument for each method, including `__init__`. When you call a method of an instance, Python automatically passes `self` as the first argument. Whew. I realize that's a lot to absorb, but you'll get the hang of it. All Python classes work the same way, so once you learn one, you've learned them all.

Instance methods are optional, but when you define one, you must remember to explicitly call the ancestor's `__init__` method (if it defines one). This is more of a convention than a requirement.

Further Reading on Python Classes

Learning to Program (<http://www.freenetpages.co.uk/hp/alan.gauld/>) has a gentler introduction to classes (<http://www.freenetpages.co.uk/hp/alan.gauld/python/learning-to-program/classes.html>).

How to Think Like a Computer Scientist (<http://www.ibiblio.org/obp/thinkCSPy/>) shows how to use classes to model compound data structures.

Python Tutorial (<http://www.python.org/doc/current/tut/tut.html>) has an in-depth look at classes, namespaces, and inheritance (<http://www.python.org/doc/current/tut/tut.html#inheritance>).

Python Knowledge Base (<http://www.faqs.com/knowledge-base/index.phtml/fid/199/>) answers common questions about classes (<http://www.faqs.com/knowledge-base/index.phtml/fid/199/>).

5.4. Instantiating Classes

Instantiating classes in Python is straightforward. To instantiate a class, simply call the class as if it were a function, passing the arguments that you want to pass to the class's `__init__` method.

Example 5.7. Creating a Class Instance

```
1 class Point:
2     def __init__(self, x, y):
3         self.x = x
4         self.y = y
5
6 p = Point(3, 4)
7
8 print(p.x, p.y)
```

- 1 You are creating an instance of the `Point` class (defined in the `__main__` module) and assigning the newly created instance to the variable `p`. You are passing the arguments `3` and `4` to the class's `__init__` method.
- 2 Every class instance has a built-in attribute, `__class__`, which is the object's class. (Note that the representation of this includes the physical address of the class object.)
- 3 You can access the instance's `__class__` just as with a function or a module. All instances of a class share the same `__class__` attribute.
- 4 Remember when the `__init__` method assigned its `self` argument to `self`. Well, here's the result. The arguments you pass when you create the class instance are used to initialize the instance's attributes.

In Python, simply call a class as if it were a function to create a new instance of the class. There is no explicit `new` operator like C++ or Java.

5.4.1. Garbage Collection

If creating new instances is easy, destroying them is even easier. In general, there is no need to explicitly free instances, because they are freed automatically when they are no longer referenced.

Example 5.8. Trying to Implement a Memory Leak

```
1 def create_instance():
2     obj = Point(1, 1)
3     return obj
4
5 obj = create_instance()
6
7 print(obj)
```

- 1 Every time the `create_instance` function is called, you are creating an instance of `Point` and assigning it to the variable `obj`, which is a local variable within the function's scope.

② No matter how many times you call the `__del__` function, it will never leak memory, because every time, Python will destroy the newly created object.

The technical term for this form of garbage collection is "reference counting". Python keeps a list of references to every instance created. In the current version of Python, this is done by reference counting.

In previous versions of Python, there were situations where reference counting failed, and Python couldn't clean up after you. If you created a large object and then deleted it, Python would sometimes not free the memory.

As a former philosophy major, it disturbs me to think that things disappear when no one is looking at them, but that's exactly what happens in Python.

Further Reading on Garbage Collection

Python Library Reference (<http://www.python.org/doc/current/lib/>) summarizes built-in attributes like `__del__` (<http://www.python.org/doc/current/lib/attributes.html>).

Python Library Reference (<http://www.python.org/doc/current/lib/>) documents the `gc` module (<http://www.python.org/doc/current/lib/module-gc.html>).

5.5. Exploring `__del__`: A Wrapper Class

As you've seen, `__del__` is a class that acts like a dictionary. To explore this further, let's look at the `__del__` class in the `__del__` module, which is the ancestor of the `__del__` class.

In the ActivePython IDE on Windows, you can quickly open any module in your library path by selecting File->Locate... (**Ctrl-L**).

Example 5.9. Defining the `__del__` Class

- ① Note that `__del__` is a base class, not inherited from any other class.
- ② This is the `__del__` method that you overrode in the `__del__` class. Note that the argument list in this ancestor class is different than the descendant. The `__del__` method in the `__del__` class has no arguments, while the `__del__` method in the `__del__` class has one argument, `self`.
- ③ Python supports data attributes (called "instance variables" in Java and Powerbuilder, and "member variables" in C++). Data attributes are attributes that are not methods. In Python, data attributes are attributes that are not methods. In Python, data attributes are attributes that are not methods.
- ④ The `__del__` method is a dictionary duplicator: it copies all the keys and values from one dictionary to another. This does *not* clear the target dictionary.
- ⑤ This is a syntax you may not have seen before (I haven't used it in the examples in this book). It's an `__del__` statement, but instead of having a colon, it has a semicolon.

Java and Powerbuilder support function overloading by argument list, *i.e.* one class can have multiple methods with the same name but a different number of arguments.

Guido, the original author of Python, explains method overriding this way: "Derived classes may override methods of their base classes. Because Python uses a dynamic (not static) class resolution scheme, overriding is done automatically."

Always assign an initial value to `__del__` of an instance's data attributes in the `__del__` method. It will save you hours of debugging later, tracking down `__del__` errors.

Example 5.10. Normal Methods

- ①
- ②
- ③
- ④
- ⑤
- ⑥
- ⑦
- ⑧
- ⑨
- ⑩

- ❶ is a normal class method; it is publicly available to be called by anyone at any time. Notice that like all class methods, has its first parameter.
- ❷ The method of a real dictionary returns a new dictionary that is an exact duplicate of the original (all the same key-value pairs). But
- ❸ You use the attribute to see if is a if so, you're golden, because you know how to copy a just create a new and give it the real dictionary.
- ❹ If is not then must be some subclass of like maybe in which case life gets trickier. Doesn't know how to make an exact copy of
- ❺ The rest of the methods are straightforward, redirecting the calls to the built-in methods on

In versions of Python prior to 2.2.7 you could not directly subclass built-in datatypes like strings, lists, and dictionaries. To compensate for this

In Python, you can inherit directly from the built-in datatype, as shown in this example. There are three differences here compared to the

Example 5.11. Inheriting Directly from Built-In Datatype

- ❶ The first difference is that you don't need to import the module, since is a built-in datatype and is always available. The second is that
- ❷ The third difference is subtle but important. Because of the way works internally, it requires you to manually call its method to properly

Further Reading on

Python Library Reference (<http://www.python.org/doc/current/lib/>) documents the module (<http://www.python.org/doc/current/lib/>)

5.6. Special Class Methods

In addition to normal class methods, there are a number of special methods that Python classes can define. Instead of being called directly by

As you saw in the previous section, normal methods go a long way towards wrapping a dictionary in a class. But normal methods alone are not

5.6.1. Getting and Setting Items

Example 5.12. The Special Method

- ❶ The special method looks simple enough. Like the normal methods and it just redirects to the dictionary to return its value. But here's
- ❷ This looks just like the syntax you would use to get a dictionary value, and in fact it returns the value you would expect. But here's the

Of course, Python has a special method to go along with as shown in the next example.

Example 5.13. The Special Method



- ❶ Like the `__getitem__` method, `__setitem__` simply redirects to the real dictionary `__dict__` to do its work. And like `__getitem__` you wouldn't ordinarily call it directly like this; Python does it for you.
- ❷ This looks like regular dictionary syntax, except of course that `__setitem__` is really a class that's trying very hard to masquerade as a dictionary, and `__setitem__` is a special class method because it gets called for you, but it's still a class method. Just as easily as the `__getitem__` method was defined in `Dict` you can redefine it in a subclass.

This concept is the basis of the entire framework you're studying in this chapter. Each file type can have a handler class that knows how to get and set data.

For example, `MP3File` is a descendant of `File`. When an `MP3File` is set, it doesn't just set the `__dict__` key (like the ancestor `File` does); it also looks in the file itself for MP3 metadata.

Example 5.14. Overriding `__setitem__` in `MP3File`



- ❶ Notice that this `__setitem__` method is defined exactly the same way as the ancestor method. This is important, since Python will be calling the method on the `MP3File` object, not the `File` object.
- ❷ Here's the crux of the entire `MP3File` class: if you're assigning a value to the `__dict__` key, you want to do something extra.
- ❸ The extra processing you do for `__dict__` is encapsulated in the `__setitem__` method. This is another class method defined in `MP3File` and when you call it, you qualify the name with `MP3File.__setitem__(self, key, value)`.
- ❹ After doing this extra processing, you want to call the ancestor method. Remember that this is never done for you in Python; you must do it yourself.

When accessing data attributes within a class, you need to qualify the attribute name: `MP3File.__dict__`. When calling other methods within a class, you need to qualify the method name: `MP3File.__setitem__(self, key, value)`.

Example 5.15. Setting an `MP3File`



- ❶ First, you create an instance of `MP3File` without passing it a filename. (You can get away with this because the `__init__` argument of the `__init__` method is optional.)

- ② Now the real fun begins. Setting the `key` of `triggers` the `__method` on `not` which notices that you're setting the `key` with a real value
- ③ Modifying the `key` will go through the same process again: Python calls `__` which calls `__` which sets all the other keys.

5.7. Advanced Special Class Methods

Python has more special methods than just `__` and `__`. Some of them let you emulate functionality that you may not even know about.

This example shows some of the other special methods in `__`

Example 5.16. More Special Methods in `__`



- ① `__` is a special method that is called when you call `__`. The `__` function is a built-in function that returns a string representation of an object. It
- ② `__` is called when you compare class instances. In general, you can compare any two Python objects, not just class instances, by using `==`
- ③ `__` is called when you call `__`. The `__` function is a built-in function that returns the length of an object. It works on any object that could reasonably
- ④ `__` is called when you call `__` which you may remember as the way to delete individual items from a dictionary. When you use `__` on a class

In Java, you determine whether `two` string variables reference the same physical memory location by using `==`. This is called *object identity*, and

At this point, you may be thinking, "All this work just to do something in a class that I can do with a built-in datatype." And it's true that life

Special methods mean that *any class* can store key/value pairs like a dictionary, just by defining the `__` method. *Any class* can act like a sequence

While other object-oriented languages only let you define the physical model of an object ("this object has a `__` method"), Python's special class

Python has a lot of other special methods. There's a whole set of them that let classes act like numbers, allowing you to add, subtract, and do

Further Reading on Special Class Methods

Python Reference Manual (<http://www.python.org/doc/current/ref/>) documents all the special class methods (<http://www.python.org>)

5.8. Introducing Class Attributes

You already know about data attributes, which are variables owned by a specific instance of a class. Python also supports class attributes, which

Example 5.17. Introducing Class Attributes





- ❶ Is the class itself, not any particular instance of the class.
- ❷ Is a class attribute: literally, an attribute of the class. It is available before creating any instances of the class.
- ❸ Class attributes are available both through direct reference to the class and through any instance of the class.

In Java, both static variables (called class attributes in Python) and instance variables (called data attributes in Python) are defined immediately.

Class attributes can be used as class-level constants (which is how you use them in `❷`) but they are not really constants. You can also change them.

There are no constants in Python. Everything can be changed if you try hard enough. This fits with one of the core principles of Python: bad

Example 5.18. Modifying Class Attributes



5
2
2
2

- 1 is a class attribute of the class.
- 2 is a built-in attribute of every class instance (of every class). It is a reference to the class that is an instance of (in this case, the class).
- 3 Because is a class attribute, it is available through direct reference to the class, before you have created any instances of the class.
- 4 Creating an instance of the class calls the method, which increments the class attribute by 1. This affects the class itself, not just the new instance.
- 5 Creating a second instance will increment the class attribute again. Notice how the class attribute is shared by the class and all instances.

5.9. Private Functions

Like most languages, Python has the concept of private elements:

- Private functions, which can't be called from outside their module.
- Private class methods, which can't be called from outside their class.
- Private attributes, which can't be accessed from outside their class.

Unlike in most languages, whether a Python function, method, or attribute is private or public is determined entirely by its name.

If the name of a Python function, class method, or attribute starts with (but doesn't end with) two underscores, it's private; everything else is public.

In this example there are two methods: `__del__` and `__init__`. As you have already discussed, `__del__` is a special method; normally, you would call it indirectly by using the dictionary `__dict__`.

In Python, all special methods (like `__del__`) and built-in attributes (like `__dict__`) follow a standard naming convention: they both start with and end with two underscores.

Example 5.19. Trying to Call a Private Method

5
2
2
2

- 1 If you try to call a private method, Python will raise a slightly misleading exception, saying that the method does not exist. Of course it does exist, but it's private.

Further Reading on Private Functions

Python Tutorial (<http://www.python.org/doc/current/tut/tut.html>) discusses the inner workings of private variables (<http://www.python.org/doc/current/tut/tut.html>).

5.10. Summary

That's it for the hard-core object trickery. You'll see a real-world application of special class methods in Chapter 12, which uses `__new__` to create a new class.

The next chapter will continue using this code sample to explore other Python concepts, such as exceptions, file objects, and loops.

Before diving into the next chapter, make sure you're comfortable doing all of these things:

Importing modules using either `import` or `from`

Defining and instantiating classes

Defining `__init__` methods and other special class methods, and understanding when they are called

Subclassing `dict` to define classes that act like dictionaries

Defining data attributes and class attributes, and understanding the differences between them

Defining private attributes and methods

Chapter 6. Exceptions and File Handling

In this chapter, you will dive into exceptions, file objects, loops, and the `os` and `sys` modules. If you've used exceptions in another programming language, you'll find them familiar.

6.1. Handling Exceptions

Like many other programming languages, Python has exception handling via `try` blocks.

Python uses `try` to handle exceptions and `raise` to generate them. Java and C++ use `try` to handle exceptions, and `throw` to generate them.

Exceptions are everywhere in Python. Virtually every module in the standard Python library uses them, and Python itself will raise them in a number of places.

- Accessing a non-existent dictionary key will raise a `KeyError`.
- Searching a list for a non-existent value will raise a `ValueError`.
- Calling a non-existent method will raise an `AttributeError`.
- Referencing a non-existent variable will raise a `NameError`.
- Mixing datatypes without coercion will raise a `TypeError`.

In each of these cases, you were simply playing around in the Python IDE: an error occurred, the exception was printed (depending on your IDE settings).

An exception doesn't need result in a complete program crash, though. Exceptions, when raised, can be *handled*. Sometimes an exception is raised, but the program continues to run.

Example 6.1. Opening a Non-Existent File

```
1 # Open a file that doesn't exist.
2 f = open('nonexistent.txt')
3 # This line will raise an exception.
4 # The following line will never be executed.
5 print('This line will never be executed.')
6 # The following line will always be executed.
7 print('This line will always be executed.')
```

- Using the built-in `open` function, you can try to open a file for reading (more on `open` in the next section). But the file doesn't exist, so this raises an exception.
- You're trying to open the same non-existent file, but this time you're doing it within a `try` block.
- When the `open` method raises an exception, you're ready for it. The `except` line catches the exception and executes your own block of code, which prints a message.
- Once an exception has been handled, processing continues normally on the first line after the `try` block. Note that this line will always print.

Exceptions may seem unfriendly (after all, if you don't catch the exception, your entire program will crash), but consider the alternative. Without exceptions, you'd have to check for errors before every operation.

6.1.1. Using Exceptions For Other Purposes

There are a lot of other uses for exceptions besides handling actual error conditions. A common use in the standard Python library is to try to do something and catch an exception if it fails.

You can also define your own exceptions by creating a class that inherits from the built-in `Exception` class, and then raise your exceptions with the `raise` statement.

The next example demonstrates how to use an exception to support platform-specific functionality. This code comes from the `os` module, a wrapper for the operating system.

Example 6.2. Supporting Platform-Specific Functionality

```
try:
    import sys
except ImportError:
    pass
else:
    # This block can have an else clause, like an if statement. If no exception is raised during the try block, the else clause is executed afterwards. In this case, the else clause is executed because no exception was raised.
```

- 1 `try` is a UNIX-specific module that provides low-level control over the input terminal. If this module is not available (because it's not on Windows), an exception is raised.
- 2 OK, you didn't have `try`, so let's try `sys` which is a Windows-specific module that provides an API to many useful functions in the Microsoft Windows operating system.
- 3 If the first two didn't work, you try to import a function from `os` which is a Mac OS-specific module that provides functions to pop up dialog boxes.
- 4 None of these platform-specific modules is available (which is possible, since Python has been ported to a lot of different platforms), so an exception is raised.
- 5 A `try` block can have an `else` clause, like an `if` statement. If no exception is raised during the `try` block, the `else` clause is executed afterwards. In this case, the `else` clause is executed because no exception was raised.

Further Reading on Exception Handling

Python Tutorial (<http://www.python.org/doc/current/tut/tut.html>) discusses defining and raising your own exceptions, and handling them.

Python Library Reference (<http://www.python.org/doc/current/lib/>) summarizes all the built-in exceptions (<http://www.python.org/doc/current/lib/exceptions.html>).

Python Library Reference (<http://www.python.org/doc/current/lib/>) documents the `getpass` (<http://www.python.org/doc/current/lib/module-getpass.html>).

Python Library Reference (<http://www.python.org/doc/current/lib/>) documents the `try` (<http://www.python.org/doc/current/lib/module-try.html>) module.

Python Reference Manual (<http://www.python.org/doc/current/ref/>) discusses the inner workings of the `try` block (<http://www.python.org/doc/current/ref/try.html>).

6.2. Working with File Objects

Python has a built-in function, `open`, for opening a file on disk. `open` returns a file object, which has methods and attributes for getting information about the file.

Example 6.3. Opening a File

```
try:
    f = open('file.txt', 'r')
except IOError:
    print('File not found')
```

- 1 The `open` method can take up to three parameters: a filename, a mode, and a buffering parameter. Only the first one, the filename, is required.
- 2 The `open` function returns an object (by now, this should not surprise you). A file object has several useful attributes.
- 3 The `mode` attribute of a file object tells you in which mode the file was opened.
- 4 The `name` attribute of a file object tells you the name of the file that the file object has open.

6.2.1. Reading Files

After you open a file, the first thing you'll want to do is read from it, as shown in the next example.

Example 6.4. Reading a File



- 1 A file object maintains state about the file it has open. The `tell()` method of a file object tells you your current position in the open file. Since the file is open at the beginning, the current position is 0.
- 2 The `seek()` method of a file object moves to another position in the open file. The second parameter specifies what the first one means; `0` means the beginning of the file.
- 3 The `seek()` method confirms that the current file position has moved.
- 4 The `read()` method reads a specified number of bytes from the open file and returns a string with the data that was read. The optional parameter is the number of bytes to read.
- 5 The `tell()` method confirms that the current position has moved. If you do the math, you'll see that after reading 128 bytes, the position has moved to 128.

6.2.2. Closing Files

Open files consume system resources, and depending on the file mode, other programs may not be able to access them. It's important to close files when you're done with them.

Example 6.5. Closing a File





- ❶ The `closed` attribute of a file object indicates whether the object has a file open or not. In this case, the file is still open (see ❶).
- ❷ To close a file, call the `close` method of the file object. This frees the lock (if any) that you were holding on the file, flushes buffered writes.
- ❸ The `closed` attribute confirms that the file is closed.
- ❹ Just because a file is closed doesn't mean that the file object ceases to exist. The variable `f` will continue to exist until it goes out of scope.
- ❺ Calling `close` on a file object whose file is already closed does *not* raise an exception; it fails silently.

6.2.3. Handling I/O Errors

Now you've seen enough to understand the file handling code in the example code from the previous chapter. This example shows how to safely

Example 6.6. File Objects in a Block

```

❶ try:
❷     f = open('example66.txt', 'r')
❸     data = f.read(128)
❹     print(data)
❺     f.close()
❻ except IOError:
❼     pass

```

- ❶ Because opening and reading files is risky and may raise an exception, all of this code is wrapped in a `try` block. (Hey, isn't standardized?)
- ❷ The `open` function may raise an `IOError` (Maybe the file doesn't exist.)
- ❸ The `read` method may raise an `IOError` (Maybe the file is smaller than 128 bytes.)
- ❹ The `print` method may raise an `IOError` (Maybe the disk has a bad sector, or it's on a network drive and the network just went down.)
- ❺ This is new: a `finally` block. Once the file has been opened successfully by the `open` function, you want to make absolutely sure that you close it, so you wrap the `close` call in a `finally` block.
- ❻ At last, you handle your `IOError` exception. This could be the `IOError` exception raised by the call to `open` or `read`. Here, you really don't care, because all you

6.2.4. Writing to Files

As you would expect, you can also write to files in much the same way that you read from them. There are two basic file modes:

- "Append" mode will add data to the end of the file.
- "write" mode will overwrite the file.

Either mode will create the file automatically if it doesn't already exist, so there's never a need for any sort of fiddly "if the log file doesn't exist" code.

Example 6.7. Writing to Files

```

❶ f = open('example67.txt', 'w')
❷ f.write('Example 6.7: Writing to Files')
❸ f.close()

```



- ❶ You start boldly by creating either the new file or overwrites the existing file, and opening the file for writing. (The second parameter
- ❷ You can add data to the newly opened file with the write method of the file object returned by open
- ❸ flush is a synonym for write. This one-liner opens the file, reads its contents, and prints them.
- ❹ You happen to know that the file exists (since you just finished writing to it), so you can open it and append to it. (The 'a' parameter means open
- ❺ As you can see, both the original line you wrote and the second line you appended are now in the file. Also note that carriage returns are not

Further Reading on File Handling

Python Tutorial (<http://www.python.org/doc/current/tut/tut.html>) discusses reading and writing files, including how to read a file on the command line. *eff-bot* (<http://www.effbot.org/guides/>) discusses efficiency and performance of various ways of reading a file (<http://www.effbot.org/guides/reading-files/>). Python Knowledge Base (<http://www.faqs.com/knowledge-base/index.phtml/fid/199/>) answers common questions about files (<http://www.faqs.com/knowledge-base/index.phtml/fid/199/>). *Python Library Reference* (<http://www.python.org/doc/current/lib/>) summarizes all the file object methods (<http://www.python.org/doc/current/lib/>).

6.3. Iterating with Loops

Like most other languages, Python has loops. The only reason you haven't seen them until now is that Python is good at so many other things.

Most other languages don't have a powerful list datatype like Python, so you end up doing a lot of manual work, specifying a start, end, and step.

Example 6.8. Introducing the Loop

```

for letter in 'abcde':
    print(letter)

```

- ❶ The syntax for a loop is similar to list comprehensions. letters is a list, and print will take the value of each element in turn, starting from the first.
- ❷ Like an if statement or any other indented block, a loop can have any number of lines of code in it.
- ❸ This is the reason you haven't seen the loop yet: you haven't needed it yet. It's amazing how often you use loops in other languages.

Doing a "normal" (by Visual Basic standards) counter loop is also simple.

Example 6.9. Simple Counters

```

for i in range(10):
    print(i)

```

```
0
1
2
3
4
5
6
7
8
9
a
b
c
d
e
```

- ❶ As you saw in Example 3.20, “Assigning Consecutive Values”, `range()` produces a list of integers, which you then loop through. I know it looks like you’re looping through a range, but you’re not. You’re looping through a list of integers.
- ❷ Don’t ever do this. This is Visual Basic-style thinking. Break out of it. Just iterate through the list, as shown in the previous example. Loops are not just for simple counters. They can iterate through all kinds of things. Here is an example of using a loop to iterate through a dictionary.

Example 6.10. Iterating Through a Dictionary

```
1 # Loop through the environment variables dictionary
2 import os
3
4 # Print the keys and values of the dictionary
5 for key, value in os.environ.items():
6     print(key, value)
7
8 # Print the keys and values of the dictionary using list comprehension
9 print([key, value] for key, value in os.environ.items())
```

- ❶ `os.environ` is a dictionary of the environment variables defined on your system. In Windows, these are your user and system variables accessible to the current process.
- ❷ `os.environ.items()` returns a list of tuples: `(key, value)`. The loop iterates through this list. The first round, it assigns `key` and `value` to `key` and `value`. In the second round, it assigns `key` and `value` to the next key and value.
- ❸ With multi-variable assignment and list comprehensions, you can replace the entire loop with a single statement. Whether you actually use this is up to you.

Now we can look at the loop in `main.py` from the sample program introduced in Chapter 5.

Example 6.11. Loop in `main.py`

```
1 # Loop through the environment variables dictionary
2 import os
3
4 # Print the keys and values of the dictionary
5 for key, value in os.environ.items():
6     print(key, value)
```



- ❶ `MP3` is a class attribute that defines the tags you're looking for in an MP3 file. Tags are stored in fixed-length fields. Once you read the last
- ❷ This looks complicated, but it's not. The structure of the `MP3` variables matches the structure of the elements of the list returned by `removes`
- ❸ Now that you've extracted all the parameters for a single MP3 tag, saving the tag data is easy. You slice `data` from `start` to `end` to get the actual data.

6.4. Using `sys`

Modules, like everything else in Python, are objects. Once imported, you can always get a reference to a module through the global dictionary

Example 6.12. Introducing `sys`



- ❶ The `sys` module contains system-level information, such as the version of Python you're running (`sys.version`) and system-level options such as t
- ❷ `sys.modules` is a dictionary containing all the modules that have ever been imported since Python was started; the key is the module name, the valu

This example demonstrates how to use `sys`

Example 6.13. Using `sys`





- ❶ As new modules are imported, they are added to `sys.modules`. This explains why importing the same module twice is very fast: Python has already imported it.
- ❷ Given the name (as a string) of any previously-imported module, you can get a reference to the module itself through the `sys.modules` dictionary.

The next example shows how to use the `__class__` attribute with the `sys.modules` dictionary to get a reference to the module in which a class is defined.

Example 6.14. The `__class__` Attribute



- ❶ Every Python class has a built-in class attribute `__class__` which is the name of the module in which the class is defined.
- ❷ Combining this with the `sys.modules` dictionary, you can get a reference to the module in which a class is defined.

Now you're ready to see how `__class__` is used in `__main__` the sample program introduced in Chapter 5. This example shows that portion of the code.

Example 6.15. `__main__`



- ❶ This is a function with two arguments; `obj` is required, but `module` is optional and defaults to the module that contains the `obj` class. This looks inefficient.
- ❷ You'll plow through this line later, after you dive into the `sys` module. For now, take it on faith that `obj.__class__` ends up as the name of a class, like `obj`.
- ❸ You already know about `sys.get_object_by_name` which gets a reference to an object by name. `sys.get_object_by_name` is a complementary function that checks whether an object has a `__class__` attribute.

Further Reading on Modules

Python Tutorial (<http://www.python.org/doc/current/tut/tut.html>) discusses exactly when and how default arguments are evaluated (`__class__` is a default argument).

Python Library Reference (<http://www.python.org/doc/current/lib/>) documents the `sys` module (<http://www.python.org/doc/current/lib/module-sys.html>).

6.5. Working with Directories

The `os` module has several functions for manipulating files and directories. Here, we're looking at handling pathnames and listing the contents of a directory.

Example 6.16. Constructing Pathnames





- ❶ is a reference to a module -- which module depends on your platform. Just as `os.path` encapsulates differences between platforms by setting `os.path.sep`.
- ❷ The function of `os.path.join` constructs a pathname out of one or more partial pathnames. In this case, it simply concatenates strings. (Note that `os.path.join` will not add a slash between two strings.)
- ❸ In this slightly less trivial case, `os.path.join` will add an extra backslash to the pathname before joining it to the filename. I was overjoyed when I discovered this.
- ❹ `os.path.expanduser` will expand a pathname that uses `~` to represent the current user's home directory. This works on any platform where users have a home directory.
- ❺ Combining these techniques, you can easily construct pathnames for directories and files under the user's home directory.

Example 6.17. Splitting Pathnames



- ❶ The `os.path.split` function splits a full pathname and returns a tuple containing the path and filename. Remember when I said you could use multi-assignment to unpack a tuple?
- ❷ You assign the return value of the `os.path.split` function into a tuple of two variables. Each variable receives the value of the corresponding element of the tuple.
- ❸ The first variable, `path`, receives the value of the first element of the tuple returned from `os.path.split`, the file path.
- ❹ The second variable, `filename`, receives the value of the second element of the tuple returned from `os.path.split`, the filename.
- ❺ `os.path.splitext` also contains a function `os.path.splitext` which splits a filename and returns a tuple containing the filename and the file extension. You use the same multi-assignment technique to unpack the tuple.

Example 6.18. Listing Directories





- ❶ The `os.listdir()` function takes a pathname and returns a list of the contents of the directory.
- ❷ It returns both files and folders, with no indication of which is which.
- ❸ You can use list filtering and the `os.path.isfile()` function of the `os.path` module to separate the files from the folders. `os.path.isfile()` takes a pathname and returns 1 if the path is a file, and 0 otherwise.
- ❹ `os.path.isdir()` also has a function which returns 1 if the path represents a directory, and 0 otherwise. You can use this to get a list of the subdirectories.

Example 6.19. Listing Directories in `os`

```

import os

# Get a list of all files and folders in the current directory
files = os.listdir('.')

# Iterate through the list and normalize the case
normalized_files = [os.path.normcase(f) for f in files]

# Iterate through the normalized list and split each filename into name and extension
filtered_files = []
for f in normalized_files:
    name, ext = os.path.splitext(f)
    # Check if the extension is in the list of file extensions you care about
    if ext in ['.txt', '.py']:
        # For each file you care about, use os.path.join to construct the full pathname of the file, and return a list of the full pathnames.
        filtered_files.append(os.path.join('.', f))

print(filtered_files)

```

- ❶ `os.listdir()` returns a list of all the files and folders in `.`
- ❷ Iterating through the list with `for`, you use `os.path.normcase()` to normalize the case according to operating system defaults. `os.path.normcase()` is a useful little function that converts a pathname to the case used by the operating system.
- ❸ Iterating through the normalized list with `for` again, you use `os.path.splitext()` to split each filename into name and extension.
- ❹ For each file, you see if the extension is in the list of file extensions you care about (`ext in ['.txt', '.py']` which was passed to the `if` function).
- ❺ For each file you care about, you use `os.path.join()` to construct the full pathname of the file, and return a list of the full pathnames.

Whenever possible, you should use the functions in `os.path` for file, directory, and path manipulations. These modules are wrappers for platform-specific functions.

There is one other way to get the contents of a directory. It's very powerful, and it uses the sort of wildcards that you may already be familiar with.

Example 6.20. Listing Directories with `glob`

```

import glob

# Get a list of all files and folders in the current directory
files = glob.glob('*.*')

# Iterate through the list and normalize the case
normalized_files = [os.path.normcase(f) for f in files]

# Iterate through the normalized list and split each filename into name and extension
filtered_files = []
for f in normalized_files:
    name, ext = os.path.splitext(f)
    # Check if the extension is in the list of file extensions you care about
    if ext in ['.txt', '.py']:
        # For each file you care about, use os.path.join to construct the full pathname of the file, and return a list of the full pathnames.
        filtered_files.append(os.path.join('.', f))

print(filtered_files)

```

- ❶ As you saw earlier, `os.listdir()` takes a directory path and lists all files and directories in that directory.
- ❷ The `os.walk()` module, on the other hand, takes a wildcard and returns the full path of all files and directories matching the wildcard. Here the wildcard is `*`.
- ❸ If you want to find all the files in a specific directory that start with "s" and end with ".mp3", you can do that too.
- ❹ Now consider this scenario: you have a `dir` directory, with several subdirectories within it, with `files` within each subdirectory. You can get a list of all files in `dir` with the following code:

Further Reading on the `os` Module

Python Knowledge Base (<http://www.faqs.com/knowledge-base/index.phtml/fid/199/>) answers questions about the `os` module (<http://www.python.org/doc/current/lib/module-os.html>)
Python Library Reference (<http://www.python.org/doc/current/lib/>) documents the `os` module (<http://www.python.org/doc/current/lib/module-os.html>)

6.6. Putting It All Together

Once again, all the dominoes are in place. You've seen how each line of code works. Now let's step back and see how it all fits together.

Example 6.21. `os.walk()`

```

import os
import sys

def walk_dir(dir, extensions):
    """Walk a directory tree and return a list of files with interesting extensions"""
    interesting_files = []
    for dirpath, dirnames, filenames in os.walk(dir):
        for filename in filenames:
            if filename.lower().endswith(extensions):
                interesting_files.append(os.path.join(dirpath, filename))
    return interesting_files

if __name__ == '__main__':
    dir = sys.argv[1]
    extensions = sys.argv[2]
    interesting_files = walk_dir(dir, extensions)
    print interesting_files

```

- ❶ `os.walk()` is the main attraction of this entire module. It takes a directory (like `dir` in my case) and a list of interesting file extensions (like `extensions` and it returns a list of tuples.
- ❷ As you saw in the previous section, this line of code gets a list of the full pathnames of all the files in `dir` that have an interesting file extension.
- ❸ Old-school Pascal programmers may be familiar with them, but most people give me a blank stare when I tell them that Python supports `os.path.splitext()`.
- ❹ Now that you've seen the `os` module, this line should make more sense. It gets the extension of the file (`os.path.splitext()` forces it to uppercase (`os.path.splitext()` slices the filename).
- ❺ Having constructed the name of the handler class that would handle this file, you check to see if that handler class actually exists in `handlers`.
- ❻ For each file in the "interesting files" list (`interesting_files`) you call `handlers` with the filename (`f`). Calling `handlers` returns a class; you don't know exactly which class

Note that `handlers` is completely generic. It doesn't know ahead of time which types of files it will be getting, or which classes are defined that could handle them.

6.7. Summary

The `handlers` program introduced in Chapter 5 should now make perfect sense.

```

import os
import sys

def walk_dir(dir, extensions):
    """Walk a directory tree and return a list of files with interesting extensions"""
    interesting_files = []
    for dirpath, dirnames, filenames in os.walk(dir):
        for filename in filenames:
            if filename.lower().endswith(extensions):
                interesting_files.append(os.path.join(dirpath, filename))
    return interesting_files

if __name__ == '__main__':
    dir = sys.argv[1]
    extensions = sys.argv[2]
    interesting_files = walk_dir(dir, extensions)
    print interesting_files

```




Before diving into the next chapter, make sure you're comfortable doing the following things:

Catching exceptions with `try`

Protecting external resources with `with`

Reading from files

Assigning multiple values at once in a `loop`

Using the `os` module for all your cross-platform file manipulation needs

Dynamically instantiating classes of unknown type by treating classes as objects and passing them around

Chapter 7. Regular Expressions

Regular expressions are a powerful and standardized way of searching, replacing, and parsing text with complex patterns of characters. If you

7.1. Diving In

Strings have methods for searching (`find` and `findall`), replacing (`replace`), and parsing (`split`) but they are limited to the simplest of cases. The search methods look

If what you're trying to do can be accomplished with string functions, you should use them. They're fast and simple and easy to read, and the

Although the regular expression syntax is tight and unlike normal code, the result can end up being *more* readable than a hand-rolled solution.

7.2. Case Study: Street Addresses

This series of examples was inspired by a real-life problem I had in my day job several years ago, when I needed to scrub and standardize street

Example 7.1. Matching at the End of a String



- 1 My goal is to standardize a street address so that `St`s always abbreviated as `St`. At first glance, I thought this was simple enough that I could
- 2 Life, unfortunately, is full of counterexamples, and I quickly discovered this one. The problem here is that `St` appears twice in the address
- 3 To solve the problem of addresses with more than one `St` substring, you could resort to something like this: only search and replace `St` at the
- 4 It's time to move up to regular expressions. In Python, all functionality related to regular expressions is contained in the `re` module.
- 5 Take a look at the first parameter: `st$`. This is a simple regular expression that matches `St` only when it occurs at the end of a string. The `$` anchors
- 6 Using the `re.sub` function, you search the string `St St` for the regular expression `st$` and replace it with `St`. This matches the `St` at the end of the string `St St`

Continuing with my story of scrubbing addresses, I soon discovered that the previous example, matching `St` at the end of the address, was not good

Example 7.2. Matching Whole Words





- ❶ What I *really* wanted was to match `\b` when it was at the end of the string *and* it was its own whole word, not a part of some larger word
- ❷ To work around the backslash plague, you can use what is called a raw string, by prefixing the string with the letter `r` This tells Python
- ❸ `*sigh*` Unfortunately, I soon found more cases that contradicted my logic. In this case, the street address contained the word `\b` as a whole
- ❹ To solve this problem, I removed the `$` character and added another `\b` Now the regular expression reads "match `\b` when it's a whole word"

7.3. Case Study: Roman Numerals

You've most likely seen Roman numerals, even if you didn't recognize them. You may have seen them in copyrights of old movies and television shows.

In Roman numerals, there are seven characters that are repeated and combined in various ways to represent numbers.

I= 1
 V= 5
 X= 10
 L= 50
 C= 100
 D= 500
 M= 1000

The following are some general rules for constructing Roman numerals:

Characters are additive. `II` is 2 and `III` is 3 `V` is 5 and `VI` is 6 (literally, "5 and 1"), `VII` is 7 and `VIII` is 8
 The tens characters (`IX` and `M`) can be repeated up to three times. At `IV` you need to subtract from the next highest fives character. You can't do `IIII`.
 Similarly, at `IX` you need to subtract from the next highest tens character: `IX` is 9 but `XI` is 11 less than `X` not `IX` since the `I` character can not be repeated.
 The fives characters can not be repeated. The number `0` is always represented as `X` never as `V` The number `1000` is always `M` never `L`
 Roman numerals are always written highest to lowest, and read left to right, so the order of the characters matters very much. `DC` is 600

7.3.1. Checking for Thousands

What would it take to validate that an arbitrary string is a valid Roman numeral? Let's take it one digit at a time. Since Roman numerals are always written highest to lowest, we can check for thousands first.

Example 7.3. Checking for Thousands



- ❶ This pattern has three parts:
`^` to match what follows only at the beginning of the string. If this were not specified, the pattern would match no matter where

`M` optionally match a single `M` character. Since this is repeated three times, you're matching anywhere from zero to three `M` characters. `$` matches what precedes only at the end of the string. When combined with the `^` character at the beginning, this means that the

- 2 The essence of the `re` module is the `re.match()` function, that takes a regular expression (`re.compile()`) and a string (`string`) to try to match against the regular expression.
- 3 `re.match()` matches because the first and second optional `M` characters match and the third `M` is ignored.
- 4 `re.match()` matches because all three `M` characters match.
- 5 `re.match()` does not match. All three `M` characters match, but then the regular expression insists on the string ending (because of the `$` character), and it doesn't.
- 6 Interestingly, an empty string also matches this regular expression, since all the `M` characters are optional.

7.3.2. Checking for Hundreds

The hundreds place is more difficult than the thousands, because there are several mutually exclusive ways it could be expressed, depending on the

- 0- C
- 0- C
- 0- C
- 0- D
- 0- M

So there are four possible patterns:

- 0
- D
- Zero to three `C` characters (zero if the hundreds place is 0)
- `D` followed by zero to three `C` characters

The last two patterns can be combined:

- an optional `D` followed by zero to three `C` characters

This example shows how to validate the hundreds place of a Roman numeral.

Example 7.4. Checking for Hundreds



- 1 This pattern starts out the same as the previous one, checking for the beginning of the string (`^`), then the thousands place (`M`). Then it has

- ② `M` matches because the first `M` matches, the second and third `M` characters are ignored, and the `D` matches (so the `C` and `D` patterns are never even evaluated).
- ③ `MD` matches because the first `M` matches, the second and third `M` characters are ignored, and the `D` pattern matches `D` each of the three `C` characters.
- ④ `MMM` matches because all three `M` characters match, and the `C` pattern matches `C` (the `D`s optional and is ignored). `MM` is the Roman numeral representing 2.
- ⑤ `MDCCC` does not match. The first `M` matches, the second and third `M` characters are ignored, and the `D` matches, but then the `C` does not match because there are only two `C`'s.
- ⑥ Interestingly, an empty string still matches this pattern, because all the `M` characters are optional and ignored, and the empty string matches the `C` and `D` patterns.

Whew! See how quickly regular expressions can get nasty? And you've only covered the thousands and hundreds places of Roman numerals!

7.4. Using the `|` Syntax

In the previous section, you were dealing with a pattern where the same character could be repeated up to three times. There is another way to do this.

Example 7.5. The Old Way: Every Character Optional

```

1 re.match('M?D?C?C?M?', 'M')
2 re.match('M?D?C?C?M?', 'MD')
3 re.match('M?D?C?C?M?', 'MDC')
4 re.match('M?D?C?C?M?', 'MDCCLXXXIX')
5 re.match('M?D?C?C?M?', 'MDCCC')
6 re.match('M?D?C?C?M?', '')
>

```

- ① This matches the start of the string, and then the first optional `M` but not the second and third `M` (but that's okay because they're optional).
- ② This matches the start of the string, and then the first and second optional `M` but not the third `M` (but that's okay because it's optional), and then the `D`.
- ③ This matches the start of the string, and then all three optional `M` and then the end of the string.
- ④ This matches the start of the string, and then all three optional `M` but then does not match the the end of the string (because there is still more to the string).

Example 7.6. The New Way: From `MDCCC` to `M{0,3}`

```

1 re.match('M{0,3}D{0,1}C{0,3}M{0,1}$', 'M')
2 re.match('M{0,3}D{0,1}C{0,3}M{0,1}$', 'MD')
3 re.match('M{0,3}D{0,1}C{0,3}M{0,1}$', 'MDC')
4 re.match('M{0,3}D{0,1}C{0,3}M{0,1}$', 'MDCCLXXXIX')
5 re.match('M{0,3}D{0,1}C{0,3}M{0,1}$', 'MDCCC')
6 re.match('M{0,3}D{0,1}C{0,3}M{0,1}$', '')
>

```

- ① This pattern says: "Match the start of the string, then anywhere from zero to three `M` characters, then the end of the string." The `0` and `3` are the minimum and maximum number of matches, respectively.
- ② This matches the start of the string, then one `M` but of a possible three, then the end of the string.
- ③ This matches the start of the string, then two `M` but of a possible three, then the end of the string.

7.6. Case study: Parsing Phone Numbers

So far you've concentrated on matching whole patterns. Either the pattern matches, or it doesn't. But regular expressions are much more powerful.

This example came from another real-world problem I encountered, again from a previous day job. The problem: parsing an American phone number.

Here are the phone numbers I needed to be able to accept:

```
000-000-0000
000-000-0000
000-000-0000
000-000-0000
000-000-0000
000-000-0000
000-000-0000
000-000-0000
000-000-0000
```

Quite a variety! In each of these cases, I need to know that the area code was `000`, the trunk was `000` and the rest of the phone number was `0000`. For those who are interested, the regular expression I used was `^000-000-0000$`.

Let's work through developing a solution for phone number parsing. This example shows the first step.

Example 7.10. Finding Numbers

```
>>> re.compile(r'^000-000-0000$')
```

- 1 Always read regular expressions from left to right. This one matches the beginning of the string, and then `000-000-0000`. Well, the `000-000-0000` means the rest of the string.
- 2 To get access to the groups that the regular expression parser remembered along the way, use the `groups()` method on the object that the `re.compile()` function returns.
- 3 This regular expression is not the final answer, because it doesn't handle a phone number with an extension on the end. For that, you'll need a different regular expression.

Example 7.11. Finding the Extension

```
>>> re.compile(r'^000-000-0000-$')
```

- 1 This regular expression is almost identical to the previous one. Just as before, you match the beginning of the string, then a `000-000-0000`, and finally a `$`.
- 2 The `groups()` method now returns a tuple of four elements, since the regular expression now defines four groups to remember.
- 3 Unfortunately, this regular expression is not the final answer either, because it assumes that the different parts of the phone number are separated by hyphens.
- 4 Oops! Not only does this regular expression not do everything you want, it's actually a step backwards, because now you can't parse phone numbers with extensions.



- 1 Other than being spread out over multiple lines, this is exactly the same regular expression as the last step, so it's no surprise that it passes.
- 2 Final sanity check. Yes, this still works. You're done.

Further Reading on Regular Expressions

Regular Expression HOWTO (<http://py-howto.sourceforge.net/regex/regex.html>) teaches about regular expressions and how to use them.
Python Library Reference (<http://www.python.org/doc/current/lib/>) summarizes the `re` module (<http://www.python.org/doc/current/lib/>)

7.7. Summary

This is just the tiniest tip of the iceberg of what regular expressions can do. In other words, even though you're completely overwhelmed by the possibilities, you should know a few things.

You should now be familiar with the following techniques:

- `^` matches the beginning of a string.
- `$` matches the end of a string.
- `\b` matches a word boundary.
- `\d` matches any numeric digit.
- `\D` matches any non-numeric character.
- `X?` matches an optional `X` character (in other words, it matches an `X` zero or one times).
- `X*` matches `X` zero or more times.
- `X+` matches `X` one or more times.
- `X{n,m}` matches an `X` character at least `n` times, but not more than `m` times.
- `X|Y` matches either `X` or `Y`.
- `(...)` in general is a *remembered group*. You can get the value of what matched by using the `group()` method of the object returned by `re.compile(...)`.

Regular expressions are extremely powerful, but they are not the correct solution for every problem. You should learn enough about them to know when to use them.

Some people, when confronted with a problem, think "I know, I'll use regular expressions." Now they have two problems.
--Jamie Zawinski, in *Regexes for Dummies*

Entity reference	An escaped character referenced by its decimal or hexadecimal equivalent, like <code>&#x27;</code> When found, <code>&</code> calls <code>'</code> with the text of the de
Comment	An HTML comment, enclosed in <code><!--</code> When found, <code><</code> calls <code>--></code> with the body of the comment.
Processing instruction	An HTML processing instruction, enclosed in <code><?></code> When found, <code><</code> calls <code>></code> with the body of the processing instruction.
Declaration	An HTML declaration, such as a <code><!DOCTYPE html></code> enclosed in <code><!--</code> When found, <code><</code> calls <code>--></code> with the body of the declaration.
Text data	A block of text. Anything that doesn't fit into the other 7 categories. When found, <code><</code> calls <code>></code> with the text.

Python 2.0 had a bug where `HTMLParser` would not recognize declarations at all (`HTMLParser` would never be called), which meant that `HTMLParser` were silently ignored. This comes with a test suite to illustrate this. You can run `test_htmlparser.py` passing the name of an HTML file on the command line, and it will print out the tags a

In the ActivePython IDE on Windows, you can specify command line arguments in the "Run script" dialog. Separate multiple arguments with

Example 8.4. Sample test of `HTMLParser`

Here is a snippet from the table of contents of the HTML version of this book. Of course your paths may vary. (If you haven't downloaded th

```

<h1>Python</h1>
<h2>Getting Started</h2>
<h3>Installing Python</h3>
<h3>Installing the Standard Library</h3>
<h3>Installing the IDE</h3>
<h3>Installing the Test Suite</h3>
<h3>Installing the Documentation</h3>
<h3>Installing the Examples</h3>
<h3>Installing the Source Code</h3>
<h3>Installing the Test Suite</h3>
<h3>Installing the Documentation</h3>
<h3>Installing the Examples</h3>
<h3>Installing the Source Code</h3>

```

Running this through the test suite of `HTMLParser` yields this output:

```

<h1>Python</h1>
<h2>Getting Started</h2>
<h3>Installing Python</h3>
<h3>Installing the Standard Library</h3>
<h3>Installing the IDE</h3>
<h3>Installing the Test Suite</h3>
<h3>Installing the Documentation</h3>
<h3>Installing the Examples</h3>
<h3>Installing the Source Code</h3>

```

Here's the roadmap for the rest of the chapter:

Subclass `HTMLParser` to create classes that extract interesting data out of HTML documents.

Subclass `HTMLParser` to create `HTMLReconstructor` which overrides all 8 handler methods and uses them to reconstruct the original HTML from the pieces.

Subclass `HTMLReconstructor` to create `HTMLReconstructor` which adds some methods to process specific HTML tags specially, and overrides the `unescape` method to provide a frame.

Subclass `HTMLReconstructor` to create classes that define text processing rules used by `HTMLReconstructor`.

Write a test suite that grabs a real web page from `urllib` and processes it.

Along the way, you'll also learn about `collections` and dictionary-based string formatting.

8.3. Extracting data from HTML documents

To extract data from HTML documents, subclass the `HTMLParser` class and define methods for each tag or entity you want to capture.

The first step to extracting data from an HTML document is getting some HTML. If you have some HTML lying around on your hard drive,

Example 8.5. Introducing `urllib`



- 1 The `urllib` module is part of the standard Python library. It contains functions for getting information about and actually retrieving data from the Internet.
- 2 The simplest use of `urllib` is to retrieve the entire text of a web page using the `urlopen` function. Opening a URL is similar to opening a file. The returned object is a file-like object.
- 3 The simplest thing to do with the file-like object returned by `urlopen` is to read the entire HTML of the web page into a single string. This is done by calling the `read` method.
- 4 When you're done with the object, make sure to `close` it, just like a normal file object.
- 5 You now have the complete HTML of the home page of `http://www.python.org` in a string, and you're ready to parse it.

Example 8.6. Introducing `HTMLParser`

If you have not already done so, you can download this and other examples (<http://diveintopython.org/download/diveintopython-examples-5>).

```
1 from html.parser import HTMLParser
2
3 class SimpleHTMLParser(HTMLParser):
4     def handle_starttag(self, tag, attrs):
5         print '<{}>'.format(tag)
6
7     def handle_endtag(self, tag):
8         print '</{}>'.format(tag)
9
10    def handle_data(self, data):
11        print data
12
13 parser = SimpleHTMLParser()
14 parser.feed('<p>Dive Into Python</p>')
```

- 1 `handle_starttag` is called by the `feed` method of `HTMLParser` and it can also be called manually once an instance of the parser has been created. So if you need to do anything special when a tag is encountered, you can override this method.
- 2 `handle_endtag` is called by `HTMLParser` whenever it finds an end tag. The tag may contain an `id` attribute, and/or other attributes, like `class` or `style`. The `attrs` parameter is a list of tuples containing the attribute name and value.
- 3 You can find out whether this tag has an attribute with a simple multi-variable list comprehension.
- 4 String comparisons like `id` are always case-sensitive, but that's safe in this case, because `HTMLParser` converts attribute names to lowercase while building the `attrs` list.

Example 8.7. Using `HTMLParser` to Find Links

```
1 from html.parser import HTMLParser
2
3 class HTMLLinkParser(HTMLParser):
4     def handle_starttag(self, tag, attrs):
5         if tag == 'a':
6             href = None
7             for attr in attrs:
8                 if attr[0] == 'href':
9                     href = attr[1]
10            if href:
11                print href
12
13 parser = HTMLLinkParser()
14 parser.feed('<p>Dive Into Python</p>')
```

- 1 `HTMLLinkParser` takes over the `HTMLParser` class and inherits its `feed` method.
- 2 Like files, you should close your URL objects as soon as you're done with them.
- 3 You should close your parser object, too, but for a different reason. You've read all the data and fed it to the parser, but the `feed` method isn't guaranteed to have finished parsing the HTML document.
- 4 Once the parser is closed, the parsing is complete, and `HTMLLinkParser` contains a list of all the linked URLs in the HTML document. (Your output may look like this.)

8.4. Introducing `HTMLParser`

`HTMLParser` doesn't produce anything by itself. It parses and parses and parses, and it calls a method for each interesting thing it finds, but the methods do the work. `HTMLParser` subclasses `BaseParser` and provides all 8 essential handler methods: `handle_start_tag`, `handle_end_tag`, `handle_start_entity`, `handle_end_entity`, `handle_charref`, `handle_entityref`, `handle_comment`, and `handle_text`.

Example 8.8. Introducing `HTMLParser`

```
class HTMLParser(BaseParser):
    def __init__(self, *args, **kwargs):
        self.__data = []
        super().__init__(*args, **kwargs)

    def handle_start_tag(self, tag, attrs):
        # Called by handle_data()
        self.__data.append("<{}>".format(tag))

    def handle_end_tag(self, tag):
        # Reconstructing end tags is much simpler; just take the tag name and wrap it in the </> brackets.
        self.__data.append("</{}>".format(tag))

    def handle_start_entity(self, entity):
        # When <entity> is found, it calls handle_entityref() with the bare reference. If the HTML document contains the reference <#entity> it will be handled by handle_entityref(). Reconstructing the original entity reference requires wrapping the entity name in &#entity;
        self.__data.append("&#{};".format(entity))

    def handle_end_entity(self, entity):
        # Entity references are similar to character references, but without the hash mark. Reconstructing the original entity reference requires wrapping the entity name in &#entity;
        self.__data.append("&#{};".format(entity))

    def handle_charref(self, ref):
        # Blocks of text are simply appended to the data list unaltered.
        self.__data.append("&#{};".format(ref))

    def handle_entityref(self, entity):
        # HTML comments are wrapped in <!--> characters.
        self.__data.append("<!--{}-->".format(entity))

    def handle_comment(self, text):
        # Processing instructions are wrapped in <?> characters.
        self.__data.append("<?{}>".format(text))

    def handle_text(self, text):
        self.__data.append(text)
```

- 1 Called by `handle_data()` initializes `self.__data` as an empty list before calling the ancestor method. `self.__data` is a data attribute which will hold the pieces of the HTML document.
- 2 Since `HTMLParser` does not define any methods for specific tags (like the `handle_start_tag` method in `BaseParser`), it will call `handle_start_tag` for every start tag. This method takes the tag (`tag`) and appends it to `self.__data`.
- 3 Reconstructing end tags is much simpler; just take the tag name and wrap it in the `</>` brackets.
- 4 When `HTMLParser` finds a character reference, it calls `handle_charref` with the bare reference. If the HTML document contains the reference `&#entity;` it will be handled by `handle_entityref`. Reconstructing the original entity reference requires wrapping the entity name in `&#entity;`.
- 5 Entity references are similar to character references, but without the hash mark. Reconstructing the original entity reference requires wrapping the entity name in `&#entity;`.
- 6 Blocks of text are simply appended to the data list unaltered.
- 7 HTML comments are wrapped in `<!-->` characters.
- 8 Processing instructions are wrapped in `<?>` characters.

The HTML specification requires that all non-HTML (like client-side JavaScript) must be enclosed in HTML comments, but not all web pages follow this rule.

Example 8.9. Output



- 1 This is the one method in `HTMLParser` that is never called by the ancestor `BaseHTMLParser`. Since the other handler methods store their reconstructed HTML in `self.fragments`, this method is the only one that doesn't.
- 2 If you prefer, you could use the `unescape` method of the `HTMLParser` module instead: `parser.unescape(html)`

Further reading

W3C (<http://www.w3.org/>) discusses character and entity references (<http://www.w3.org/TR/REC-html40/charset.html#entities>).
Python Library Reference (<http://www.python.org/doc/current/lib/>) confirms your suspicions that the `HTMLParser` module (<http://www.python.org/doc/current/lib/module-urllib.html>)

8.5. and

Let's digress from HTML processing for a minute and talk about how Python handles variables. Python has two built-in functions, `and` and `or`, which

Remember `and`. You first saw it here:



No, wait, you can't learn about `and` yet. First, you need to learn about namespaces. This is dry stuff, but it's important, so pay attention.

Python uses what are called namespaces to keep track of variables. A namespace is just like a dictionary where the keys are names of variables.

At any particular point in a Python program, there are several namespaces available. Each function has its own namespace, called the local namespace.

When a line of code asks for the value of a variable `x`, Python will search for that variable in all the available namespaces, in order:

- local namespace - specific to the current function or class method. If the function defines a local variable `x` or has an argument `x`, Python will use that value.
- global namespace - specific to the current module. If the module has defined a variable, function, or class called `x`, Python will use that value.
- built-in namespace - global to all modules. As a last resort, Python will assume that `x` is the name of built-in function or variable.

If Python doesn't find `x` in any of these namespaces, it gives up and raises a `NameError` with the message `NameError: name 'x' is not defined`, which you saw back in Example 3.18, "Reference Error".

Python 2.2 introduced a subtle but important change that affects the namespace search order: nested scopes. In versions of Python prior to 2.2, the search order was:



Are you confused yet? Don't despair! This is really cool, I promise. Like many things in Python, namespaces are *directly accessible at run-time*.

Example 8.10. Introducing



- 1 The function `vars` has two variables in its local namespace: `args` whose value is passed in to the function, and `kwargs` which is defined within the function.
- 2 `vars` returns a dictionary of name/value pairs. The keys of this dictionary are the names of the variables as strings; the values of the dictionary are the values of the variables.
- 3 Remember, Python has dynamic typing, so you could just as easily pass a string in for `args` the function (and the call to `vars` would still work).

What does `vars` do for the module namespace? `vars` is a function defined in the `builtins` module, so it is available in the global namespace. It includes all the functions and classes defined in the module.

Remember the difference between `import` and `from`. With `import` the module itself is imported, but it retains its own namespace, which is why you need to use `module.name`.

Example 8.11. Introducing `vars`

Look at the following block of code at the bottom of `vars.py`:

```

import sys
import os

```

- 1 Just so you don't get intimidated, remember that you've seen all this before. The `vars` function returns a dictionary, and you're iterating through it.

Now running the script from the command line gives this output (note that your output may be slightly different, depending on your platform):

```

$ python vars.py
{'args': [], 'kwargs': {}, 'os': <module 'os' from 'os.py'>, 'sys': <module 'sys' from 'sys.py'>}
```

- 1 `os` was imported from `vars.py` using `import`. That means that it was imported directly into the module's namespace, and here it is.
- 2 Contrast this with `sys` which was imported using `from`. That means that the `sys` module itself is in the namespace, but the `sys` variable defined within `vars.py` is a reference to the `sys` module.
- 3 This module only defines one class, `vars`, and here it is. Note that the value here is the class itself, not a specific instance of the class.
- 4 Remember the `__name__` trick? When running a module (as opposed to importing it from another module), the built-in `__name__` attribute is a special value, `'__main__'`.

Using the `vars` and `globals` functions, you can get the value of arbitrary variables dynamically, providing the variable name as a string. This mirrors the `vars` function.

There is one other important difference between the `vars` and `globals` functions, which you should learn now before it bites you. It will bite you anyway.

Example 8.12. `vars` is read-only, `globals` is not

```

import sys
import os

class MyClass:
    pass

def main():
    vars['MyClass'] = MyClass
    print(vars)

```

- ❶ Since `dict` is called with `z` this will print `z`. This should not be a surprise.
- ❷ `dict` is a function that returns a dictionary, and here you are setting a value in that dictionary. You might think that this would change the value of `z`.
- ❸ This prints `z`, not `z`.
- ❹ After being burned by `dict` you might think that this *wouldn't* change the value of `z` but it does. Due to internal differences in how Python handles dictionaries.
- ❺ This prints `z`, not `z`.

8.6. Dictionary-based string formatting

Why did you learn about `dict` and `dict`? So you can learn about dictionary-based string formatting. As you recall, regular string formatting provides a way to format strings using tuples of values.

There is an alternative form of string formatting that uses dictionaries instead of tuples of values.

Example 8.13. Introducing dictionary-based string formatting

```

1 dict = {'a': 1, 'b': 2}
2 s = 'a: {} b: {}'.format(**dict)
3 print(s)

```

- ❶ Instead of a tuple of explicit values, this form of string formatting uses a dictionary, `dict`. And instead of a simple `%` marker in the string, the `format` method uses curly braces with keys.
- ❷ Dictionary-based string formatting works with any number of named keys. Each key must exist in the given dictionary, or the formatting will fail.
- ❸ You can even specify the same key twice; each occurrence will be replaced with the same value.

So why would you use dictionary-based string formatting? Well, it does seem like overkill to set up a dictionary of keys and values simply to format a string.

Example 8.14. Dictionary-based string formatting in `print`

```

1 dict = {'a': 1, 'b': 2}
2 print('a: {} b: {}'.format(**dict))

```

- ❶ Using the built-in `print` function is the most common use of dictionary-based string formatting. It means that you can use the names of local variables directly in the format string.

Example 8.15. More dictionary-based string formatting

```

1 dict = {'a': 1, 'b': 2}
2 s = 'a: {} b: {}'.format(**dict)
3 print(s)

```

- ❶ When this method is called, `dict` is a list of key/value tuples, just like the `dict` of a dictionary, which means you can use multi-variable assignment to unpack the list. Suppose `dict = [('a', 1), ('b', 2)]`. In the first round of the list comprehension, `a` will get `1` and `b` will get `2`. The string formatting `'a: {} b: {}'` will resolve to `'a: 1 b: 2'`. This string becomes the first element of the list comprehension's return value. In the second round, `a` will get `2` and `b` will get `1`. The string formatting will resolve to `'a: 2 b: 1'`. The list comprehension returns a list of these two resolved strings, and `join` will join both elements of this list together to form `'a: 1 b: 2 a: 2 b: 1'`.

Example 8.17. Handling specific tags



- 1 is called every time finds a tag in the HTML source. (In a minute, you'll see exactly how this happens.) The method takes a single parameter.
- 2 In the method, you initialize a data attribute that serves as a counter for tags. Every time you hit a tag, you increment the counter; eventually, you'll see how this is used.
- 3 That's it, that's the only special processing you do for tags. Now you pass the list of attributes along to so it can do the default processing.
- 4 is called every time finds a tag. Since end tags can not contain attributes, the method takes no parameters.
- 5 First, you want to do the default processing, just like any other end tag.
- 6 Second, you decrement your counter to signal that this block has been closed.

At this point, it's worth digging a little further into I've claimed repeatedly (and you've taken it on faith so far) that looks for and calls specific methods.

Example 8.18.



- 1 At this point, has already found a start tag and parsed the attribute list. The only thing left to do is figure out whether there is a specific method defined for this tag.
- 2 The "magic" of is nothing more than your old friend, getattr. What you may not have realized before is that getattr will find methods defined in descendant classes.
- 3 getattr raises an AttributeError if the method it's looking for doesn't exist in the object (or any of its descendants), but that's okay, because you wrapped the call in a try/except block.
- 4 Since you didn't find a method, you'll also look for a method before giving up. This alternate naming scheme is generally used for style sheets.
- 5 Another method which means that the call to getattr failed with an AttributeError. Since you found neither a method nor a method for this tag, you catch the exception and return None.
- 6 Remember, blocks can have an else clause, which is called if no exception is raised during the block. Logically, that means that you did find a method.
- 7 By the way, don't worry about these different return values; in theory they mean something, but they're never actually used. Don't worry about the return value of getattr either.
- 8 and methods are not called directly; the tag, method, and attributes are passed to this function, so that descendants can override it and do something different.

Now back to our regularly scheduled program: When you left, you were in the process of defining specific handler methods for and flags. The

Example 8.19. Overriding the method

- 1 is called with only one argument, the text to process.
- 2 In the ancestor the method simply appended the text to the output buffer, Here the logic is only slightly more complicated. If you're

You're close to completely understanding The only missing link is the nature of the text substitutions themselves. If you know any Perl, you

8.9. Putting it all together

It's time to put everything you've learned so far to good use. I hope you were paying attention.

Example 8.20. The function, part 1

- 1 The function has an optional argument, which is a string that specifies the dialect you'll be using. You'll see how this is used in a minute.
- 2 Hey, wait a minute, there's an statement in this function! That's perfectly legal in Python. You're used to seeing statements at the top of a function.
- 3 Now you get the source of the given URL.

Example 8.21. The function, part 2: curiouuser and curiouuser

- 1 is a string method you haven't seen before; it simply capitalizes the first letter of a string and forces everything else to lowercase. Consider
- 2 You have the name of a class as a string (C) and you have the global namespace as a dictionary (G). Combined, you can get a reference to the class.
- 3 Finally, you have a class object (C) and you want an instance of the class. Well, you already know how to do that: call the class like a function.

Why bother? After all, there are only 3 classes; why not just use a statement? (Well, there's no statement in Python, but why not just use a statement?)

Even better, imagine putting C in a separate module, and importing it with G. You've already seen that this includes it in G, so C would still work.

Now imagine that the name of the dialect is coming from somewhere outside the program, maybe from a database or from a user-inputted value.

Finally, imagine a framework with a plug-in architecture. You could put each class in a separate file, leaving only the function in G. Assuming

Example 8.22. The function, part 3

- 1
- 2
- 3

- ❶ After all that imagining, this is going to seem pretty boring, but the `render` function is what does the entire transformation. You had the entire page rendered in a single string.
- ❷ Because `Parser` maintains an internal buffer, you should always call the parser's `flush` method when you're done (even if you fed it all at once, like we did).
- ❸ Remember, `render` is the function you defined on `Parser` that joins all the pieces of output you've buffered and returns them in a single string.

And just like that, you've "translated" a web page, given nothing but a URL and the name of a dialect.

Further reading

You thought I was kidding about the server-side scripting idea. So did I, until I found this web-based dialectizer (<http://rinkworks.com>).

8.10. Summary

Python provides you with a powerful tool, `BeautifulSoup`, to manipulate HTML by turning its structure into an object model. You can use this tool in many ways:

- parsing the HTML looking for something specific
- aggregating the results, like the URL lister
- altering the structure along the way, like the attribute quoter
- transforming the HTML into something else by manipulating the text while leaving the tags alone, like the `HTML2Text` class

Along with these examples, you should be comfortable doing all of the following things:

- Using `from` and `import` to access namespaces
- Formatting strings using dictionary-based substitutions

[The technical term for a parser like `BeautifulSoup` is a *consumer*: it consumes HTML and breaks it down. Presumably, the name `BeautifulSoup` was chosen to fit into the `BeautifulSoup` namespace.

[The reason Python is better at lists than strings is that lists are mutable but strings are immutable. This means that appending to a list just adds an element, while appending to a string just creates a new string.

[Don't get out much.

[All right, it's not that common a question. It's not up there with "What editor should I use to write Python code?" (answer: Emacs) or "Is Python 3 better than Python 2?"

Example 9.2. 








This is, of course, complete gibberish. Well, not complete gibberish. It is syntactically and grammatically correct (although very verbose -- K

Let me repeat that this is much, much funnier if you are now or have ever been a philosophy major.

The interesting thing about this program is that there is nothing Kant-specific about it. All the content in the previous example was derived fr

Example 9.4. Simpler output from



You will take a closer look at the structure of the grammar file later in this chapter. For now, all you need to know is that the grammar file de

9.2. Packages

Actually parsing an XML document is very simple: one line of code. However, before you get to that line of code, you need to take a short d

Example 9.5. Loading an XML document (a sneak peek)





① Since the `find` method is defined in the `Element` class, it is available on any XML node, not just the `Element`.

Example 9.11. Child nodes can be text



- ① Looking at the XML in `example9_11.py` you might think that the `root` has only two child nodes, the two `Element`s. But you're missing something: the carriage return after the `<?xml` tag and before the first `<tag`.
- ② The first child is a `Text` object representing the carriage return after the `<?xml` tag and before the first `<tag`.
- ③ The second child is an `Element` object representing the first `Element`.
- ④ The fourth child is an `Element` object representing the second `Element`.
- ⑤ The last child is a `Text` object representing the carriage return after the `</tag` tag and before the `</tag` tag.

Example 9.12. Drilling down all the way to text





- ❶ As you saw in the previous example, the first element is `<?xml` since `childNodes[0]` is a `Text` node for the carriage return.
- ❷ The `<?xml` element has its own set of child nodes, one for the carriage return, a separate one for the spaces, one for the `<?xml` element, and so forth.
- ❸ You can even use the `find` method here, deeply nested within the document.
- ❹ The `<?xml` element has only one child node (you can't tell that from this example, but look at `childNodes` if you don't believe me), and it is a `Text` node for the carriage return.
- ❺ The `data` attribute of a `Text` node gives you the actual string that the text node represents. But what is that `
` in front of the string? The answer to that is...

9.4. Unicode

Unicode is a system to represent characters from all the world's different languages. When Python parses an XML document, all data is stored in Unicode. You'll get to all that in a minute, but first, some background.

Historical note. Before unicode, there were separate character encoding systems for each language, each using the same numbers (0-255) to represent characters. For example, the number 65 represented the character 'A' in ASCII, but in other systems it represented different characters.

Of course, there is still the matter of all these legacy encoding systems. 7-bit ASCII, for instance, which stores English characters as numbers 0-127. When dealing with unicode data, you may at some point need to convert the data back into one of these other legacy encoding systems. For instance, to store it in a database. And on that note, let's get back to Python.

Python has had unicode support throughout the language since version 2.0. The XML package uses unicode to store all parsed XML data, but...

Example 9.13. Introducing unicode

- ❶ To create a unicode string instead of a regular ASCII string, add the letter "u" before the string. Note that this particular string doesn't have any non-ASCII characters.
- ❷ When printing a string, Python will attempt to convert it to your default encoding, which is usually ASCII. (More on this in a minute.)

Example 9.14. Storing non-ASCII characters

...



- ❶ The real advantage of unicode, of course, is its ability to store non-ASCII characters, like the Spanish "ñ" (n with a tilde over it). The un
- ❷ Remember I said that the print function attempts to convert a unicode string to ASCII so it can print it? Well, that's not going to work here.
- ❸ Here's where the conversion-from-unicode-to-other-encoding-schemes comes in. s is a unicode string, but print can only print a regular string

Remember I said Python usually converted unicode to ASCII whenever it needed to make a regular string out of a unicode string? Well, this

Example 9.15. 📌



- ❶ is a special script; Python will try to import it on startup, so any code in it will be run automatically. As the comment mentions, it can
- ❷ function sets, well, the default encoding. This is the encoding scheme that Python will try to use whenever it needs to auto-coerce a un

Example 9.16. Effects of setting the default encoding



- ❶ This example assumes that you have made the changes listed in the previous example to your file, and restarted Python. If your default
- ❷ Now that the default encoding scheme includes all the characters you use in your string, Python has no problem auto-coercing the string

Example 9.17. Specifying encoding in files

If you are going to be storing non-ASCII strings within your Python code, you'll need to specify the encoding of each individual file by putting



Now, what about XML? Well, every XML document is in a specific encoding. Again, ISO-8859-1 is a popular encoding for data in Western

Example 9.18. 📌





It has two `find` methods either a `find` and a `findall`

Example 9.21. Introducing `find`



1 `find` takes one argument, the name of the element you wish to find. It returns a list of `Element` objects, corresponding to the XML elements that have the given name.

Example 9.22. Every element is searchable



- 1 Continuing from the previous example, the first object in your list is the `Element`.
- 2 You can use the same `find` method on this `Element` to find all the `Element` objects within the `Element`.
- 3 Just as before, the `find` method returns a list of all the elements it found. In this case, you have two, one for each bit.

Example 9.23. Searching is actually recursive





- 1 Note carefully the difference between this and the previous example. Previously, you were searching for elements within `books` but here you are searching for elements within `books[0]`.
- 2 The first two elements are within the first `book` element.
- 3 The last element is the one within the second `book` element.

9.6. Accessing element attributes

XML elements can have one or more attributes, and it is incredibly simple to access them once you have parsed an XML document.

For this section, you'll be using the `books.xml` grammar file that you saw in the previous section.

This section may be a little confusing, because of some overlapping terminology. Elements in an XML document have attributes, and Python objects have attributes.

Example 9.24. Accessing element attributes



- 1 Each `Element` object has an attribute called `attrib` which is a `dict` object. This sounds scary, but it's not, because a `dict` is an object that acts like a dictionary.
- 2 Treating the `attrib` as a dictionary, you can get a list of the names of the attributes of this element by using `attrib.keys()`. This element has only one attribute.
- 3 Attribute names, like all other text in an XML document, are stored in unicode.
- 4 Again treating the `attrib` as a dictionary, you can get a list of the values of the attributes by using `attrib.values()`. The values are themselves objects, of type `Element`.
- 5 Still treating the `attrib` as a dictionary, you can access an individual attribute by name, using normal dictionary syntax. (Readers who have been following the book will recognize this as the same syntax used to access dictionary elements.)

Example 9.25. Accessing individual attributes



- 1 The `Attribute` object completely represents a single XML attribute of a single XML element. The name of the attribute (the same name as you used in the XML) is stored in `name`.
- 2 The actual text value of this XML attribute is stored in `value`.

Like a dictionary, attributes of an XML element have no ordering. Attributes may *happen to be* listed in a certain order in the original XML document.

9.7. Segue

OK, that's it for the hard-core XML stuff. The next chapter will continue to use these same example programs, but focus on other aspects than XML.

Before moving on to the next chapter, you should be comfortable doing all of these things:

- Parsing XML documents using `lxml`
- Searching through the parsed document, and accessing arbitrary element attributes and element children
- Organizing complex libraries into packages
- Converting unicode strings to different character encodings

This, sadly, is *still* an oversimplification. Unicode now has been extended to handle ancient Chinese, Korean, and Japanese texts, which had

Chapter 10. Scripts and Streams

10.1. Abstracting input sources

One of Python's greatest strengths is its dynamic binding, and one powerful use of dynamic binding is the *file-like object*.

Many functions which require an input source could simply take a filename, go open the file for reading, read it, and close it when they're done.

In the simplest case, a *file-like object* is any object with a `read()` method with an optional `size` parameter, which returns a string. When called with no parameter, it returns a string of up to 1024 bytes.

This is how reading from real files works; the difference is that you're not limiting yourself to real files. The input source could be anything: a socket, a database cursor, or even a string.

In case you were wondering how this relates to XML processing, `xml.etree.ElementTree.parse()` is one such function which can take a file-like object.

Example 10.1. Parsing XML from a file

```
1 from xml.etree.ElementTree import parse
2
3 def parse_xml(filename):
4     f = open(filename, 'r')
5     tree = parse(f)
6     f.close()
7     return tree
8
9 if __name__ == '__main__':
10    tree = parse_xml('example.xml')
11    print(tree)
```

- 1 First, you open the file on disk. This gives you a file object.
- 2 You pass the file object to `parse()` which calls the `read()` method of `f` and reads the XML document from the file on disk.
- 3 Be sure to call the `close()` method of the file object after you're done with it. `parse()` will not do this for you.
- 4 Calling the `print()` method on the returned XML document prints out the entire thing.

Well, that all seems like a colossal waste of time. After all, you've already seen that `parse()` can simply take the filename and do all the opening and closing for you.

Example 10.2. Parsing XML from a URL

```
1 from xml.etree.ElementTree import parse
2
3 def parse_xml(url):
4     f = open(url, 'r')
5     tree = parse(f)
6     f.close()
7     return tree
8
9 if __name__ == '__main__':
10    tree = parse_xml('http://example.com/example.xml')
11    print(tree)
```



- ❶ As you saw in a previous chapter, `urllib.urlopen` takes a web page URL and returns a file-like object. Most importantly, this object has a `read` method which...
- ❷ Now you pass the file-like object to `xml.etree.ElementTree.parse` which obediently calls the `read` method of the object and parses the XML data that the `read` method returns.
- ❸ As soon as you're done with it, be sure to close the file-like object that `urllib.urlopen` gives you.
- ❹ By the way, this URL is real, and it really is XML. It's an XML representation of the current headlines on Slashdot (<http://slashdot.org>).

Example 10.3. Parsing XML from a string (the easy but inflexible way)



❶ `ElementTree.parse` has a method, `parse_from_string`, which takes an entire XML document as a string and parses it. You can use this instead of `parse` if you know you already have a string.

OK, so you can use the `parse` function for parsing both local files and remote URLs, but for parsing strings, you use... a different function. That means you have to use `urllib.urlopen` to get a file-like object, then pass it to `parse`. If there were a way to turn a string into a file-like object, then you could simply pass this object to `parse`. And in fact, there is a module specifically for this purpose.

Example 10.4. Introducing `io.StringIO`





- ❶ The `io` module contains a single class, also called `IO` which allows you to turn a string into a file-like object. The `IO` class takes the string as a parameter.
- ❷ Now you have a file-like object, and you can do all sorts of file-like things with it. Like `read()` which returns the original string.
- ❸ Calling `read()` again returns an empty string. This is how real file objects work too; once you read the entire file, you can't read any more without seeking.
- ❹ You can explicitly seek to the beginning of the string, just like seeking through a file, by using the `seek()` method of the `IO` object.
- ❺ You can also read the string in chunks, by passing a `size` parameter to the `read()` method.
- ❻ At any time, `read()` will return the rest of the string that you haven't read yet. All of this is exactly how file objects work; hence the term *file-like object*.

Example 10.5. Parsing XML from a string (the file-like object way)



- ❶ Now you can pass the file-like object (really a `IO`) to `parse()` which will call the object's `read()` method and happily parse away, never knowing that it's just a string.

So now you know how to use a single function, `parse()` to parse an XML document stored on a web page, in a local file, or in a hard-coded string. Fun!

Example 10.6. `urlretrieve()`



- ❶ The `urlretrieve()` function takes a single parameter, `url`, and returns a file-like object. `url` is a string of some sort; it can either be a URL (like `http://www.python.org`) or a path to a file on disk.
- ❷ First, you see if `url` is a URL. You do this through brute force: you try to open it as a URL and silently ignore errors caused by trying to connect to a non-existent host.
- ❸ On the other hand, if `url` is a path to a file on disk, you assume it's a path to a file on disk and try to open it. As you can see, `urlretrieve()` is a bit of a hack.

4 By this point, you need to assume that `data` is a string that has hard-coded data in it (since nothing else worked), so you use `print` to create a file

Now you can use this `print_xml` function in conjunction with `open` to make a function that takes a `path` that refers to an XML document somehow (either as a URL

Example 10.7. Using `print_xml`

```
#!/usr/bin/env python
import sys
import xml.etree.ElementTree as ET

def print_xml(path):
    tree = ET.parse(path)
    root = tree.getroot()
    print(root)

if __name__ == '__main__':
    print_xml(sys.argv[1])
```

10.2. Standard input, output, and error

UNIX users are already familiar with the concept of standard input, standard output, and standard error. This section is for the rest of you.

Standard output and standard error (commonly abbreviated `stdout` and `stderr`) are pipes that are built into every UNIX system. When you do something, it goes

Example 10.8. Introducing `stdout` and `stderr`

```
#!/usr/bin/env python
import sys
import os

def main():
    # Print to stdout
    print('Hello, world!')

    # Print to stderr
    print('Error message', file=sys.stderr)

if __name__ == '__main__':
    main()
```

1 As you saw in Example 6.9, “Simple Counters”, you can use Python’s built-in `range` function to build simple counter loops that repeat some

2 `print` is a file-like object; calling its `write` function will print out whatever string you give it. In fact, this is what the `print` function really does; it adds

3 In the simplest case, `stdout` and `stderr` send their output to the same place: the Python IDE (if you’re in one), or the terminal (if you’re running Python

`stdout` and `stderr` are both file-like objects, like the ones you discussed in Section 10.1, “Abstracting input sources”, but they are both write-only. They have

Example 10.9. Redirecting output

```
#!/usr/bin/env python
import sys
import os

def main():
    # Redirect stdout to a file
    sys.stdout = open('output.txt', 'w')
    print('Hello, world!')

if __name__ == '__main__':
    main()
```

(On Windows, you can use `print` instead of `print` to display the contents of a file.)

If you have not already done so, you can download this and other examples (<http://diveintopython.org/download/diveintopython-examples-5>).

```
#
```



- 1 This will print to the IDE "Interactive Window" (or the terminal, if running the script from the command line).
- 2 Always save before redirecting it, so you can set it back to normal later.
- 3 Open a file for writing. If the file doesn't exist, it will be created. If the file does exist, it will be overwritten.
- 4 Redirect all further output to the new file you just opened.
- 5 This will be "printed" to the log file only; it will not be visible in the IDE window or on the screen.
- 6 Set back to the way it was before you mucked with it.
- 7 Close the log file.

Redirecting works exactly the same way, using instead of

Example 10.10. Redirecting error information



If you have not already done so, you can download this and other examples (<http://diveintopython.org/download/diveintopython-examples-5>).



- 1 Open the log file where you want to store debugging information.
- 2 Redirect standard error by assigning the file object of the newly-opened log file to
- 3 Raise an exception. Note from the screen output that this does *not* print anything on screen. All the normal traceback information has
- 4 Also note that you're not explicitly closing your log file, nor are you setting back to its original value. This is fine, since once the prog

Since it is so common to write error messages to standard error, there is a shorthand syntax that can be used instead of going through the hass

Example 10.11. Printing to




```
#
```

- 1 `__call__` is only ever called once, since there is only one `Node` in an XML document, and only one `Node` object in the parsed XML representation. I
- 2 `__call__` is called on nodes that represent bits of text. The function itself does some special processing to handle automatic capitalization of the
- 3 `__call__` is just a `__call__` since you don't care about embedded comments in the grammar files. Note, however, that you still need to define the function
- 4 The `__call__` method is actually itself a dispatcher, based on the name of the element's tag. The basic idea is the same: take what distinguishes

In this example, the dispatch functions `__call__` and `__call__` simply find other methods in the same class. If your processing is very complex (or you have ma

10.6. Handling command-line arguments

Python fully supports creating programs that can be run on the command line, complete with command-line arguments and either short- or lo

It's difficult to talk about command-line processing without understanding how command-line arguments are exposed to your Python program

Example 10.20. Introducing `sys.argv`

If you have not already done so, you can download this and other examples (<http://diveintopython.org/download/diveintopython-examples-5>).

```
#!/usr/bin/env python
```

```
import sys
```

```
for arg in sys.argv[1:]:
```

```
    print arg
```

- 1 Each command-line argument passed to the program will be in `sys.argv` which is just a list. Here you are printing each argument on a separate

Example 10.21. The contents of `sys.argv`

```
#!/usr/bin/env python
```

```
import sys
```

```
print sys.argv
```

```
print sys.argv[0]
```

```
print sys.argv[1]
```

```
print sys.argv[2]
```

```
print sys.argv[3]
```

```
print sys.argv[4]
```

```
print sys.argv[5]
```

```
print sys.argv[6]
```

```
print sys.argv[7]
```

```
print sys.argv[8]
```

- 1 The first thing to know about `sys.argv` is that it contains the name of the script you're calling. You will actually use this knowledge to your adv
- 2 Command-line arguments are separated by spaces, and each shows up as a separate element in the `sys.argv` list.
- 3 Command-line flags, like `-h` also show up as their own element in the `sys.argv` list.
- 4 To make things even more interesting, some command-line flags themselves take arguments. For instance, here you have a flag (`-m`) whi

So as you can see, you certainly have all the information passed on the command line, but then again, it doesn't look like it's going to be all t

Example 10.22. Introducing `argparse`

```
#!/usr/bin/env python
import argparse

parser = argparse.ArgumentParser(
    description='A simple parser for command-line arguments'
)
parser.add_argument(
    '-g', '--grammar',
    type=str,
    help='use specified grammar file or URL'
)
parser.add_argument(
    '-d', '--debug',
    action='store_true',
    help='show debugging information while parsing'
)

args = parser.parse_args()

# ...

if __name__ == '__main__':
    # ...
    parser.print_usage()
    # ...
```

- 1 First off, look at the bottom of the example and notice that you're calling the `main` function with `parser.parse_args()`. Remember, `main` is the name of the script that you're running.
- 2 This is where all the interesting processing happens. The `main` function of the `main` module takes three parameters: the argument list (which you get from `parser.parse_args()`), the `parser` object, and the `args` object.
- 3 If anything goes wrong trying to parse these command-line flags, `parser.parse_args()` will raise an exception, which you catch. You told `parser` all the flags you expect to see.
- 4 As is standard practice in the UNIX world, when the script is passed flags it doesn't understand, you print out a summary of proper usage.

So what are all those parameters you pass to the `main` function? Well, the first one is simply the raw list of command-line flags and arguments (not

```
h
    print usage summary
g
    use specified grammar file or URL
d
    show debugging information while parsing
```

The first and third flags are simply standalone flags; you specify them or you don't, and they do things (print help) or change state (turn on debug).

To further complicate things, the script accepts either short flags (like `-g`) or long flags (like `--grammar`) and you want them to do the same thing. This is

```
h
    print usage summary
g
    use specified grammar file or URL
```

Three things of note here:

- All long flags are preceded by two dashes on the command line, but you don't include those dashes when calling `parser.parse_args()`. They are understood by `parser`.
- The `-g` flag must always be followed by an additional argument, just like the `--grammar` flag. This is notated by an equals sign, `=str`.
- The list of long flags is shorter than the list of short flags, because the `-d` flag does not have a corresponding long version. This is fine;

Confused yet? Let's look at the actual code and see if it makes sense in context.

Example 10.23. Handling command-line arguments in `argparse`

```
#!/usr/bin/env python
import argparse

parser = argparse.ArgumentParser(
    description='A simple parser for command-line arguments'
)
parser.add_argument(
    '-g', '--grammar',
    type=str,
    help='use specified grammar file or URL'
)
parser.add_argument(
    '-d', '--debug',
    action='store_true',
    help='show debugging information while parsing'
)

args = parser.parse_args()

# ...

if __name__ == '__main__':
    # ...
    parser.print_usage()
    # ...
```


Oh, and along the way, you take advantage of your knowledge of the structure of the XML document to set up a little cache of references, wh

```


```

If you specified some source material on the command line, you use that; otherwise you rip through the grammar looking for the "top-level" r

```


```

Now you rip through the source material. The source material is also XML, and you parse it one node at a time. To keep the code separated a

```


```

You bounce through the grammar, parsing all the children of each element,

```


```

replacing elements with a random child,

```


```

and replacing elements with a random child of the corresponding element, which you previously cached.

```


```

Eventually, you parse your way down to plain text,

```


```

which you print out.

```


```

10.8. Summary

Python comes with powerful libraries for parsing and manipulating XML documents. The `lxml` takes an XML file and parses it into Python objects.

Before moving on to the next chapter, you should be comfortable doing all of these things:

- Chaining programs with standard input and output
- Defining dynamic dispatchers with `dispatch`
- Using command-line flags and validating them with `argparse`

Chapter 11. HTTP Web Services

11.1. Diving in

You've learned about HTML processing and XML processing, and along the way you saw how to download a web page and how to parse XML.

Simply stated, HTTP web services are programmatic ways of sending and receiving data from remote servers using the operations of HTTP.

The main advantage of this approach is simplicity, and its simplicity has proven popular with a lot of different sites. Data -- usually XML data -- is sent over HTTP.

Examples of pure XML-over-HTTP web services:

Amazon API (<http://www.amazon.com/webservices>) allows you to retrieve product information from the Amazon.com online store.

National Weather Service (<http://www.nws.noaa.gov/alerts/>) (United States) and Hong Kong Observatory (<http://demo.xml.weather.gov.hk/>)

Atom API (<http://atomenabled.org/>) for managing web-based content.

Syndicated feeds (<http://syndic8.com/>) from weblogs and news sites bring you up-to-the-minute news from a variety of sites.

In later chapters, you'll explore APIs which use HTTP as a transport for sending and receiving data, but don't map application semantics to the data.

Here is a more advanced version of the `urllib` module that you saw in the previous chapter:

Example 11.1. `urllib`

If you have not already done so, you can download this and other examples (<http://diveintopython.org/download/diveintopython-examples-5.0/>).

```
#!/usr/bin/python
import sys
import urllib
import urllib2
import urllib3
import urllib4
import urllib5
import urllib6
import urllib7
import urllib8
import urllib9
import urllib10
import urllib11
import urllib12
import urllib13
import urllib14
import urllib15
import urllib16
import urllib17
import urllib18
import urllib19
import urllib20
import urllib21
import urllib22
import urllib23
import urllib24
import urllib25
import urllib26
import urllib27
import urllib28
import urllib29
import urllib30
import urllib31
import urllib32
import urllib33
import urllib34
import urllib35
import urllib36
import urllib37
import urllib38
import urllib39
import urllib40
import urllib41
import urllib42
import urllib43
import urllib44
import urllib45
import urllib46
import urllib47
import urllib48
import urllib49
import urllib50
import urllib51
import urllib52
import urllib53
import urllib54
import urllib55
import urllib56
import urllib57
import urllib58
import urllib59
import urllib60
import urllib61
import urllib62
import urllib63
import urllib64
import urllib65
import urllib66
import urllib67
import urllib68
import urllib69
import urllib70
import urllib71
import urllib72
import urllib73
import urllib74
import urllib75
import urllib76
import urllib77
import urllib78
import urllib79
import urllib80
import urllib81
import urllib82
import urllib83
import urllib84
import urllib85
import urllib86
import urllib87
import urllib88
import urllib89
import urllib90
import urllib91
import urllib92
import urllib93
import urllib94
import urllib95
import urllib96
import urllib97
import urllib98
import urllib99
import urllib100
```




Further reading

Paul Prescod believes that pure HTTP web services are the future of the Internet (<http://webservices.xml.com/pub/a/ws/2002/02/06/>)

11.2. How not to fetch data over HTTP

Let's say you want to download a resource over HTTP, such as a syndicated Atom feed. But you don't just want to download it once; you want

Example 11.2. Downloading a feed the quick-and-dirty way



❶ Downloading anything over HTTP is incredibly easy in Python; in fact, it's a one-liner. The `urllib` module has a handy `urlopen` function that takes t

So what's wrong with this? Well, for a quick one-off during testing or development, there's nothing wrong with it. I do it all the time. I want

Let's talk about some of the basic features of HTTP.

11.3. Features of HTTP

There are five important features of HTTP which you should support.

11.3.1. User-Agent

The `User-Agent` is simply a way for a client to tell a server who it is when it requests a web page, a syndicated feed, or any sort of web service over HTTP.

By default, Python sends a generic `User-Agent`. In the next section, you'll see how to change this to something more specific.

11.3.2. Redirects

Sometimes resources move around. Web sites get reorganized, pages move to new addresses. Even web services can reorganize. A syndicated feed is a good example of this. Every time you request any kind of resource from an HTTP server, the server includes a status code in its response. Status code 301 means "permanent redirect". HTTP has two different ways of signifying that a resource has moved. Status code 302 is a *temporary redirect*; it means "oops, that got moved over there". A browser will automatically "follow" redirects when it receives the appropriate status code from the HTTP server, but unfortunately, it doesn't tell you where it went.

11.3.3. Last-Modified

Some data changes all the time. The home page of CNN.com is constantly updating every few minutes. On the other hand, the home page of a newspaper is updated only once a day. If you ask for the same data a second time (or third, or fourth), you can tell the server the last-modified date that you got last time: you send a `Last-Modified` header. All modern web browsers support last-modified date checking. If you've ever visited a page, re-visited the same page a day later and found that it had changed, that's because the browser checked the last-modified date. Python's URL library has no built-in support for last-modified date checking, but since you can add arbitrary headers to each request and read the response headers, you can do it.

11.3.4. ETags

ETags are an alternate way to accomplish the same thing as the last-modified date checking: don't re-download data that hasn't changed. The server sends an `Etag` header with the data, and you send an `If-None-Match` header with the ETag you got last time. Python's URL library has no built-in support for ETags, but you'll see how to add it later in this chapter.

11.3.5. Compression

The last important HTTP feature is gzip compression. When you talk about HTTP web services, you're almost always talking about moving data around. Python's URL library has no built-in support for gzip compression per se, but you can add arbitrary headers to the request. And Python comes with a `gzip` module that can decompress data. Note that our little one-line script to download a syndicated feed did not support any of these HTTP features. Let's see how you can improve it.

11.4. Debugging HTTP web services

First, let's turn on the debugging features of Python's HTTP library and see what's being sent over the wire. This will be useful throughout the chapter.

Example 11.3. Debugging HTTP





- ❶ relies on another standard Python library, urllib. Normally you don't need to import it directly (urllib does that automatically), but you will here so you can use the urllib.urlopen() function.
- ❷ Now that the debugging flag is set, information on the the HTTP request and response is printed out in real time. The first thing it tells you is the URL you're requesting.
- ❸ When you request the Atom feed, urllib sends three lines to the server. The first line specifies the HTTP verb you're using, and the path of the resource you're requesting.
- ❹ The second line is the Host header, which specifies the domain name of the service you're accessing. This is important, because a single Host header can be used to request multiple resources from the same server.
- ❺ The third line is the User-Agent header. What you see here is the generic User-Agent that the urllib library adds by default. In the next section, you'll see how to customize it.
- ❻ The server replies with a status code and a bunch of headers (and possibly some data, which got stored in the response variable). The status code is 200, which means the request was successful.
- ❼ The server tells you when this Atom feed was last modified (in this case, about 13 minutes ago). You can send this date back to the server to request only the new data.
- ❽ The server also tells you that this Atom feed has an ETag hash of '123456789'. The hash doesn't mean anything by itself; there's nothing you can do with it.

11.5. Setting the User-Agent

The first step to improving your HTTP web services client is to identify yourself properly with a User-Agent. To do that, you need to move beyond the basic urllib.urlopen() function.

Example 11.4. Introducing urllib2



- ❶ If you still have your Python IDE open from the previous section's example, you can skip this, but this turns on HTTP debugging so you can see what's going on.
- ❷ Fetching an HTTP resource with urllib2 is a three-step process, for good reasons that will become clear shortly. The first step is to create a urllib2.Request object.
- ❸ The second step is to build a URL opener. This can take any number of handlers, which control how responses are handled. But you can use the default opener.

4 The final step is to tell the opener to open the URL, using the `obj` you created. As you can see from all the debugging information t

Example 11.5. Adding headers with the `obj`



- 1 You're continuing from the previous example; you've already created a `obj` with the URL you want to access.
- 2 Using the `obj` method on the `obj`, you can add arbitrary HTTP headers to the request. The first argument is the header, the second is th
- 3 The `obj` you created before can be reused too, and it will retrieve the same feed again, but with your custom `obj` header.
- 4 And here's you sending your custom `obj` in place of the generic one that Python sends by default. If you look closely, you'll notice that y

11.6. Handling `obj` and `obj`

Now that you know how to add custom HTTP headers to your web service requests, let's look at adding support for `obj` and `obj` headers.

These examples show the output with debugging turned off. If you still have it turned on from the previous section, you can turn it off by sett

Example 11.6. Testing `obj`





- ❶ Remember all those HTTP headers you saw printed out when you turned on debugging? This is how you can get access to them programmatically.
- ❷ On the second request, you add the `headers` with the last-modified date from the first request. If the data hasn't changed, the server should return a 304 status code.
- ❸ Sure enough, the data hasn't changed. You can see from the traceback that `urllib.error.HTTPError` is raised in response to the 304 status code. `urllib.error.HTTPError` also raises an `urllib.error.URLError` for conditions that you would think of as errors, such as 404 (page not found). In fact, it will raise `urllib.error.HTTPError` for *any* status code other than 200.

Example 11.7. Defining URL handlers

This custom URL handler is part of `urllib2`:



- ❶ `urllib2` is designed around URL handlers. Each handler is just a class that can define any number of methods. When something happens -- like a 304 status code -- `urllib2` searches through the defined handlers and calls the `http_error_304` method when it encounters a 304 status code from the server. By defining a custom handler, you can intercept the 304 status code and return the data from the first request.
- ❷ `urllib2` searches through the defined handlers and calls the `http_error_304` method when it encounters a 304 status code from the server. By defining a custom handler, you can intercept the 304 status code and return the data from the first request.
- ❸ This is the key part: before returning, you save the status code returned by the HTTP server. This will allow you easy access to it from `urllib2`.

Example 11.8. Using custom URL handlers





- ❶ You're continuing the previous example, so the `obj` is already set up, and you've already added the `header`.
- ❷ This is the key: now that you've defined your custom URL handler, you need to tell `obj` to use it. Remember how I said that `obj` broke up the
- ❸ Now you can quietly open the resource, and what you get back is an object that, along with the usual headers (use `obj` to access them), also
- ❹ Note that when the server sends back a `status` code, it doesn't re-send the data. That's the whole point: to save bandwidth by not re-do

Handling `obj` works much the same way, but instead of checking for `obj` and sending `obj` you check for `obj` and send `obj`. Let's start with a fresh IDE session

Example 11.9. Supporting `obj`



- ❶ Using the `obj` pseudo-dictionary, you can get the `obj` returned from the server. (What happens if the server didn't send back an `obj`? Then this line
- ❷ OK, you got the data.
- ❸ Now set up the second call by setting the `header` to the `obj` you got from the first call.
- ❹ The second call succeeds quietly (without throwing an exception), and once again you see that the server has sent back a `status` code.
- ❺ Regardless of whether the `obj` is triggered by `obj` late checking or `obj` hash matching, you'll never get the data along with the `obj`. That's the whole

In these examples, the HTTP server has supported both `Location` and `URI` headers, but not all servers do. As a web services client, you should be prepared

11.7. Handling redirects

You can support permanent and temporary redirects using a different kind of custom URL handler.

First, let's see why a redirect handler is necessary in the first place.

Example 11.10. Accessing web services without a redirect handler





- ❶ You'll be better able to see what's happening if you turn on debugging.
- ❷ This is a URL which I have set up to permanently redirect to my Atom feed at `http://www.snoke.com/atom.xml`
- ❸ Sure enough, when you try to download the data at that address, the server sends back a `301` status code, telling you that the resource has moved.
- ❹ The server also sends back a `Location` header that gives the new address of this data.
- ❺ `urllib2` notices the redirect status code and automatically tries to retrieve the data at the new location specified in the `Location` header.
- ❻ The object you get back from the `urllib2.urlopen` contains the new permanent address and all the headers returned from the second request (retrieved from `http://www.snoke.com/atom.xml`).

This is suboptimal, but easy to fix. `urllib2` doesn't behave exactly as you want it to when it encounters a `301`. So let's override its behavior. How? Well, we can define a custom redirect handler.

Example 11.11. Defining the redirect handler

This class is defined in `urllib2.py`:



- ❶ Redirect behavior is defined in `urllib2.py` in a class called `RedirectHandler`. You don't want to completely override the behavior, you just want to extend it a little.
- ❷ When it encounters a `301` status code from the server, `RedirectHandler` will search through its handlers and call the `handle_301` method. The first thing ours does is just call the ancestor's `handle_301` method.
- ❸ Here's the key: before you return, you store the status code (`code`) so that the calling program can access it later.
- ❹ Temporary redirects (status code `302`) work the same way: override the `handle_302` method, call the ancestor, and save the status code before returning.

So what has this bought us? You can now build a URL opener with the custom redirect handler, and it will still automatically follow redirects.

Example 11.12. Using the redirect handler to detect permanent redirects





- ❶ First, build a URL opener with the redirect handler you just defined.
- ❷ You sent off a request, and you got a 3 status code in response. At this point, the `do_request` method gets called. You call the ancestor method, which
- ❸ This is the payoff: now, not only do you have access to the new URL, but you have access to the redirect status code, so you can tell the

The same redirect handler can also tell you that you *shouldn't* update your address book.

Example 11.13. Using the redirect handler to detect temporary redirects





- ❶ This is a sample URL I've set up that is configured to tell clients to *temporarily* redirect to `http://www.example.com/`.
- ❷ The server sends back a `302` status code, indicating a temporary redirect. The temporary new location of the data is given in the `Location` header.
- ❸ `requests.get()` calls your `get` method, which calls the ancestor method of the same name in `requests.Session` which follows the redirect to the new location. Then your `get` method returns the response object.
- ❹ And here you are, having successfully followed the redirect to `http://www.example.com/`. `response.history` tells you that this was a temporary redirect, which means that you should follow it.

11.8. Handling compressed data

The last important HTTP feature you want to support is compression. Many web services have the ability to send data compressed, which can save a lot of bandwidth.

Servers won't give you compressed data unless you tell them you can handle it.

Example 11.14. Telling the server you would like compressed data





- 1 This is the key: once you've created your `Request` object, add an `Accept` header to tell the server you can accept gzip-encoded data. `headers` is the name of the dictionary.
- 2 There's your header going across the wire.
- 3 And here's what the server sends back: the `Content-Length` header means that the data you're about to receive has been gzip-compressed.
- 4 The `Content-Length` header is the length of the compressed data, not the uncompressed data. As you'll see in a minute, the actual length of the uncompressed data is much longer.

Example 11.15. Decompressing the data



- 1 Continuing from the previous example, `resp` is the file-like object returned from the URL opener. Using its `read()` method would ordinarily get you the compressed data.
- 2 OK, this step is a little bit of messy workaround. Python has a `gzip` module, which reads (and actually writes) gzip-compressed files on disk.
- 3 Now you can create an instance of `GzipFile` and tell it that its "file" is the file-like object `resp`.
- 4 This is the line that does all the actual work: "reading" from `gz` will decompress the data. Strange? Yes, but it makes sense in a twisted kind of way.
- 5 Look ma, real data. (15955 bytes of it, in fact.)

"But wait!" I hear you cry. "This could be even easier!" I know what you're thinking. You're thinking that `r.read()` returns a file-like object, so why not

Example 11.16. Decompressing the data directly from the server



- 1 Continuing from the previous example, you already have a `Request` object set up with an `Header`.
- 2 Simply opening the request will get you the headers (though not download any data yet). As you can see from the returned `Header`, this is a gzipped response.
- 3 Since `r.read()` returns a file-like object, and you know from the headers that when you read it, you're going to get gzip-compressed data, why not

11.9. Putting it all together

You've seen all the pieces for building an intelligent HTTP web services client. Now let's see how they all fit together.

Example 11.17. The `urlretrieve` function

This function is defined in `urllib2.py`



- 1 `urllib2` is a handy utility module for, you guessed it, parsing URLs. Its primary function, also called `urllib2.urlopen`, takes a URL and splits it into a tuple of (scheme, netloc, path, query, fragment).
- 2 You identify yourself to the HTTP server with the `user-agent` passed in by the calling function. If no `user-agent` was specified, you use a default one defined in `urllib2.py`.
- 3 If an `auth` was given, send it in the `header`.
- 4 If a last-modified date was given, send it in the `header`.
- 5 Tell the server you would like compressed data if possible.
- 6 Build a URL opener that uses *both* of the custom URL handlers: `urllib2.HTTPSHandler` for handling `https` and `urllib2.HTTPRedirectHandler` for handling `301` and other error conditions.

- 7 That's it! Open the URL and return a file-like object to the caller.

Example 11.18. The `urlretrieve` function

This function is defined in `urllib.py`:

```
def urlretrieve(url, filename, reporthook=None, reporthookargs=(), proxies=None):
```

- 1 First, you call the `urlretrieve` function with a URL, `filename`, `reporthook`, and `reporthookargs`.
- 2 Read the actual data returned from the server. This may be compressed; if so, you'll decompress it later.
- 3 Save the `hash` returned from the server, so the calling application can pass it back to you next time, and you can pass it on to `urlretrieve` which can pass it on to `urlretrieve`.
- 4 Save the `date` too.
- 5 If the server says that it sent compressed data, decompress it.
- 6 If you got a URL back from the server, save it, and assume that the status code is `200` until you find out otherwise.
- 7 If one of the custom URL handlers captured a status code, then save that too.

Example 11.19. Using `urlretrieve`

```
urlretrieve(url, filename, reporthook=None, reporthookargs=(), proxies=None)
```



- 1 The very first time you fetch a resource, you don't have an `if-modified-since` or `if-none-match`, so you'll leave those out. (They're optional parameters.)
- 2 What you get back is a dictionary of several useful headers, the HTTP status code, and the actual data returned from the server. `handle_redirect`
- 3 If you ever get a `301` status code, that's a permanent redirect, and you need to update your URL to the new address.
- 4 The second time you fetch the same resource, you have all sorts of information to pass back: a (possibly updated) URL, the `if-modified-since` from the first fetch, and the `if-none-match` from the first fetch.
- 5 What you get back is again a dictionary, but the data hasn't changed, so all you got was a `304` status code and no data.

11.10. Summary

The `requests` and its functions should now make perfect sense.

There are 5 important features of HTTP web services that every client should support:

- Identifying your application by setting a proper `Accept` header.
- Handling permanent redirects properly.
- Supporting `ETag` checking to avoid re-downloading data that hasn't changed.
- Supporting `Cache-Control` to avoid re-downloading data that hasn't changed.
- Supporting `gzip` compression to reduce bandwidth even when data *has* changed.

Chapter 12. SOAP Web Services

Chapter 11 focused on document-oriented web services over HTTP. The "input parameter" was the URL, and the "return value" was an actual document.

This chapter will focus on SOAP web services, which take a more structured approach. Rather than dealing with HTTP requests and XML documents, SOAP uses a more structured approach.

SOAP is a complex specification, and it is somewhat misleading to say that SOAP is all about calling remote functions. Some people would prefer to think of it as a way to call remote functions.

12.1. Diving In

You use Google, right? It's a popular search engine. Have you ever wished you could programmatically access Google search results? Now you can.

Example 12.1. `gs.py`

```
#!/usr/bin/env python
import sys
import urllib
import urllib2
import re

url = "http://www.google.com/search?q=%s" % sys.argv[1]
response = urllib2.urlopen(url)
html = response.read()

# Find the first search result
m = re.search('<div class="search-result">', html)
if m:
    # Find the title
    m = re.search('<h3>(.*)</h3>', m.group(0))
    title = m.group(1)

    # Find the URL
    m = re.search('<a href="(.*)>', m.group(0))
    url = m.group(1)

    # Find the snippet
    m = re.search('<p>(.*)</p>', m.group(0))
    snippet = m.group(1)

    # Print the result
    print title
    print url
    print snippet
```

You can import this as a module and use it from a larger program, or you can run the script from the command line. On the command line, you can use the following syntax:

Here is the sample output for a search for the word "python".

Example 12.2. Sample Usage of `gs.py`

```
python gs.py python
Python (programming language)
http://www.python.org/
Python is a high-level, general-purpose programming language. It is designed to be easy to learn and use, and is suitable for a wide range of applications. It has a simple, clear syntax and is interpreted, so it can be run without the need for a compiler. Python is also open source and has a large, active community of users and developers.
```



Further Reading on SOAP

<http://www.xmethods.net/> is a repository of public access SOAP web services.

The SOAP specification (<http://www.w3.org/TR/soap/>) is surprisingly readable, if you like that sort of thing.

12.2. Installing the SOAP Libraries

Unlike the other code in this book, this chapter relies on libraries that do not come pre-installed with Python.

Before you can dive into SOAP web services, you'll need to install three libraries: PyXML, fpconst, and SOAPpy.

12.2.1. Installing PyXML

The first library you need is PyXML, an advanced set of XML libraries that provide more functionality than the built-in XML libraries we st

Procedure 12.1.

Here is the procedure for installing PyXML:

1. Go to <http://pyxml.sourceforge.net/>, click Downloads, and download the latest version for your operating system.
2. If you are using Windows, there are several choices. Make sure to download the version of PyXML that matches the version of Pyth
3. Double-click the installer. If you download PyXML 0.8.3 for Windows and Python 2.3, the installer program will be 
4. Step through the installer program.
5. After the installation is complete, close the installer. There will not be any visible indication of success (no programs installed on the

To verify that you installed PyXML correctly, run your Python IDE and check the version of the XML libraries you have installed, as shown

Example 12.3. Verifying PyXML Installation



This version number should match the version number of the PyXML installer program you downloaded and ran.

12.2.2. Installing fpconst

The second library you need is fpconst, a set of constants and functions for working with IEEE754 double-precision special values. This provides

Procedure 12.2.

Here is the procedure for installing fpconst:

1. Download the latest version of fpconst from <http://www.analytics.washington.edu/statcomp/projects/rzope/fpconst/>.
2. There are two downloads available, one in `tar` format, the other in `zip` format. If you are using Windows, download the `zip` file; otherwise, download the `tar` file.
3. Decompress the downloaded file. On Windows XP, you can right-click on the file and choose Extract All; on earlier versions of Windows, you can use WinZip or 7-Zip.
4. Open a command prompt and navigate to the directory where you decompressed the fpconst files.
5. Type `python setup.py install` to run the installation program.

To verify that you installed fpconst correctly, run your Python IDE and check the version number.

Example 12.4. Verifying fpconst Installation



This version number should match the version number of the fpconst archive you downloaded and installed.

12.2.3. Installing SOAPpy

The third and final requirement is the SOAP library itself: SOAPpy.

Procedure 12.3.

Here is the procedure for installing SOAPpy:

Go to <http://pywebsvcs.sourceforge.net/> and select Latest Official Release under the SOAPpy section.

2. There are two downloads available. If you are using Windows, download the `pywebsvcs-0.1.10-win32.exe`; otherwise, download the `pywebsvcs-0.1.10.tar.gz`.
3. Decompress the downloaded file, just as you did with `fpconst`.
4. Open a command prompt and navigate to the directory where you decompressed the SOAPpy files.
5. Type `python setup.py install` to run the installation program.

To verify that you installed SOAPpy correctly, run your Python IDE and check the version number.

Example 12.5. Verifying SOAPpy Installation



This version number should match the version number of the SOAPpy archive you downloaded and installed.

12.3. First Steps with SOAP

The heart of SOAP is the ability to call remote functions. There are a number of public access SOAP servers that provide simple functions for

The most popular public access SOAP server is <http://www.xmethods.net/>. This example uses a demonstration function that takes a United States

Example 12.6. Getting the Current Temperature



- 1 You access the remote SOAP server through a proxy class, `SOAPProxy`. The proxy handles all the internals of SOAP for you, including creating the SOAP envelope.
- 2 Every SOAP service has a URL which handles all the requests. The same URL is used for all function calls. This particular service one is `http://www.xmethods.net/soap/temperature`.
- 3 You're creating the `SOAPProxy` with the service URL and the service namespace. This doesn't make any connection to the SOAP server; it simply sets up the proxy.
- 4 Now with everything configured properly, you can actually call remote SOAP methods as if they were local functions. You pass arguments to the method call.

Let's peek under those covers.

Most of the XML request document that gets sent to the server is just boilerplate. Ignore all the namespace declarations; they're going to be t



- 1 The element name is the function name, `calculator` is a dispatcher. Instead of calling separate local methods based on the method name, i
- 2 The function's XML element is contained in a specific namespace, which is the namespace you specified when you created the `calculator` object.
- 3 The arguments of the function also got translated into XML. `calculator` introspects each argument to determine its datatype (in this case it's a str

The XML return document is equally easy to understand, once you know what to ignore. Focus on this fragment within the `calculator`



- 1 The server wraps the function return value within a `calculator` element. By convention, this wrapper element is the name of the function, plus `Response`
- 2 The server returns the response in the same namespace we used in the request, the same namespace we specified when we first create
- 3 The return value is specified, along with its datatype (it's a float). `calculator` uses this explicit datatype to create a Python object of the correct na

12.5. Introducing WSDL

The `Calculator` class proxies local method calls and transparently turns them into invocations of remote SOAP methods. As you've seen, this is a lot of v

Consider this: the previous two sections showed an example of calling a simple remote SOAP method with one argument and one return value.

That shouldn't come as a big surprise. If I wanted to call a local function, I would need to know what package or module it was in (the equivalent

The big difference is introspection. As you saw in Chapter 4, Python excels at letting you discover things about modules and functions at run

WSDL lets you do that with SOAP web services. WSDL stands for "Web Services Description Language". Although designed to be flexible

A WSDL file is just that: a file. More specifically, it's an XML file. It usually lives on the same server you use to access the SOAP web servi

A WSDL file contains a description of everything involved in calling a SOAP web service:

- The service URL and namespace.
- The type of web service (probably function calls using SOAP, although as I mentioned, WSDL is flexible enough to describe a wide
- The list of available functions.
- The arguments for each function.
- The datatype of each argument.
- The return values of each function, and the datatype of each return value.

In other words, a WSDL file tells you everything you need to know to be able to call a SOAP web service.

12.6. Introspecting SOAP Web Services with WSDL

Like many things in the web services arena, WSDL has a long and checkered history, full of political strife and intrigue. I will skip over this

The most fundamental thing that WSDL allows you to do is discover the available methods offered by a SOAP server.

Example 12.8. Discovering The Available Methods



- 1 SOAPpy includes a WSDL parser. At the time of this writing, it was labeled as being in the early stages of development, but I had no
- 2 To use a WSDL file, you again use a proxy class, which takes a single argument: the WSDL file. Note that in this case you are passing
- 3 The WSDL proxy class exposes the available functions as a Python dictionary, so getting the list of available methods is as simple as

Okay, so you know that this SOAP server offers a single method: But how do you call it? The WSDL proxy object can tell you that too.

Example 12.9. Discovering A Method's Arguments



- 1 The dictionary is filled with a SOAPpy-specific structure called `Function`. A `Function` object contains information about one specific function, including
- 2 The function arguments are stored in `args`, which is a Python list of `Parameter` objects that hold information about each parameter.
- 3 Each `Parameter` object contains a `name` attribute, which is the argument name. You are not required to know the argument name to call the function though.
- 4 Each parameter is also explicitly typed, using datatypes defined in XML Schema. You saw this in the wire trace in the previous section.

WSDL also lets you introspect into a function's return values.

Example 12.10. Discovering A Method's Return Values



Beyond the underlying HTTP layer, there are a number of things that can go wrong. SOAPpy does an admirable job hiding the SOAP syntax

Here are a few examples of common mistakes that I've made in using SOAP web services, and the errors they generated.

Example 12.15. Calling a Method With an Incorrectly Configured Proxy

```
#!/usr/bin/python
import sys
import SOAPpy

url = 'http://www.example.com/soap'
proxy = SOAPpy.Proxy(url)

# This is the mistake: proxy is not configured with a proxy
proxy.someMethod(1, 2, 3)
```

- 1 Did you spot the mistake? You're creating a `Proxy` manually, and you've correctly specified the service URL, but you haven't specified the proxy.
- 2 The server responds by sending a SOAP Fault, which SOAPpy turns into a Python exception of type `SOAPpy.SOAPFault`. All errors returned from any SOAP service are of this type.

Misconfiguring the basic elements of the SOAP service is one of the problems that WSDL aims to solve. The WSDL file contains the service

Example 12.16. Calling a Method With the Wrong Arguments

```
#!/usr/bin/python
import sys
import SOAPpy

url = 'http://www.example.com/soap'
proxy = SOAPpy.Proxy(url)

# This is the mistake: calling someMethod with an integer instead of a string
proxy.someMethod(1, 2, 3)
```

- 1 Did you spot the mistake? It's a subtle one: you're calling `someMethod` with an integer instead of a string. As you saw from introspecting the WSDL file, the method expects a string.
- 2 Again, the server returns a SOAP Fault, and the human-readable part of the error gives a clue as to the problem: you're calling a function with the wrong arguments.

It's also possible to write Python code that expects a different number of return values than the remote function actually returns.

Example 12.17. Calling a Method and Expecting the Wrong Number of Return Values



❶ Did you spot the mistake? `only` returns one value, a float, but you've written code that assumes you're getting two values and trying to

What about Google's web service? The most common problem I've had with it is that I forget to set the application key properly.

Example 12.18. Calling a Method With An Application-Specific Error





- ❶ Can you spot the mistake? There's nothing wrong with the calling syntax, or the number of arguments, or the datatypes. The problem is...
- ❷ The Google server responds with a SOAP Fault and an incredibly long error message, which includes a complete Java stack trace. Read...

Further Reading on Troubleshooting SOAP

New developments for SOAPpy (<http://www-106.ibm.com/developerworks/webservices/library/ws-pyth17.html>) steps through trying...

12.9. Summary

SOAP web services are very complicated. The specification is very ambitious and tries to cover many different use cases for web services. The...

Before diving into the next chapter, make sure you're comfortable doing all of these things:

- Connecting to a SOAP server and calling remote methods.

- Loading a WSDL file and introspecting remote methods.

- Debugging SOAP calls with wire traces.

- Troubleshooting common SOAP-related errors.

Chapter 13. Unit Testing

13.1. Introduction to Roman numerals

In previous chapters, you "dived in" by immediately looking at code and trying to understand it as quickly as possible. Now that you have some

In the next few chapters, you're going to write, debug, and optimize a set of utility functions to convert to and from Roman numerals. You saw

The rules for Roman numerals lead to a number of interesting observations:

There is only one correct way to represent a particular number as Roman numerals.1.

The converse is also true: if a string of characters is a valid Roman numeral, it represents only one number (*i.e.* it can only be read one

There is a limited range of numbers that can be expressed as Roman numerals, specifically 1 through 9. (The Romans did have several

There is no way to represent 0 in Roman numerals. (Amazingly, the ancient Romans had no concept of 0 as a number. Numbers were

There is no way to represent negative numbers in Roman numerals.5.

There is no way to represent fractions or non-integer numbers in Roman numerals.6.

Given all of this, what would you expect out of a set of functions to convert to and from Roman numerals?

Requirements

Should return the Roman numeral representation for all integers 1 to 9.

Should fail when given an integer outside the range 1 to 9.

Should fail when given a non-integer number.3.

Should take a valid Roman numeral and return the number that it represents.4.

Should fail when given an invalid Roman numeral.5.

If you take a number, convert it to Roman numerals, then convert that back to a number, you should end up with the number you started

Should always return a Roman numeral using uppercase letters.7.

Should only accept uppercase Roman numerals (*i.e.* it should fail when given lowercase input).8.

Further reading

This site (<http://www.wilkiecollins.demon.co.uk/roman/front.htm>) has more on Roman numerals, including a fascinating history (http://

13.2. Diving in

Now that you've completely defined the behavior you expect from your conversion functions, you're going to do something a little unexpected

This is called unit testing, since the set of two conversion functions can be written and tested as a unit, separate from any larger program they

is included with Python 2.1 and later. Python 2.0 users can download it from <http://pyunit.sourceforge.net/>).

Unit testing is an important part of an overall testing-centric development strategy. If you write unit tests, it is important to write them early (

Before writing code, it forces you to detail your requirements in a useful fashion.

While writing code, it keeps you from over-coding. When all the test cases pass, the function is complete.

When refactoring code, it assures you that the new version behaves the same way as the old version.

When maintaining code, it helps you cover your ass when someone comes screaming that your latest change broke their old code. ("

When writing code in a team, it increases confidence that the code you're about to commit isn't going to break other peoples' code,

13.3. Introducing `itertools`

This is the complete test suite for your Roman numeral conversion functions, which are yet to be written but will eventually be in `itertools`. It is not in `itertools`.

Example 13.1. `itertools`

If you have not already done so, you can download this and other examples (<http://diveintopython.org/download/diveintopython-examples-5>).

```
import sys
import unittest
import itertools

def test():
    test_suite = unittest.TestSuite()
    test_suite.addTest(unittest.TestLoader().loadTestsFromModule(sys.modules[__name__]))
    runner = unittest.TextTestRunner(sys.stdout)
    runner.run(test_suite)
```




Further reading

The PyUnit home page (<http://pyunit.sourceforge.net/>) has an in-depth discussion of using the framework (<http://pyunit.sourceforge.net/>).
The PyUnit FAQ (<http://pyunit.sourceforge.net/pyunit.html>) explains why test cases are stored separately (<http://pyunit.sourceforge.net/>).
Python Library Reference (<http://www.python.org/doc/current/lib/>) summarizes the <http://www.python.org/doc/current/lib/module-unittest.html>.
ExtremeProgramming.org (<http://www.extremeprogramming.org/>) discusses why you should write unit tests (<http://www.extremeprogramming.org/>).
The Portland Pattern Repository (<http://www.c2.com/cgi/wiki/>) has an ongoing discussion of unit tests (<http://www.c2.com/cgi/wiki/>).

13.4. Testing for success

The most fundamental part of unit testing is constructing individual test cases. A test case answers a single question about the code it is testing.

A test case should be able to...

- ...run completely by itself, without any human input. Unit testing is about automation.
- ...determine by itself whether the function it is testing has passed or failed, without a human interpreting the results.
- ...run in isolation, separate from any other test cases (even if they test the same functions). Each test case is an island.

Given that, let's build the first test case. You have the following requirement:

It should return the Roman numeral representation for all integers from 1 to 9.

Example 13.2. ■

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100



- ❶ To write a test case, first subclass the `TestCase` class of the `unittest` module. This class provides many useful methods which you can use in your test cases.
- ❷ This is a list of integer/numeral pairs that I verified manually. It includes the lowest ten numbers, the highest number, every number between 10 and 20, and a few more.
- ❸ Every individual test is its own method, which must take no parameters and return no value. If the method exits normally without raising an exception, the test passes.
- ❹ Here you call the actual `isPrime` function. (Well, the function hasn't been written yet, but once it is, this is the line that will call it.) Notice that you are not trapping any exceptions when you call `isPrime`. This is intentional. `isPrime` shouldn't raise an exception when you call it with a valid integer.
- ❺ Also notice that you are not trapping any exceptions when you call `isPrime`. This is intentional. `isPrime` shouldn't raise an exception when you call it with a valid integer.
- ❻ Assuming the `isPrime` function was defined correctly, called correctly, completed successfully, and returned a value, the last step is to check whether the value is `True` or `False`.

13.5. Testing for failure

It is not enough to test that functions succeed when given good input; you must also test that they fail when given bad input. And not just any bad input.

Remember the other requirements for `isPrime`:

- `isPrime` should fail when given an integer outside the range 1 to 9.
- `isPrime` should fail when given a non-integer number.

In Python, functions indicate failure by raising exceptions, and the `unittest` module provides methods for testing whether a function raises a particular exception.

Example 13.3. Testing bad input to `isPrime`



- ❶ The `test` class of the `Test` provides the `test` method, which takes the following arguments: the exception you're expecting, the function you're testing.
- ❷ Along with testing numbers that are too large, you need to test numbers that are too small. Remember, Roman numerals cannot express zero.
- ❸ Requirement #3 specifies that `test` cannot accept a non-integer number, so here you test to make sure that `test` raises a `ValueError` exception when called with a non-integer.

The next two requirements are similar to the first three, except they apply to `test` instead of `test`.

- ❹ `test` should take a valid Roman numeral and return the number that it represents.
- ❺ `test` should fail when given an invalid Roman numeral.

Requirement #4 is handled in the same way as requirement #1, iterating through a sampling of known values and testing each in turn. Requirement #5 is handled in the same way as requirement #3, testing a sampling of invalid values.

Example 13.4. Testing bad input to `test`

```

class Test:
    def test(self, expected_exception, func, *args, **kwargs):
        """Test a function against an expected exception.

        Args:
            expected_exception: The exception you expect to be raised.
            func: The function to test.
            *args: Arguments to pass to the function.
            **kwargs: Keyword arguments to pass to the function.

        Returns:
            True if the function raises the expected exception, False otherwise.
        """
        try:
            func(*args, **kwargs)
        except expected_exception:
            return True
        else:
            return False

    def test_invalid(self, func, *args, **kwargs):
        """Test a function against invalid input.

        Args:
            func: The function to test.
            *args: Arguments to pass to the function.
            **kwargs: Keyword arguments to pass to the function.

        Returns:
            True if the function raises a ValueError, False otherwise.
        """
        return self.test(ValueError, func, *args, **kwargs)

    def test_valid(self, func, *args, **kwargs):
        """Test a function against valid input.

        Args:
            func: The function to test.
            *args: Arguments to pass to the function.
            **kwargs: Keyword arguments to pass to the function.

        Returns:
            True if the function returns the expected value, False otherwise.
        """
        return self.test(None, func, *args, **kwargs)

```

- ❶ Not much new to say about these; the pattern is exactly the same as the one you used to test bad input to `test`. I will briefly note that you should use `self` to refer to the instance of the class.

13.6. Testing for sanity

Often, you will find that a unit of code contains a set of reciprocal functions, usually in the form of conversion functions where one converts from one unit to another.

Consider this requirement:

If you take a number, convert it to Roman numerals, then convert that back to a number, you should end up with the number you started with.

Example 13.5. Testing `test` against `test`

```

class Test:
    def test(self, expected_exception, func, *args, **kwargs):
        """Test a function against an expected exception.

        Args:
            expected_exception: The exception you expect to be raised.
            func: The function to test.
            *args: Arguments to pass to the function.
            **kwargs: Keyword arguments to pass to the function.

        Returns:
            True if the function raises the expected exception, False otherwise.
        """
        try:
            func(*args, **kwargs)
        except expected_exception:
            return True
        else:
            return False

    def test_invalid(self, func, *args, **kwargs):
        """Test a function against invalid input.

        Args:
            func: The function to test.
            *args: Arguments to pass to the function.
            **kwargs: Keyword arguments to pass to the function.

        Returns:
            True if the function raises a ValueError, False otherwise.
        """
        return self.test(ValueError, func, *args, **kwargs)

    def test_valid(self, func, *args, **kwargs):
        """Test a function against valid input.

        Args:
            func: The function to test.
            *args: Arguments to pass to the function.
            **kwargs: Keyword arguments to pass to the function.

        Returns:
            True if the function returns the expected value, False otherwise.
        """
        return self.test(None, func, *args, **kwargs)

```

- ❶ You've seen the `test` function before, but here it is called with two arguments, which returns a list of integers starting at the first argument.
- ❷ I just wanted to mention in passing that `self` is not a keyword in Python; here it's just a variable name like any other.
- ❸ The actual testing logic here is straightforward: take a number (`num`), convert it to a Roman numeral (`roman`), then convert it back to a number (`num2`).





- ❶ Running the script runs `test.py` which runs each test case, which is to say each method defined in each class within `test.py`. For each test case, it prints the result.
- ❷ For each failed test case, `test.py` displays the trace information showing exactly what happened. In this case, the call to `test()` also called `test_raise()` raised a `ValueError`.
- ❸ After the detail, `test.py` displays a summary of how many tests were performed and how long it took.
- ❹ Overall, the unit test failed because at least one test case did not pass. When a test case doesn't pass, `test.py` distinguishes between failures and errors.

14.2. Stage 2

Now that you have the framework of the `unittest` module laid out, it's time to start writing code and passing test cases.

Example 14.3. `test.py`

This file is available in `examples` in the examples directory.

If you have not already done so, you can download this and other examples (<http://diveintopython.org/download/diveintopython-examples-5>).



```
1
2
3
```

- 1 `roman` is a tuple of tuples which defines three things:
 - The character representations of the most basic Roman numerals. Note that this is not just the single-character Roman numeral.
 - The order of the Roman numerals. They are listed in descending value order, from `M` all the way down to `I`.
 - The value of each Roman numeral. Each inner tuple is a pair of `(char, value)`.
- 2 Here's where your rich data structure pays off, because you don't need any special logic to handle the subtraction rule. To convert to Roman...

Example 14.4. How `roman` works

If you're not clear how `roman` works, add a `print` statement to the end of the `while` loop:

```
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
```

So `roman` appears to work, at least in this manual spot check. But will it pass the unit testing? Well no, not entirely.

Example 14.5. Output of `roman` against `test`

Remember to run `test` with the `-v` command-line flag to enable verbose mode.

```
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
```

- 1 `roman` does, in fact, always return uppercase, because `roman` defines the Roman numeral representations as uppercase. So this test passes already.
- 2 Here's the big news: this version of the `roman` function passes the known values test. Remember, it's not comprehensive, but it does put the...
- 3 However, the function does not "work" for bad values; it fails every single bad input test. That makes sense, because you didn't include...



14.3. stage 3

Now that `is_int` behaves correctly with good input (integers from `1` to `9`) it's time to make it behave correctly with bad input (everything else).

Example 14.6. `is_int`

This file is available in `ch14` in the examples directory.

If you have not already done so, you can download this and other examples (<http://diveintopython.org/download/diveintopython-examples-5>).

```
def romanize(n):
    """Convert an integer to a Roman numeral.

    The range of integers that can be converted is from 1 to 3999.
    The range of Roman numerals that can be converted is from I to
    MMMCMXCIX.

    Raises:
    ValueError: If n is not an integer or if n is not in the range
    1 to 3999.
    """
    if not isinstance(n, int):
        raise ValueError("n must be an integer")
    if n < 1 or n > 3999:
        raise ValueError("n must be in the range 1 to 3999")
    # The rest of the function is unchanged.
    ...
```

- ❶ This is a nice Pythonic shortcut: multiple comparisons at once. This is equivalent to `n < 1 or n > 3999` but it's much easier to read. This is the range check.
- ❷ You raise exceptions yourself with the `raise` statement. You can raise any of the built-in exceptions, or you can raise any of your custom exceptions.
- ❸ This is the non-integer check. Non-integers can not be converted to Roman numerals.
- ❹ The rest of the function is unchanged.

Example 14.7. Watching and handling bad input

```
def romanize(n):
    """Convert an integer to a Roman numeral.

    The range of integers that can be converted is from 1 to 3999.
    The range of Roman numerals that can be converted is from I to
    MMMCMXCIX.

    Raises:
    ValueError: If n is not an integer or if n is not in the range
    1 to 3999.
    """
    if not isinstance(n, int):
        raise ValueError("n must be an integer")
    if n < 1 or n > 3999:
        raise ValueError("n must be in the range 1 to 3999")
    # The rest of the function is unchanged.
    ...
```



Example 14.8. Output of `test_against`



- 1 `test` still passes the known values test, which is comforting. All the tests that passed in stage 2 still pass, so the latest code hasn't broken anything.
- 2 More exciting is the fact that all of the bad input tests now pass. This test, `test_bad_input`, passes because of the `isinstance` check. When a non-integer is passed to `isinstance`, it returns `False`.
- 3 This test, `test_exception`, passes because of the `try` check, which raises an `AssertionError`, which is what `test` is looking for.





❶ You're down to 6 failures, and all of them involve the known values test, the three separate bad input tests, the case check, and the sa

The most important thing that comprehensive unit testing can tell you is when to stop coding. When all the unit tests for a function pass, stop

14.4. stage 4

Now that is done, it's time to start coding. Thanks to the rich data structure that maps individual Roman numerals to integer values, this is n

Example 14.9.

This file is available in the examples directory.

If you have not already done so, you can download this and other examples (<http://diveintopython.org/download/diveintopython-examples-5>).

```
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
```

❶ The pattern here is the same as `roman`. You iterate through your Roman numeral data structure (a tuple of tuples), and instead of matching the

Example 14.10. How `roman` works

If you're not clear how `roman` works, add a `print` statement to the end of the `while` loop:

```
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
```



Example 14.11. Output of `is_prime` against `is_prime`



- ❶ Two pieces of exciting news here. The first is that `is_prime` works for good input, at least for all the known values you test.
- ❷ The second is that the sanity check also passed. Combined with the known values tests, you can be reasonably sure that both `is_prime` and `is_prime` work.



14.5. Stage 5

Now that it works properly with good input, it's time to fit in the last piece of the puzzle: making it work properly with bad input. That means

If you're not familiar with regular expressions and didn't read Chapter 7, *Regular Expressions*, now would be a good time.

As you saw in Section 7.3, "Case Study: Roman Numerals", there are several simple rules for constructing a Roman numeral, using the letters

Characters are additive. **I**s **1**, **II**s **2** and **III**s **3**. **V**s **5** and **X**s **10**, **L**s **50** and **C**s **100**.

The tens characters (**X**, **L** and **C**) can be repeated up to three times. At **4** you need to subtract from the next highest fives character. You

Similarly, at **9** you need to subtract from the next highest tens character: **8**s **VIII** but **9**s **IX** "less than **10**", not **IIIIIIII** since the **I** character can not be

The fives characters can not be repeated. **0**s always represented as **X** never as **VL** **0**s always **C** never **L**.

Roman numerals are always written highest to lowest, and read left to right, so order of characters matters very much. **0**s **00**s a comp

Example 14.12.

This file is available in `examples` directory.

If you have not already done so, you can download this and other examples (<http://diveintopython.org/download/diveintopython-examples-5>).

```


```



- 1 This is just a continuation of the pattern you discussed in Section 7.3, “Case Study: Roman Numerals”. The tens places is either ~~IX~~ ~~IX~~
- 2 Having encoded all that logic into a regular expression, the code to check for invalid Roman numerals becomes trivial. If `is_invalid` returns an o

At this point, you are allowed to be skeptical that that big ugly regular expression could possibly catch all the types of invalid Roman numerals.

Example 14.13. Output of `is_invalid` against `is_valid`





- ❶ One thing I didn't mention about regular expressions is that, by default, they are case-sensitive. Since the regular expression `was` expr
- ❷ More importantly, the bad input tests pass. For instance, the malformed antecedents test checks cases like `As` you've seen, this does n
- ❸ In fact, all the bad input tests pass. This regular expression catches everything you could think of when you made your test cases.
- ❹ And the anticlimax award of the year goes to the word `"` which is printed by the `module` when all the tests pass.

When all of your tests pass, stop coding.



Now you can fix the bug.

Example 15.4. Fixing the bug (1)

This file is available in `01` in the examples directory.



- 1 Only two lines of code are required: an explicit check for an empty string, and a `return` statement.

Example 15.5. Output of `is_palindrome` against `test.py`



- 1 The blank string test case now passes, so the bug is fixed.
- 2 All the other test cases still pass, which means that this bug fix didn't break anything else. Stop coding.





- ❶ The existing known values don't change (they're all still reasonable values to test), but you need to add a few more in the **Q** range. Here
- ❷ The definition of "large input" has changed. This test used to call **D** with **Q** and expect an error; now that **Q** are good values, you need to bu

- ③ The definition of "too many repeated numerals" has also changed. This test used to call `is_valid` with `M` and expect an error; now that `M`s consider
- ④ The sanity check and case checks loop through every number in the range, from `1` to `9`. Since the range has now expanded, these loops m

Now your test cases are up to date with the new requirements, but your code is not, so you expect several of the test cases to fail.

Example 15.7. Output of `is_valid` against `is_valid`



- ① Our case checks now fail because they loop from `1` to `9` but `is_valid` only accepts numbers from `1` to `9` so it will fail as soon the test case hits `0`
- ② The `known values` test will fail as soon as it hits `M` because `is_valid` still thinks this is an invalid Roman numeral.
- ③ The `known values` test will fail as soon as it hits `0` because `is_valid` still thinks this is out of range.
- ④ The sanity check will also fail as soon as it hits `0` because `is_valid` still thinks this is out of range.





Now that you have test cases that fail due to the new requirements, you can think about fixing the code to bring it in line with the test cases. (

Example 15.8. Coding the new requirements (

This file is available in in the examples directory.





- ① Only needs one small change, in the range check. Where you used to check `0` you now check `1`. And you change the error message that...
- ② You don't need to make any changes to `it` at all. The only change is to `it` if you look closely, you'll notice that you added another option...

You may be skeptical that these two small changes are all that you need. Hey, don't take my word for it; see for yourself:

Example 15.9. Output of `it` against `it`





- 1 All the test cases pass. Stop coding.

Comprehensive unit testing means never having to rely on a programmer who says "Trust me."

15.3. Refactoring

The best thing about comprehensive unit testing is not the feeling you get when all your test cases finally pass, or even the feeling you get when...

Refactoring is the process of taking working code and making it work better. Usually, "better" means "faster", although it can also mean "using less code".

Here, "better" means "faster". Specifically, the function is slower than it needs to be, because of that big nasty regular expression that you use to...

Example 15.10. Compiling regular expressions



- 1 This is the syntax you've seen before: `re.match(pattern, string)` takes a regular expression as a string (`pattern`) and a string to match against it (`string`). If the pattern matches, it returns a match object; otherwise, it returns `None`.
- 2 This is the new syntax: `re.compile(pattern)` takes a regular expression as a string and returns a pattern object. Note there is no string to match here. Compiling a regular expression is a one-time cost.
- 3 The compiled pattern object returned from `re.compile()` has several useful-looking functions, including several (like `match` and `search`) that are available directly on the pattern object.
- 4 Calling the compiled pattern object's `match` function with the string `string` accomplishes the same thing as calling `re.match` with both the regular expression and the string.

Whenever you are going to use a regular expression more than once, you should compile it to get a pattern object, then call the methods on the pattern object.

Example 15.11. Compiled regular expressions in Python

This file is available in `examples` in the examples directory.

If you have not already done so, you can download this and other examples (<http://diveintopython.org/download/diveintopython-examples-5.0.tar.gz>).



- 1 This looks very similar, but in fact a lot has changed. `re.compile()` is no longer a string; it is a pattern object which was returned from `re.compile()`.
- 2 That means that you can call methods on `re.compile()` directly. This will be much, much faster than calling `re.compile()` every time. The regular expression is c...

So how much faster is it to compile regular expressions? See for yourself:

Example 15.12. Output of `re.compile()` against `re.compile()`



- 1 Just a note in passing here: this time, I ran the unit test *without* the `-v` option, so instead of the full `re.compile()` for each test, you only get a dot for ea...
- 2 You ran `re.compile()` tests in `re.compile()` seconds, compared to `re.compile()` seconds without precompiling the regular expressions. That's an `re.compile()` improvement overall, and re...
- 3 Oh, and in case you were wondering, precompiling the regular expression didn't break anything, and you just proved it.

There is one other performance optimization that I want to try. Given the complexity of regular expression syntax, it should come as no surpr...

Example 15.13. `re.compile()`

This file is available in `re.compile()` in the examples directory.

If you have not already done so, you can download this and other examples (<http://diveintopython.org/download/diveintopython-examples-5>).





① You have replaced `re.match` with `re.fullmatch`. Both mean the same thing: "match 0 to 4 characters". Similarly, `re.compile` became `re.compile` ("match 0 to 3 characters") and so on.

This form of the regular expression is a little shorter (though not any more readable). The big question is, is it any faster?

Example 15.14. Output of `perf` against `re`



- ① Overall, the unit tests run 2% faster with this form of regular expression. That doesn't sound exciting, but remember that the `re` function is used everywhere.
- ② More important than any performance boost is the fact that the module still works perfectly. This is the freedom I was talking about earlier.

One other tweak I would like to make, and then I promise I'll stop refactoring and put this module to bed. As you've seen repeatedly, regular

Example 15.15. `re`

This file is available in `examples` directory.

If you have not already done so, you can download this and other examples (<http://diveintopython.org/download/diveintopython-examples-5>).



① The `re` function can take an optional second argument, which is a set of one or more flags that control various options about the compilation.

Example 15.16. Output of `perf` against `re`





- ❶ This new, "verbose" version runs at exactly the same speed as the old version. In fact, the compiled pattern objects are the same, since
- ❷ This new, "verbose" version passes all the same tests as the old version. Nothing has changed, except that the programmer who comes

15.4. Postscript

A clever reader read the previous section and took it to the next level. The biggest headache (and performance drain) in the program as it is c

And best of all, he already had a complete set of unit tests. He changed over half the code in the module, but the unit tests stayed the same, so

Example 15.17. 🐍

This file is available in 📁 in the examples directory.

If you have not already done so, you can download this and other examples (<http://diveintopython.org/download/diveintopython-examples-5>).



```

.
-
█
█

```

So how fast is it?

Example 15.18. Output of `perf` against `perf`

```

.
-
█
█

```

Remember, the best performance you ever got in the original version was 13 tests in 3.315 seconds. Of course, it's not entirely a fair comparison.

The moral of the story?

Simplicity is a virtue.

Especially when regular expressions are involved.

And unit tests can give you the confidence to do large-scale refactoring... even if you didn't write the original code.

15.5. Summary

Unit testing is a powerful concept which, if properly implemented, can both reduce maintenance costs and increase flexibility in any long-term

This chapter covered a lot of ground, and much of it wasn't even Python-specific. There are unit testing frameworks for many languages, all

- Designing test cases that are specific, automated, and independent.

- Writing test cases *before* the code they are testing.

- Writing tests that test good input and check for proper results.

- Writing tests that test bad input and check for proper failures.

- Writing and updating test cases to illustrate bugs or reflect new requirements.

- Refactoring mercilessly to improve performance, scalability, readability, maintainability, or whatever other -ility you're lacking.

Additionally, you should be comfortable doing all of the following Python-specific things:

- Subclassing `unittest.TestCase` and writing methods for individual test cases.

- Using `assertEqual` to check that a function returns a known value.

- Using `assertRaises` to check that a function raises a known exception.

- Calling `unittest.main()` in your `__main__` clause to run all your test cases at once.

- Running unit tests in verbose or regular mode.

Further reading

XProgramming.com (<http://www.xprogramming.com/>) has links to download unit testing frameworks (<http://www.xprogramming.com/>)



- ❶ The first 5 tests are from `test_introspection.py` which tests the example script from Chapter 4, *The Power Of Introspection*.
- ❷ The next 5 tests are from `test_first_program.py` which tests the example script from Chapter 2, *Your First Python Program*.
- ❸ The rest are from `test_testing.py` which you studied in depth in Chapter 13, *Unit Testing*.

16.2. Finding the path

When running Python scripts from the command line, it is sometimes useful to know where the currently running script is located on disk.

This is one of those obscure little tricks that is virtually impossible to figure out on your own, but simple to remember once you see it. The key

Example 16.3. `__file__`

If you have not already done so, you can download this and other examples (<http://diveintopython.org/download/diveintopython-examples-5>).



- ❶ Regardless of how you run a script, `__file__` will always contain the name of the script, exactly as it appears on the command line. This may or may not include a path.
- ❷ `os.path.basename()` takes a filename as a string and returns the directory path portion. If the given filename does not include any path information, `os.path.basename()` returns the filename.
- ❸ `os.path.abspath()` is the key here. It takes a pathname, which can be partial or even blank, and returns a fully qualified pathname.

deserves further explanation. It is very flexible; it can take any kind of pathname.

Example 16.4. Further explanation of `os.path.abspath()`

```
1 #!/usr/bin/env python
2 """
3 Example 16.4. Further explanation of
4 os.path.abspath().
5 """
6 import os
7
8 # Returns the current working directory.
9 print(os.path.abspath('.'))
10
11 # Calling os.path.abspath() with an empty string
12 # returns the current working directory, same as
13 # os.getcwd().
14 print(os.path.abspath(''))
15
16 # Calling os.path.abspath() with a partial
17 # pathname constructs a fully qualified
18 # pathname out of it, based on the current
19 # working directory.
20 print(os.path.abspath('foo'))
21
22 # Calling os.path.abspath() with a full
23 # pathname simply returns it.
24 print(os.path.abspath('/usr/bin/python'))
25
26 # Also normalizes the pathname it returns.
27 # Note that this example worked even though
28 # I don't actually have a 'foo' directory.
29 # os.path.abspath() never checks if the
30 # directory exists.
```

1 Returns the current working directory.

2 Calling `os.path.abspath()` with an empty string returns the current working directory, same as `os.getcwd()`.

3 Calling `os.path.abspath()` with a partial pathname constructs a fully qualified pathname out of it, based on the current working directory.

4 Calling `os.path.abspath()` with a full pathname simply returns it.

5 Also *normalizes* the pathname it returns. Note that this example worked even though I don't actually have a 'foo' directory. `os.path.abspath()` never checks if the directory exists.

The pathnames and filenames you pass to `os.path.abspath()` do not need to exist.

`os.path.abspath()` not only constructs full path names, it also normalizes them. That means that if you are in the `foo` directory, `os.path.abspath('foo')` will return `foo`. It normalizes the path.

Example 16.5. Sample output from `os.path.dirname()`

```
1 #!/usr/bin/env python
2 """
3 Example 16.5. Sample output from
4 os.path.dirname().
5 """
6 import os
7
8 # Returns the full path of the script.
9 print(os.path.dirname(os.path.abspath(__file__)))
10
11 # If the script is run by using a partial
12 # pathname, os.path.dirname() will still
13 # contain exactly what appears on the
14 # command line. os.path.dirname() will
15 # then give you a partial path.
16 print(os.path.dirname('foo'))
17
18 # If the script is run from the current
19 # directory without giving any path,
20 # os.path.dirname() will simply return
21 # an empty string. Given an empty
22 # string, os.path.dirname() returns
23 # an empty string.
```

1 In the first case, `os.path.dirname()` includes the full path of the script. You can then use the `os.path.basename()` function to strip off the script name and return the full directory.

2 If the script is run by using a partial pathname, `os.path.dirname()` will still contain exactly what appears on the command line. `os.path.dirname()` will then give you a partial path.

3 If the script is run from the current directory without giving any path, `os.path.dirname()` will simply return an empty string. Given an empty string, `os.path.dirname()` returns an empty string.

Like the other functions in the `os` module, `chdir` is cross-platform. Your results will look slightly different than my examples if you're running **Addendum**. One reader was dissatisfied with this solution, and wanted to be able to run all the unit tests in the current directory, not the directory where the script is located.

Example 16.6. Running scripts in the current directory

```

import os
import sys

current_dir = os.getcwd()
sys.path.append(current_dir)

# ... rest of the function ...

```

- ❶ Instead of setting `dir` to the directory where the currently running script is located, you set it to the current working directory instead. This is the directory where you are running the script from.
- ❷ Append this directory to the Python library search path, so that when you dynamically import the unit test modules later, Python can find them.
- ❸ The rest of the function is the same.

This technique will allow you to re-use this script on multiple projects. Just put the script in a common directory, then change to the project's directory and run the script.

16.3. Filtering lists revisited

You're already familiar with using list comprehensions to filter lists. There is another way to accomplish this same thing, which some people find more intuitive. The function `filter` takes a function and a list and returns a list of elements from the original list for which the function returns a true value.

Got all that? It's not as difficult as it sounds.

Example 16.7. Introducing `filter`

```

def is_even(n):
    return n % 2 == 0

my_list = [1, 2, 3, 4, 5, 6, 7, 8, 9, 10]

filtered_list = filter(is_even, my_list)

print(list(filtered_list))

```

- ❶ Uses the built-in mod function `%` to return `True` if `n` is odd and `False` if `n` is even.
- ❷ Takes two arguments, a function (`is_even`) and a list (`my_list`). It loops through the list and calls `is_even` with each element. If `is_even` returns a true value (remember, `True` is considered true), the element is included in the returned list.
- ❸ You could accomplish the same thing using list comprehensions, as you saw in Section 4.5, "Filtering Lists".
- ❹ You could also accomplish the same thing with a `for` loop. Depending on your programming background, this may seem more "straightforward".

Example 16.8. `filter` in action

```

# ... code for Example 16.8 ...

```

❶

- ❶ As you saw in Section 16.2, “Finding the path”, `os.path` may contain the full or partial pathname of the directory of the currently running script.
- ❷ This is a compiled regular expression. As you saw in Section 15.3, “Refactoring”, if you’re going to use the same regular expression on a list, you should compile it once and reuse the compiled object.
- ❸ For each element in the list, you’re going to call the `match` method of the compiled regular expression object. If the regular expression matches, you’ll append the element to the new list.

Historical note. Versions of Python prior to 2.0 did not have list comprehensions, so you couldn’t filter using list comprehensions; the `filter` function was used instead.

Example 16.9. Filtering using list comprehensions instead

```

def filter_paths(paths):
    """Return a list of paths that contain the string 'foo'"""
    return [path for path in paths if 'foo' in path]

```

- ❶ This will accomplish exactly the same result as using the `filter` function. Which way is more expressive? That’s up to you.

16.4. Mapping lists revisited

You’re already familiar with using list comprehensions to map one list into another. There is another way to accomplish the same thing, using the `map` function.

Example 16.10. Introducing `map`

```

def double(x):
    """Return twice the value of x"""
    return 2 * x

paths = ['foo', 'bar', 'baz', 'qux', 'quux', 'quuz', 'quaa', 'quoo', 'quuu', 'quuz']
double_paths = map(double, paths)

```

- ❶ `map` is a function that takes a function and a list as arguments. It returns a list of the results of calling the function with each element of the list in order. In this case, the function simply multiplies each element by 2.
- ❷ You could accomplish the same thing with a list comprehension. List comprehensions were first introduced in Python 2.0; `map` has been around since Python 1.0.
- ❸ You could, if you insist on thinking like a Visual Basic programmer, use a `for` loop to accomplish the same thing.

Example 16.11. `map` with lists of mixed datatypes

```

def double(x):
    """Return twice the value of x"""
    return 2 * x

paths = ['foo', 'bar', 'baz', 'qux', 'quux', 'quuz', 'quaa', 'quoo', 'quuu', 'quuz']
double_paths = map(double, paths)

```

- ❶ As a side note, I’d like to point out that `map` works just as well with lists of mixed datatypes, as long as the function you’re using correctly handles the different types.

All right, enough play time. Let’s look at some real code.

Example 16.12. `pin`



- 1 As you saw in Section 4.7, “Using lambda Functions”, `def` defines an inline function. And as you saw in Example 6.17, “Splitting Pathnames”, `os.path.splitext` splits a filename into a base name and an extension.
- 2 Calling `map` takes each filename listed in `files` and passes it to the function `os.path.splitext` and returns a list of the return values of each of those function calls. In this case, the return values are tuples of (base name, extension).

As you’ll see in the rest of the chapter, you can extend this type of data-centric thinking all the way to the final goal, which is to define and use a function that takes a list of data and returns a list of data.

16.5. Data-centric programming

By now you’re probably scratching your head wondering why this is better than using `for` loops and straight function calls. And that’s a perfectly reasonable question.

In this case, you started with no data at all; the first thing you did was get the directory path of the current script, and got a list of files in that directory.

However, you knew you didn’t care about all of those files, only the ones that were actually test suites. You had *too much data*, so you needed to filter it.

Now you had the filenames of each of the test suites (and only the test suites, since everything else had been filtered out), but you really wanted to know which ones were actually test suites.

Instead of `if` you could have used a `for` loop with an `if` statement. Instead of `map` you could have used a `for` loop with a function call. But using `for` loops like this is a pain.

Oh sure, you could try to be fancy and delete elements in place without creating a new list. But you’ve been burned by that before. Trying to do that is a pain.

I resisted list comprehensions when I first learned Python, and I resisted `map` and `filter` even longer. I insisted on making my life more difficult, sticking with `for` loops and `if` statements.

Let it all go. Busywork code is not important. Data is important. And data is not difficult. It’s only data. If you have too much, filter it. If it’s too small, use `map` and `filter`.

16.6. Dynamically importing modules

OK, enough philosophizing. Let’s talk about dynamically importing modules.

First, let’s look at how you normally import modules. The `import` syntax looks in the search path for the named module and imports it by name. You can also use `from` to import specific attributes from a module.

Example 16.13. Importing multiple modules at once



- 1 This imports four modules at once: `sys` (for system functions and access to the command line parameters), `os` (for operating system functions), `re` (for regular expressions), and `random` (for random number generation).

Now let’s do the same thing, but with dynamic imports.

Example 16.14. Importing modules dynamically



- 1 The built-in `import` function accomplishes the same goal as using the `from` statement, but it's an actual function, and it takes a string as an argument.
- 2 The variable `sys` is now the `sys` module, just as if you had said `import sys`. The variable `os` is now the `os` module, and so forth.

So `import` imports a module, but takes a string argument to do it. In this case the module you imported was just a hard-coded string, but it could just as easily be a variable.

Example 16.15. Importing a list of modules dynamically

```

1 import sys
2 import os
3
4 modules = ['sys', 'os']
5
6 for module in modules:
7     import module
8
9 sys.path.append('..')
10
11 from . import *

```

- 1 `modules` is just a list of strings. Nothing fancy, except that the strings happen to be names of modules that you could import, if you wanted to.
- 2 Surprise, you wanted to import them, and you did, by mapping the `import` function onto the list. Remember, this takes each element of the list and imports it.
- 3 So now from a list of strings, you've created a list of actual modules. (Your paths may be different, depending on your operating system.)
- 4 To drive home the point that these are real modules, let's look at some module attributes. Remember, `sys` is the `sys` module, so `sys.__name__` is `'sys'`. All the other modules have their own names.

Now you should be able to put this all together and figure out what most of this chapter's code sample is doing.

16.7. Putting it all together

You've learned enough now to deconstruct the first seven lines of this chapter's code sample: reading a directory and importing selected modules.

Example 16.16. The `import` function

```

1 import sys
2 import os
3
4 modules = ['sys', 'os']
5
6 for module in modules:
7     import module
8
9 sys.path.append('..')
10
11 from . import *

```

Let's look at it line by line, interactively. Assume that the current directory is `..` which contains the examples that come with this book, including `example16.15.py`.

Example 16.17. Step 1: Get all the files



- ❶ As you saw in Section 16.6, “Dynamically importing modules”, you can use a combination of `importlib` and `importlib.util` to map a list of module names (as in `modules`) to a list of module objects.
- ❷ `modules` is now a list of modules, fully accessible like any other module.
- ❸ The last module in the list is the `__main__` module, just as if you had said `import *`.

Example 16.21. Step 5: Loading the modules into a test suite



- ❶ These are real module objects. Not only can you access them like any other module, instantiate classes and call functions, you can also introspect them.
- ❷ Finally, you wrap the list of module objects into one big test suite. The `unittest` module has no problem traversing this tree of nested test suites within `unittest`.

This introspection process is what the `unittest` module usually does for us. Remember that magic-looking `__import__` function that our individual test modules call.

Example 16.22. Step 6: Telling `unittest` to use your test suite



- ❶ Instead of letting the `unittest` module do all its magic for us, you’ve done most of it yourself. You’ve created a function (`run_tests`) that imports the modules and runs the tests.

16.8. Summary

The `run_tests` program and its output should now make perfect sense.

You should now feel comfortable doing all of these things:

- Manipulating path information from the command line.
- Filtering lists using `filter` instead of list comprehensions.
- Mapping lists using `map` instead of list comprehensions.
- Dynamically importing modules.

[Technically, the second argument to `unittest` can be any sequence, including lists, tuples, and custom classes that act like lists by defining the `__iter__` special method.]

Again, I should point out that `len` can take a list, a tuple, or any object that acts like a sequence. See previous footnote about `len`

Chapter 17. Dynamic functions

17.1. Diving in

I want to talk about plural nouns. Also, functions that return other functions, advanced regular expressions, and generators. Generators are new

If you haven't read Chapter 7, *Regular Expressions*, now would be a good time. This chapter assumes you understand the basics of regular expressions.

English is a schizophrenic language that borrows from a lot of other languages, and the rules for making singular nouns into plural nouns are

If you grew up in an English-speaking country or learned English in a formal school setting, you're probably familiar with the basic rules:

If a word ends in S, X, or Z, add ES. "Bass" becomes "basses", "fax" becomes "faxes", and "waltz" becomes "waltzes".1.

If a word ends in a noisy H, add ES; if it ends in a silent H, just add S. What's a noisy H? One that gets combined with other letters to

If a word ends in Y that sounds like I, change the Y to IES; if the Y is combined with a vowel to sound like something else, just add S.

If all else fails, just add S and hope for the best.4.

(I know, there are a lot of exceptions. "Man" becomes "men" and "woman" becomes "women", but "human" becomes "humans". "Mouse" becomes "mice".)

Other languages are, of course, completely different.

Let's design a module that pluralizes nouns. Start with just English nouns, and just these four rules, but keep in mind that you'll inevitably need

17.2. Stage 1

So you're looking at words, which at least in English are strings of characters. And you have rules that say you need to find different combinations

Example 17.1.

```
import re
```

```
def pluralize(word):
```

```
    return re.sub(r'([sxz])$', r'\1es', word)
```

```
    return re.sub(r'([sh])$', r'\1s', word)
```

```
    return re.sub(r'([aiouy])$', r'\1ies', word)
```

```
    return word + 's'
```

```
pluralize('bass')
```

```
pluralize('fax')
```

```
pluralize('waltz')
```

```
pluralize('man')
```

```
pluralize('woman')
```

1 OK, this is a regular expression, but it uses a syntax you didn't see in Chapter 7, *Regular Expressions*. The square brackets mean "match any of these characters".

2 This function performs regular expression-based string substitutions. Let's look at it in more detail.

Example 17.2. Introducing

```
import re
```

```
def pluralize(word):
```

```
    return re.sub(r'([sxz])$', r'\1es', word)
```

```
    return re.sub(r'([sh])$', r'\1s', word)
```

```
    return re.sub(r'([aiouy])$', r'\1ies', word)
```

```
    return word + 's'
```

```
pluralize('bass')
```



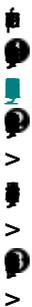
- 1 Does the string `Mountain` contain `ab` or `c`? Yes, it contains `a`
- 2 OK, now find `ab` or `c` and replace it with `0`. It becomes `M0untain`
- 3 The same function turns `0` into `0`
- 4 You might think this would turn `0` into `0` but it doesn't. `re.sub` replaces *all* of the matches, not just the first one. So this regular expression turns `0` into `0`

Example 17.3. Back to `re.sub`



- 1 Back to the `re.sub` function. What are you doing? You're replacing the end of string with `0`. In other words, adding `0` to the string. You could also use `re.sub(r'$', '0', s)`
- 2 Look closely, this is another new variation. The `^` as the first character inside the square brackets means something special: negation. `[^aeiou]` means "not a, e, i, o, or u".
- 3 Same pattern here: match words that end in `Y`, where the character before the `Y` is *not* `a`, `e`, `i`, `o`, or `u`. You're looking for words that end in `Y` where the character before the `Y` is not `a`, `e`, `i`, `o`, or `u`.

Example 17.4. More on negation regular expressions



- 1 `my` matches this regular expression, because it ends in `y` and `ds` is not `a`, `e`, `i`, `o`, or `u`
- 2 `by` does not match, because it ends in `y` and you specifically said that the character before the `y` could not be `o`. `dy` does not match, because it ends in `y` and you specifically said that the character before the `y` could not be `o`.
- 3 `py` does not match, because it does not end in `y`

Example 17.5. More on `re.sub`





- ❶ This regular expression turns `foo` into `foo` and `foo` into `foo` which is what you wanted. Note that it would also turn `foo` into `foo` but that will never happen
- ❷ Just in passing, I want to point out that it is possible to combine these two regular expressions (one to find out if the rule applies, and a

Regular expression substitutions are extremely powerful, and the `re` syntax makes them even more powerful. But combining the entire operation

17.3. Stage 2

Now you're going to add a level of abstraction. You started by defining a list of rules: if this, then do that, otherwise go to the next rule. Let's

Example 17.6.



- ❶ This version looks more complicated (it's certainly longer), but it does exactly the same thing: try to match four different rules, in order



- ❶ This is the same set of rules as you defined in stage 2. The only difference is that instead of defining named functions like `land` and `by` you have...
- ❷ Note that the `rule` function hasn't changed at all. It iterates through a set of rule functions, checks the first rule, and if it returns a true value...

Now to add a new rule, all you need to do is define the functions directly in the `rules` list itself: one match rule, and one apply rule. But defining the...

17.5. Stage 4

Let's factor out the duplication in the code so that defining new rules can be easier.

Example 17.9.



- ❶ `rule` is a function that builds other functions dynamically. It takes `rule` and `apply` (actually it takes a tuple, but more on that in a minute), and you can...
- ❷ Building the apply function works the same way. The apply function is a function that takes one parameter, and calls `rule` with the `land` and `by`...
- ❸ Finally, the `rule` function returns a tuple of two values: the two functions you just created. The constants you defined within those functions...

If this is incredibly confusing (and it should be, this is weird stuff), it may become clearer when you see how to use it.

Example 17.10. Continued



- ❶ Our pluralization rules are now defined as a series of strings (not functions). The first string is the regular expression that you would use...
- ❷ This line is magic. It takes the list of strings in `rules` and turns them into a list of functions. How? By mapping the strings to the `rule` function, which...

I swear I am not making this up: `rules` ends up with exactly the same list of functions as the previous example. Unroll the `rule` definition, and you'll get...

Example 17.11. Unrolling the rules definition



```
(
)
(
)
(
)
)
```

Example 17.12. Finishing up

```


```

Since the list is the same as the previous example, it should come as no surprise that the function hasn't changed. Remember, it's com

Just in case that wasn't mind-blowing enough, I must confess that there was a subtlety in the definition of that I skipped over. Let's go back

Example 17.13. Another look at

```


```

Notice the double parentheses? This function doesn't actually take three parameters; it actually takes one parameter, a tuple of three el

Example 17.14. Expanding tuples when calling functions

```


```

The proper way to call the function `os` with a tuple of three elements. When the function is called, the elements are assigned to differ

Now let's go back and see why this auto-tuple-expansion trick was necessary. `os` was a list of tuples, and each tuple had three elements. When y

17.6. Stage 5

You've factored out all the duplicate code and added enough abstractions so that the pluralization rules are defined in a list of strings. The ne

First, let's create a text file that contains the rules you want. No fancy data structures, just space- (or tab-)delimited strings in three columns. T

Example 17.15. `pluralize.py`

```
def pluralize(word, count):  
    if count == 1:  
        return word  
    else:  
        return word + 's'
```

Now let's see how you can use this rules file.

Example 17.16. `pluralize.py`

```
def pluralize(word, count, rules):  
    if count == 1:  
        return word  
    else:  
        for rule in rules:  
            if rule.matches(word):  
                return rule.pluralize(word, count)  
        return word + 's'
```

- 1 You're still using the closures technique here (building a function dynamically that uses variables defined outside the function), but now you're building a list of functions.
- 2 Our `pluralize` function now takes an optional second parameter, `rules`, which defaults to `PLURAL_RULES`.
- 3 You use the `rules` parameter to construct a filename, then open the file and read the contents into a list. If `rules` is `None`, then you'll open the `rules.py` file, read the contents, and parse them into a list of `Rule` objects.
- 4 As you saw, each line in the file really has three values, but they're separated by whitespace (tabs or spaces, it makes no difference). `split` returns a list of strings, so you need to convert each string to a `Rule` object.
- 5 If `rules` is a list of tuples, then `pluralize` will be a list of the functions created dynamically by each call to `Rule`. Calling `pluralize` returns a function that takes a string and a count and returns a pluralized string.
- 6 Because you're now building a combined match-and-apply function, you need to call it differently. Just call the function, and if it returns `None`, then try the next rule.

So the improvement here is that you've completely separated the pluralization rules into an external file. Not only can the file be maintained separately, but you can also reuse it in other programs.

The downside here is that you're reading that file every time you call the `pluralize` function. I thought I could get through this entire book without using `cached_property`.

17.7. Stage 6

Now you're ready to talk about generators.

Example 17.17. `pluralize.py`

```
def pluralize(word, count, rules):  
    if count == 1:  
        return word  
    else:  
        for rule in rules:  
            if rule.matches(word):  
                return rule.pluralize(word, count)  
        return word + 's'
```



This uses a technique called generators, which I'm not even going to try to explain until you look at a simpler example first.

Example 17.18. Introducing generators



- ❶ The presence of the `yield` keyword in `fib` means that this is not a normal function. It is a special kind of function which generates values one at a time.
- ❷ To create an instance of the `fib` generator, just call it like any other function. Note that this does not actually execute the function code. You just get a generator object.
- ❸ The `fib` function returns a generator object.
- ❹ The first time you call the `next` method on the generator object, it executes the code in `fib` up to the first `yield` statement, and then returns the value of the expression after `yield`.
- ❺ Repeatedly calling `next` on the generator object *resumes where you left off* and continues until you hit the next `yield` statement. The next line of code after the `yield` statement is executed, and the value of the expression after `yield` is returned.
- ❻ The second time you call `next` you do all the same things again, but this time `x` is now `4` and so forth. Since `yield` sets up an infinite loop, you can call `next` as many times as you want.

Example 17.19. Using generators instead of recursion



- ❶ The Fibonacci sequence is a sequence of numbers where each number is the sum of the two numbers before it. It starts with `0` and `1`, goes on forever, and each number is the sum of the two numbers before it.
- ❷ `0` is the current number in the sequence, so yield it.
- ❸ `1` is the next number in the sequence, so assign that to `a` but also calculate the next value (`a + b`) and assign that to `b` for later use. Note that this is the same as the recursive version.

So you have a function that spits out successive Fibonacci numbers. Sure, you could do that with recursion, but this way is easier to read. Also, you can use the `yield` statement to generate a sequence of numbers.

Example 17.20. Generators in loops



- ❶ You can use a generator like `lines` in a `for` loop directly. The `for` loop will create the generator object and successively call the `next` method to get values.
- ❷ Each time through the `for` loop, `next` gets a new value from the `yield` statement in `lines` and all you do is print it out. Once `lines` runs out of numbers (`next` gets `StopIteration`),

OK, let's go back to the `lines` function and see how you're using this.

Example 17.21. Generators that generate dynamic functions



- ❶ `lines` is a common idiom for reading lines from a file, one line at a time. It works because `lines` actually returns a generator whose `next` method returns the next line.
- ❷ No magic here. Remember that the lines of the rules file have three values separated by whitespace, so `lines` returns a tuple of 3 values, and `next` returns a tuple of 3 values.
- ❸ And then you yield. What do you yield? A function, built dynamically with `eval` that is actually a closure (it uses the local variables `lines` and `next`).
- ❹ Since `lines` is a generator, you can use it directly in a `for` loop. The first time through the `for` loop, you will call the `next` function, which will open the file.

What have you gained over stage 5? In stage 5, you read the entire rules file and built a list of all the possible rules before you even tried the first one.

Further reading

- PEP 255 (<http://www.python.org/peps/pep-0255.html>) defines generators.
- Python Cookbook (<http://www.activestate.com/ASPN/Python/Cookbook/>) has many more examples of generators (<http://www.google.com/search?q=python+cookbook+generators>).

17.8. Summary

You talked about several different advanced techniques in this chapter. Not all of them are appropriate for every situation.

You should now be comfortable with all of these techniques:

- Performing string substitution with regular expressions.
- Treating functions as objects, storing them in lists, assigning them to variables, and calling them through those variables.
- Building dynamic functions with `eval`.
- Building closures, dynamic functions that contain surrounding variables as constants.
- Building generators, resumable functions that perform incremental logic and return different values each time you call them.

Adding abstractions, building functions dynamically, building closures, and using generators can all make your code simpler, more readable,

Chapter 18. Performance Tuning

Performance tuning is a many-splendored thing. Just because Python is an interpreted language doesn't mean you shouldn't worry about code

18.1. Diving in

There are so many pitfalls involved in optimizing your code, it's hard to know where to start.

Let's start here: *are you sure you need to do it at all?* Is your code really so bad? Is it worth the time to tune it? Over the lifetime of your app

Second, *are you sure you're done coding?* Premature optimization is like spreading frosting on a half-baked cake. You spend hours or days (

This is not to say that code optimization is worthless, but you need to look at the whole system and decide whether it's the best use of your tim

Oh yes, unit tests. It should go without saying that you need a complete set of unit tests before you begin performance tuning. The last thing y

With these caveats in place, let's look at some techniques for optimizing Python code. The code in question is an implementation of the Sound

There are several subtle variations of the Soundex algorithm. This is the one used in this chapter:

Keep the first letter of the name as-is.1.

Convert the remaining letters to digits, according to a specific table:

B, F, P, and V become 1.

C, G, J, K, Q, S, X, and Z become 2.

D and T become 3.

L becomes 4.

M and N become 5.

R becomes 6.

All other letters become 9.

Remove consecutive duplicates.3.

Remove all 9s altogether.4.

If the result is shorter than four characters (the first letter plus three digits), pad the result with trailing zeros.5.

if the result is longer than four characters, discard everything after the fourth character.6.

For example, my name, **B** becomes P942695. That has no consecutive duplicates, so nothing to do there. Then you remove the 9s, leaving P426

Another example: **W** becomes W99, which becomes W9, which becomes W, which gets padded with zeros to become W000.

Here's a first attempt at a Soundex function:

Example 18.1.

If you have not already done so, you can download this and other examples (<http://diveintopython.org/download/diveintopython-examples-5>).

```
def soundex(s):
```

```
    s = s.upper()
```

```
    s = s[0]
```

```
    s = s + s[1:]
```

```
    s = s[0]
```

```
    s = s + s[1:]
```

```
    s = s[0]
```

```
    s = s + s[1:]
```

Q
R
S
T
U
V
W
X
Y
Z

[
]

{
}

^
_

`
~

@

\$
%
^
_
`
~

<
>

=



Further Reading on Soundex

Soundexing and Genealogy (<http://www.avotaynu.com/soundex.html>) gives a chronology of the evolution of the Soundex and its re

18.2. Using the `timeit` Module

The most important thing you need to know about optimizing Python code is that you shouldn't write your own timing function.

Timing short pieces of code is incredibly complex. How much processor time is your computer devoting to running this code? Are there things

And then there's the matter of the variations introduced by the timing framework itself. Does the Python interpreter cache method name look

The Python community has a saying: "Python comes with batteries included." Don't write your own timing framework. Python 2.3 comes wi

Example 18.2. Introducing `timeit`

If you have not already done so, you can download this and other examples (<http://diveintopython.org/download/diveintopython-examples-5>).



- ❶ The `timeit` module defines one class, `Timeit`, which takes two arguments. Both arguments are strings. The first argument is the statement you wish
- ❷ Once you have the `Timeit` object, the easiest thing to do is call `timeit` which calls your function 1 million times and returns the number of seconds
- ❸ The other major method of the `Timeit` object is `repeat` which takes two optional arguments. The first argument is the number of times to repeat the

You can use the `timeit` module on the `command` line to test an existing Python program, without modifying the code. See <http://docs.python.org/lib>

Note that `timeit` returns a list of times. The times will almost never be identical, due to slight variations in how much processor time the Python inte

In fact, that's almost certainly wrong. The tests that took longer didn't take longer because of variations in your code or in the Python interpre

Python has a handy `min` function that takes a list and returns the smallest value:



The `isalpha` module only works if you already know what piece of code you need to optimize. If you have a larger Python program and don't know

18.3. Optimizing Regular Expressions

The first thing the `Soundex` function checks is whether the input is a non-empty string of letters. What's the best way to do this?

If you answered "regular expressions", go sit in the corner and contemplate your bad instincts. Regular expressions are almost never the right

This code fragment from `isalpha` checks whether the function argument `s` is a word made entirely of letters, with at least one letter (not the empty string)



How does `isalpha` perform? For convenience, the `isalpha` section of the script contains this code that calls the `isalpha` module, sets up a timing test with three differ



So how does `isalpha` perform with this regular expression?



As you might expect, the algorithm takes significantly longer when called with longer names. There will be a few things we can do to narrow

The other thing to keep in mind is that we are testing a representative sample of names. `isalpha` is a kind of trivial case, in that it gets shorted down to

So what about that regular expression? Well, it's inefficient. Since the expression is testing for ranges of characters (`[Aa]` uppercase, and `[a-z]` lowercase)



`isalpha` is slightly faster than `isalpha` but nothing to get terribly excited about:



We saw in Section 15.3, “Refactoring” that regular expressions can be compiled and reused for faster results. Since this regular expression ne

```
import re
s = 'The quick brown fox jumps over the lazy dog'
re.compile(r'fox').search(s)
```

Using a compiled regular expression in `is` significantly faster:

```
import re
s = 'The quick brown fox jumps over the lazy dog'
re.compile(r'fox').search(s)
```

But is this the wrong path? The logic here is simple: the input `s` needs to be non-empty, and it needs to be composed entirely of letters. Would

Here is `is`

```
def is_letters(s):
    return s and s.isalpha()
```

It turns out that this technique in `is` *not* faster than using a compiled regular expression (although it is faster than using a non-compiled regul

```
import re
s = 'The quick brown fox jumps over the lazy dog'
re.compile(r'fox').search(s)
```

Why isn't `is` faster? The answer lies in the interpreted nature of Python. The regular expression engine is written in C, and compiled to run nati

It turns out that Python offers an obscure string method. You can be excused for not knowing about it, since it's never been mentioned in this

This is `is`

```
def is_letters(s):
    return s and s.isalpha()
```

How much did we gain by using this specific method in `is`? Quite a bit.

```
import re
s = 'The quick brown fox jumps over the lazy dog'
re.compile(r'fox').search(s)
```




18.4. Optimizing Dictionary Lookups

The second step of the Soundex algorithm is to convert characters to digits in a specific pattern. What's the best way to do this?

The most obvious solution is to define a dictionary with individual characters as keys and their corresponding digits as values, and do dictionary lookups for each character.



```
B
```

```
R
```

```
D
```

```
S
```

```
F
```

```
G
```

```
C
```

```
K
```

```
J
```

```
H
```

```
M
```

```
N
```

```
O
```

```
P
```

```
Q
```

```
A
```

```
S
```

```
T
```

```
U
```

```
V
```

```
W
```

```
X
```

```
Y
```

```
Z
```

```
0
```

```
1
```

```
2
```

```
3
```

```
4
```

```
5
```

```
6
```

```
7
```

```
8
```

```
9
```

```
10
```

```
11
```

```
12
```

```
13
```

```
14
```

```
15
```

```
16
```

```
17
```

```
18
```

```
19
```

```
20
```

```
21
```

```
22
```

```
23
```

You timed it already; this is how it performs:

```
1
```

```
2
```

```
3
```

```
4
```

```
5
```

```
6
```

```
7
```

```
8
```

```
9
```

```
10
```

```
11
```

```
12
```

This code is straightforward, but is it the best solution? Calling `dict.get()` on each individual character seems inefficient; it would probably be better to use a `dict.get()` method that takes a default value.

Then there's the matter of incrementally building the string. Incrementally building strings like this is horribly inefficient; internally, the Python is good at lists, though. It can treat a string as a list of characters automatically. And lists are easy to combine into strings again, using

Here is which converts letters to digits by using and

```
#  
#  
#
```

Surprisingly, is not faster:

```
#  
#  
#  
#
```

The overhead of the anonymous function kills any performance you gain by dealing with the string as a list of characters.

uses a list comprehension instead of and

```
#  
#
```

Using a list comprehension in is faster than using and in but still not faster than the original code (incrementally building a string in

```
#  
#  
#  
#
```

It's time for a radically different approach. Dictionary lookups are a general purpose tool. Dictionary keys can be any length string (or many

This is

```
#  
#  
#  
#  
#
```

What the heck is going on here? creates a translation matrix between two strings: the first argument and the second argument. In this case, the

shows that is significantly faster than defining a dictionary and looping through the input and building the output incrementally:

```
#  
#  
#  
#
```

You're not going to get much better than that. Python has a specialized function that does exactly what you want to do; use it and move on.

Example 18.4. Best Result So Far:

```
0
```

```
1
```

```
2
```

```
3
```

```
4
```

```
5
```

```
6
```

```
7
```

```
8
```

```
9
```

```
0
```

```
1
```

```
2
```

```
3
```

```
4
```

```
5
```

```
6
```

```
7
```

```
8
```

```
9
```

```
0
```

```
1
```

```
2
```

18.5. Optimizing List Operations

The third step in the Soundex algorithm is eliminating consecutive duplicate digits. What's the best way to do this?

Here's the code we have so far, in

```
0
```

```
1
```

```
2
```

```
3
```

Here are the performance results for

```
0
```

```
1
```

```
2
```

```
3
```

The first thing to consider is whether it's efficient to check each time through the loop. Are list indexes expensive? Would we be better off n

To answer this question, here is

```
def is_palindrome(s):  
    return s == s[::-1]
```

Does not run any faster than and may even be slightly slower (although it's not enough of a difference to say for sure):

```
def is_palindrome(s):  
    return s == s[::-1]
```

Why isn't faster? It turns out that list indexes in Python are extremely efficient. Repeatedly accessing is no problem at all. On the other hand,

Let's try something radically different. If it's possible to treat a string as a list of characters, it should be possible to use a list comprehension

However, it is possible to create a list of index numbers using the built-in function, and use those index numbers to progressively search through

Here is

```
def is_palindrome(s):  
    return s == s[::-1]
```

Is this faster? In a word, no.

```
def is_palindrome(s):  
    return s == s[::-1]
```

It's possible that the techniques so far as have been "string-centric". Python can convert a string into a list of characters with a single command

Here is which modifies a list in place to remove consecutive duplicate elements:

```
def dedupe(s):  
    seen = set()  
    result = []  
    for char in s:  
        if char not in seen:  
            result.append(char)  
            seen.add(char)
```

Is this faster than or No, in fact it's the slowest method yet:

```
def dedupe(s):  
    seen = set()  
    result = []  
    for char in s:  
        if char not in seen:  
            result.append(char)  
            seen.add(char)
```



We haven't made any progress here at all, except to try and rule out several "clever" techniques. The fastest code we've seen so far was the o

Example 18.5. Best Result So Far:



18.6. Optimizing String Manipulation

The final step of the Soundex algorithm is padding short results with zeros, and truncating long results. What is the best way to do this?

This is what we have so far, taken from 



These are the results for 





The first thing to consider is replacing that regular expression with a loop. This code is from



Is it faster? Yes it is:



But wait a minute. A loop to remove characters from a string? We can use a simple string method for that. Here's



Is it faster? That's an interesting question. It depends on the input:



The string method in is faster than the loop for most names, but it's actually slightly slower than in the trivial case (of a very short name). Pe

Last but not least, let's examine the final two steps of the algorithm: padding short results with zeros, and truncating long results to four chara



Why do we need a loop to pad out the result? We know in advance that we're going to truncate the result to four characters, and we know th

How much speed do we gain in by dropping the loop? It's significant:



Finally, there is still one more thing you can do to these three lines of code to make them faster: you can combine them into one line. Take a



Putting all this code on one line in `is` barely faster than `is`



It is also significantly less readable, and for not much performance gain. Is that worth it? I hope you have good comments. Performance isn't

18.7. Summary

This chapter has illustrated several important aspects of performance tuning in Python, and performance tuning in general.

If you need to choose between regular expressions and writing a loop, choose regular expressions. The regular expression engine is

If you need to choose between regular expressions and string methods, choose string methods. Both are compiled in C, so choose the

General-purpose dictionary lookups are fast, but specialty functions such as `is` and string methods such as `is` are faster. If Python has a

Don't be too clever. Sometimes the most obvious algorithm is also the fastest.

Don't sweat it too much. Performance isn't everything.

I can't emphasize that last point strongly enough. Over the course of this chapter, you made this function three times faster and saved 20 seconds

Appendix A. Further reading

Chapter 1. Installing Python

Chapter 2. Your First Python Program

2.3. Documenting Functions

PEP 257 (<http://www.python.org/peps/pep-0257.html>) defines conventions.

Python Style Guide (<http://www.python.org/doc/essays/styleguide.html>) discusses how to write a good

Python Tutorial (<http://www.python.org/doc/current/tut/tut.html>) discusses conventions for spacing in (<http://www.python.org/doc/current/tut/tut.html>)

2.4.2. What's an Object?

Python Reference Manual (<http://www.python.org/doc/current/ref/>) explains exactly what it means to say that everything is an object.

eff-bot (<http://www.effbot.org/guides/>) summarizes Python objects (<http://www.effbot.org/guides/python-objects.htm>).

2.5. Indenting Code

Python Reference Manual (<http://www.python.org/doc/current/ref/>) discusses cross-platform indentation issues and shows

Python Style Guide (<http://www.python.org/doc/essays/styleguide.html>) discusses good indentation style.

2.6. Testing Modules

Python Reference Manual (<http://www.python.org/doc/current/ref/>) discusses the low-level details of importing modules (<http://www.python.org/doc/current/ref/>)

Chapter 3. Native Datatypes

3.1.3. Deleting Items From Dictionaries

How to Think Like a Computer Scientist (<http://www.ibiblio.org/obp/thinkCSPy/>) teaches about dictionaries and shows how to delete items.

Python Knowledge Base (<http://www.faqs.com/knowledge-base/index.phtml/fid/199/>) has a lot of example code using dictionaries.

Python Cookbook (<http://www.activestate.com/ASPN/Python/Cookbook/>) discusses how to sort the values of a dictionary.

Python Library Reference (<http://www.python.org/doc/current/lib/>) summarizes all the dictionary methods (<http://www.python.org/doc/current/lib/>)

3.2.5. Using List Operators

How to Think Like a Computer Scientist (<http://www.ibiblio.org/obp/thinkCSPy/>) teaches about lists and makes an important distinction between lists and tuples.

Python Tutorial (<http://www.python.org/doc/current/tut/tut.html>) shows how to use lists as stacks and queues (<http://www.python.org/doc/current/tut/tut.html>)

Python Knowledge Base (<http://www.faqs.com/knowledge-base/index.phtml/fid/199/>) answers common questions about lists.

Python Library Reference (<http://www.python.org/doc/current/lib/>) summarizes all the list methods (<http://www.python.org/doc/current/lib/>)

3.3. Introducing Tuples

How to Think Like a Computer Scientist (<http://www.ibiblio.org/obp/thinkCSPy/>) teaches about tuples and shows how to compare them.

Python Knowledge Base (<http://www.faqs.com/knowledge-base/index.phtml/fid/199/>) shows how to sort a tuple (<http://www.faqs.com/knowledge-base/index.phtml/fid/199/>)

Python Tutorial (<http://www.python.org/doc/current/tut/tut.html>) shows how to define a tuple with one element (<http://www.python.org/doc/current/tut/tut.html>)

3.4.2. Assigning Multiple Values at Once

Python Reference Manual (<http://www.python.org/doc/current/ref/>) shows examples of when you can skip the line continuation character.

How to Think Like a Computer Scientist (<http://www.ibiblio.org/obp/thinkCSPy/>) shows how to use multi-variable assignment.

3.5. Formatting Strings

Python Library Reference (<http://www.python.org/doc/current/lib/>) summarizes all the string formatting format characters.

Effective AWK Programming ([http://www-gnats.gnu.org:8080/cgi-bin/info2www?\(gawk\)Top](http://www-gnats.gnu.org:8080/cgi-bin/info2www?(gawk)Top)) discusses all the format characters.

3.6. Mapping Lists

Python Tutorial (<http://www.python.org/doc/current/tut/tut.html>) discusses another way to map lists using the built-in `map` function.

Python Tutorial (<http://www.python.org/doc/current/tut/tut.html>) shows how to do nested list comprehensions (<http://www.python.org/doc/current/tut/tut.html>)

3.7. Joining Lists and Splitting Strings

Python Knowledge Base (<http://www.faqs.com/knowledge-base/index.phtml/fid/199/>) answers common questions about strings.

Python Library Reference (<http://www.python.org/doc/current/lib/>) summarizes all the string methods (<http://www.python.org/doc/current/lib/>)

Python Library Reference (<http://www.python.org/doc/current/lib/>) documents the `unicodedata` module (<http://www.python.org/doc/current/lib/>)

The Whole Python FAQ (<http://www.python.org/doc/FAQ.html>) explains why `is` is a string method (<http://www.python.org/doc/FAQ.html>)

Chapter 4. The Power Of Introspection

4.2. Using Optional and Named Arguments

- Python Tutorial* (<http://www.python.org/doc/current/tut/tut.html>) discusses exactly when and how default arguments are evaluated.
- 4.3.3. Built-In Functions
 - Python Library Reference* (<http://www.python.org/doc/current/lib/>) documents all the built-in functions (<http://www.python.org/doc/current/lib/#built-in-functions>).
- 4.5. Filtering Lists
 - Python Tutorial* (<http://www.python.org/doc/current/tut/tut.html>) discusses another way to filter lists using the built-in `filter` function (<http://www.python.org/doc/current/tut/tut.html#filter>).
- 4.6.1. Using the and-or Trick
 - Python Cookbook* (<http://www.activestate.com/ASPN/Python/Cookbook/>) discusses alternatives to the `and-or` trick (<http://www.activestate.com/ASPN/Python/Cookbook/recipe/1182.html>).
- 4.7.1. Real-World lambda Functions
 - Python Knowledge Base* (<http://www.faqs.com/knowledge-base/index.phtml/fid/199/>) discusses using `lambda` to call functions in `__main__` (<http://www.faqs.com/knowledge-base/index.phtml/fid/199/#lambda>).
 - Python Tutorial* (<http://www.python.org/doc/current/tut/tut.html>) shows how to access outside variables from inside a `lambda` function (<http://www.python.org/doc/current/tut/tut.html#lambda>).
 - The Whole Python FAQ* (<http://www.python.org/doc/FAQ.html>) has examples of obfuscated one-liners using `lambda` (<http://www.python.org/doc/FAQ.html#lambda>).

Chapter 5. Objects and Object-Orientation

- 5.2. Importing Modules Using `from module import`
 - eff-bot* (<http://www.effbot.org/guides/>) has more to say on `from` imports (<http://www.effbot.org/guides/import-confusion.htm>).
 - Python Tutorial* (<http://www.python.org/doc/current/tut/tut.html>) discusses advanced import techniques, including `from` imports (<http://www.python.org/doc/current/tut/tut.html#from>).
- 5.3.2. Knowing When to Use `self` and `__init__`
 - Learning to Program* (<http://www.freenetpages.co.uk/hp/alan.gauld/>) has a gentler introduction to classes (<http://www.freenetpages.co.uk/hp/alan.gauld/learn-to-program/10-classes.html>).
 - How to Think Like a Computer Scientist* (<http://www.ibiblio.org/obp/thinkCSPy/>) shows how to use classes to model computation (<http://www.ibiblio.org/obp/thinkCSPy/10-classes.html>).
 - Python Tutorial* (<http://www.python.org/doc/current/tut/tut.html>) has an in-depth look at classes, namespaces, and inheritance (<http://www.python.org/doc/current/tut/tut.html#classes>).
 - Python Knowledge Base* (<http://www.faqs.com/knowledge-base/index.phtml/fid/199/>) answers common questions about classes (<http://www.faqs.com/knowledge-base/index.phtml/fid/199/#classes>).
- 5.4.1. Garbage Collection
 - Python Library Reference* (<http://www.python.org/doc/current/lib/>) summarizes built-in attributes like `__del__` (<http://www.python.org/doc/current/lib/#del>).
 - Python Library Reference* (<http://www.python.org/doc/current/lib/>) documents the `gc` module (<http://www.python.org/doc/current/lib/#gc-module>).
- 5.5. Exploring `UserDict`: A Wrapper Class
 - Python Library Reference* (<http://www.python.org/doc/current/lib/>) documents the `UserDict` module (<http://www.python.org/doc/current/lib/#userdict-module>).
- 5.7. Advanced Special Class Methods
 - Python Reference Manual* (<http://www.python.org/doc/current/ref/>) documents all the special class methods (<http://www.python.org/doc/current/ref/#special-methods>).
- 5.9. Private Functions
 - Python Tutorial* (<http://www.python.org/doc/current/tut/tut.html>) discusses the inner workings of private variables (<http://www.python.org/doc/current/tut/tut.html#private>).

Chapter 6. Exceptions and File Handling

- 6.1.1. Using Exceptions For Other Purposes
 - Python Tutorial* (<http://www.python.org/doc/current/tut/tut.html>) discusses defining and raising your own exceptions, and `try` blocks (<http://www.python.org/doc/current/tut/tut.html#try>).
 - Python Library Reference* (<http://www.python.org/doc/current/lib/>) summarizes all the built-in exceptions (<http://www.python.org/doc/current/lib/#exceptions>).
 - Python Library Reference* (<http://www.python.org/doc/current/lib/>) documents the `getpass` module (<http://www.python.org/doc/current/lib/#getpass-module>).
 - Python Library Reference* (<http://www.python.org/doc/current/lib/>) documents the `random` module (<http://www.python.org/doc/current/lib/#random-module>).
 - Python Reference Manual* (<http://www.python.org/doc/current/ref/>) discusses the inner workings of the `try` block (<http://www.python.org/doc/current/ref/#try>).
- 6.2.4. Writing to Files
 - Python Tutorial* (<http://www.python.org/doc/current/tut/tut.html>) discusses reading and writing files, including how to read and write files in binary mode (<http://www.python.org/doc/current/tut/tut.html#files>).
 - eff-bot* (<http://www.effbot.org/guides/>) discusses efficiency and performance of various ways of reading a file (<http://www.effbot.org/guides/reading-files.html>).
 - Python Knowledge Base* (<http://www.faqs.com/knowledge-base/index.phtml/fid/199/>) answers common questions about file objects (<http://www.faqs.com/knowledge-base/index.phtml/fid/199/#file-objects>).
 - Python Library Reference* (<http://www.python.org/doc/current/lib/>) summarizes all the file object methods (<http://www.python.org/doc/current/lib/#file-object-methods>).
- 6.4. Using `sys` modules
 - Python Tutorial* (<http://www.python.org/doc/current/tut/tut.html>) discusses exactly when and how default arguments are evaluated (<http://www.python.org/doc/current/tut/tut.html#sys>).
 - Python Library Reference* (<http://www.python.org/doc/current/lib/>) documents the `sys` module (<http://www.python.org/doc/current/lib/#sys-module>).
- 6.5. Working with Directories
 - Python Knowledge Base* (<http://www.faqs.com/knowledge-base/index.phtml/fid/199/>) answers questions about the `os` module (<http://www.faqs.com/knowledge-base/index.phtml/fid/199/#os>).
 - Python Library Reference* (<http://www.python.org/doc/current/lib/>) documents the `os` module (<http://www.python.org/doc/current/lib/#os-module>).

Chapter 7. Regular Expressions

7.6. Case study: Parsing Phone Numbers

Regular Expression HOWTO (<http://py-howto.sourceforge.net/regex/regex.html>) teaches about regular expressions and how to use them.
Python Library Reference (<http://www.python.org/doc/current/lib/>) summarizes the `re` module (<http://www.python.org/doc/current/lib/module-re.html>).

Chapter 8. HTML Processing

8.4. Introducing BaseHTMLProcessor.py

W3C (<http://www.w3.org/>) discusses character and entity references (<http://www.w3.org/TR/REC-html40/charset.html#ent>).
Python Library Reference (<http://www.python.org/doc/current/lib/>) confirms your suspicions that the `html` module (<http://www.python.org/doc/current/lib/module-html.html>).

8.9. Putting it all together

You thought I was kidding about the server-side scripting idea. So did I, until I found this web-based dialectizer (<http://rink.com/>).

Chapter 9. XML Processing

9.4. Unicode

Unicode.org (<http://www.unicode.org/>) is the home page of the unicode standard, including a brief technical introduction (Unicode Tutorial (http://www.reportlab.com/i18n/python_unicode_tutorial.html) has some more examples of how to use Python's Unicode support.
PEP 263 (<http://www.python.org/peps/pep-0263.html>) goes into more detail about how and when to define a character encoding.

Chapter 10. Scripts and Streams

Chapter 11. HTTP Web Services

11.1. Diving in

Paul Prescod believes that pure HTTP web services are the future of the Internet (<http://webservices.xml.com/pub/a/ws/2000/>).

Chapter 12. SOAP Web Services

12.1. Diving In

<http://www.xmethods.net/> is a repository of public access SOAP web services.

The SOAP specification (<http://www.w3.org/TR/soap/>) is surprisingly readable, if you like that sort of thing.

12.8. Troubleshooting SOAP Web Services

New developments for SOAPpy (<http://www-106.ibm.com/developerworks/webservices/library/ws-pyth17.html>) steps through some common problems.

Chapter 13. Unit Testing

13.1. Introduction to Roman numerals

This site (<http://www.wilkiecollins.demon.co.uk/roman/front.htm>) has more on Roman numerals, including a fascinating history.

13.3. Introducing romantest.py

The PyUnit home page (<http://pyunit.sourceforge.net/>) has an in-depth discussion of using the `unittest` framework (<http://pyunit.sourceforge.net/>).

The PyUnit FAQ (<http://pyunit.sourceforge.net/pyunit.html>) explains why test cases are stored separately (<http://pyunit.sourceforge.net/>).

Python Library Reference (<http://www.python.org/doc/current/lib/>) summarizes the `unittest` module (<http://www.python.org/doc/current/lib/module-unittest.html>).

ExtremeProgramming.org (<http://www.extremeprogramming.org/>) discusses why you should write unit tests (<http://www.extremeprogramming.org/>).

The Portland Pattern Repository (<http://www.c2.com/cgi/wiki/>) has an ongoing discussion of unit tests (<http://www.c2.com/cgi/wiki/UnitTesting>).

Chapter 14. Test-First Programming

Chapter 15. Refactoring

15.5. Summary

XProgramming.com (<http://www.xprogramming.com/>) has links to download unit testing frameworks (<http://www.xprogramming.com/>).

Chapter 16. Functional Programming

Chapter 17. Dynamic functions

17.7. plural.py, stage 6

PEP 255 (<http://www.python.org/peps/pep-0255.html>) defines generators.

Python Cookbook (<http://www.activestate.com/ASPN/Python/Cookbook/>) has many more examples of generators (<http://w>

Chapter 18. Performance Tuning

18.1. Diving in

Soundexing and Genealogy (<http://www.avotaynu.com/soundex.html>) gives a chronology of the evolution of the Soundex

Appendix B. A 5-minute review

Chapter 1. Installing Python

1.1. Which Python is right for you?

The first thing you need to do with Python is install it. Or do you?

1.2. Python on Windows

On Windows, you have a couple choices for installing Python.

1.3. Python on Mac OS X

On Mac OS X, you have two choices for installing Python: install it, or don't install it. You probably want to install it.

1.4. Python on Mac OS 9

Mac OS 9 does not come with any version of Python, but installation is very simple, and there is only one choice.

1.5. Python on RedHat Linux

Download the latest Python RPM by going to <http://www.python.org/ftp/python/> and selecting the highest version number

1.6. Python on Debian GNU/Linux

If you are lucky enough to be running Debian GNU/Linux, you install Python through the **apt** command.

1.7. Python Installation from Source

If you prefer to build from source, you can download the Python source code from <http://www.python.org/ftp/python/>. Sele

1.8. The Interactive Shell

Now that you have Python installed, what's this interactive shell thing you're running?

1.9. Summary

You should now have a version of Python installed that works for you.

Chapter 2. Your First Python Program

2.1. Diving in

Here is a complete, working Python program.

2.2. Declaring Functions

Python has functions like most other languages, but it does not have separate header files like C++ or sections like Pascal.

2.3. Documenting Functions

You can document a Python function by giving it a 

2.4. Everything Is an Object

A function, like everything else in Python, is an object.

2.5. Indenting Code

Python functions have no explicit `begin` and no curly braces to mark where the function code starts and stops. The only delimit

2.6. Testing Modules

Python modules are objects and have several useful attributes. You can use this to easily test your modules as you write the

Chapter 3. Native Datatypes

3.1. Introducing Dictionaries

One of Python's built-in datatypes is the dictionary, which defines one-to-one relationships between keys and values.

3.2. Introducing Lists

Lists are Python's workhorse datatype. If your only experience with lists is arrays in Visual Basic or (God forbid) the data

3.3. Introducing Tuples

A tuple is an immutable list. A tuple can not be changed in any way once it is created.

3.4. Declaring variables

Python has local and global variables like most other languages, but it has no explicit variable declarations. Variables sprin

3.5. Formatting Strings

Python supports formatting values into strings. Although this can include very complicated expressions, the most basic usa

3.6. Mapping Lists

One of the most powerful features of Python is the list comprehension, which provides a compact way of mapping a list into

3.7. Joining Lists and Splitting Strings

You have a list of key-value pairs in the form `key:value` and you want to join them into a single string. To join any list of strings into

3.8. Summary

The `program` and its output should now make perfect sense.

Chapter 4. The Power Of Introspection

4.1. Diving In

Here is a complete, working Python program. You should understand a good deal about it just by looking at it. The number

4.2. Using Optional and Named Arguments

Python allows function arguments to have default values; if the function is called without the argument, the argument gets

4.3. Using type, str, dir, and Other Built-In Functions

Python has a small set of extremely useful built-in functions. All other functions are partitioned off into modules. This was

4.4. Getting Object References With getattr

You already know that Python functions are objects. What you don't know is that you can get a reference to a function with

4.5. Filtering Lists

As you know, Python has powerful capabilities for mapping lists into other lists, via list comprehensions (Section 3.6, "Ma

4.6. The Peculiar Nature of and and or

In Python, `and` and `or` perform boolean logic as you would expect, but they do not return boolean values; instead, they return one

4.7. Using lambda Functions

Python supports an interesting syntax that lets you define one-line mini-functions on the fly. Borrowed from Lisp, these so-

4.8. Putting It All Together

The last line of code, the only one you haven't deconstructed yet, is the one that does all the work. But by now the work is

4.9. Summary

The program and its output should now make perfect sense.

Chapter 5. Objects and Object-Orientation

5.1. Diving In

Here is a complete, working Python program. Read the `__doc__` of the module, the classes, and the functions to get an overview of

5.2. Importing Modules Using from module import

Python has two ways of importing modules. Both are useful, and you should know when to use each. One way, `from module import`, you've alr

5.3. Defining Classes

Python is fully object-oriented: you can define your own classes, inherit from your own or built-in classes, and instantiate t

5.4. Instantiating Classes

Instantiating classes in Python is straightforward. To instantiate a class, simply call the class as if it were a function, passin

5.5. Exploring UserDict: A Wrapper Class

As you've seen, `UserDict` is a class that acts like a dictionary. To explore this further, let's look at the `dict` class in the `collections` module, which is

5.6. Special Class Methods

In addition to normal class methods, there are a number of special methods that Python classes can define. Instead of being

5.7. Advanced Special Class Methods

Python has more special methods than just `__and__` and `__is__`. Some of them let you emulate functionality that you may not even know a

5.8. Introducing Class Attributes

You already know about data attributes, which are variables owned by a specific instance of a class. Python also supports c

5.9. Private Functions

Unlike in most languages, whether a Python function, method, or attribute is private or public is determined entirely by its

5.10. Summary

That's it for the hard-core object trickery. You'll see a real-world application of special class methods in Chapter 12, which

Chapter 6. Exceptions and File Handling

6.1. Handling Exceptions

Like many other programming languages, Python has exception handling via `try` blocks.

6.2. Working with File Objects

Python has a built-in function, `open`, for opening a file on disk. `open` returns a file object, which has methods and attributes for getting

6.3. Iterating with `for` Loops

Like most other languages, Python has `for` loops. The only reason you haven't seen them until now is that Python is good at so

6.4. Using `sys` modules

Modules, like everything else in Python, are objects. Once imported, you can always get a reference to a module through th

6.5. Working with Directories

The `os` module has several functions for manipulating files and directories. Here, we're looking at handling pathnames and lis

6.6. Putting It All Together

Once again, all the dominoes are in place. You've seen how each line of code works. Now let's step back and see how it al

6.7. Summary

The `try` program introduced in Chapter 5 should now make perfect sense.

Chapter 7. Regular Expressions

7.1. Diving In

If what you're trying to do can be accomplished with string functions, you should use them. They're fast and simple and ea

7.2. Case Study: Street Addresses

This series of examples was inspired by a real-life problem I had in my day job several years ago, when I needed to scrub a

7.3. Case Study: Roman Numerals

You've most likely seen Roman numerals, even if you didn't recognize them. You may have seen them in copyrights of old

7.4. Using the {n,m} Syntax

In the previous section, you were dealing with a pattern where the same character could be repeated up to three times. There

7.5. Verbose Regular Expressions

So far you've just been dealing with what I'll call "compact" regular expressions. As you've seen, they are difficult to read

7.6. Case study: Parsing Phone Numbers

So far you've concentrated on matching whole patterns. Either the pattern matches, or it doesn't. But regular expressions are

7.7. Summary

This is just the tiniest tip of the iceberg of what regular expressions can do. In other words, even though you're completely

Chapter 8. HTML Processing

8.1. Diving in

I often see questions on comp.lang.python (<http://groups.google.com/groups?group=comp.lang.python>) like "How can I list

8.2. Introducing sgmlib.py

HTML processing is broken into three steps: breaking down the HTML into its constituent pieces, fiddling with the pieces,

8.3. Extracting data from HTML documents

To extract data from HTML documents, subclass the `HTMLParser` class and define methods for each tag or entity you want to capture.

8.4. Introducing BaseHTMLProcessor.py

`BaseHTMLProcessor` doesn't produce anything by itself. It parses and parses and parses, and it calls a method for each interesting thing it finds,

8.5. locals and globals

Let's digress from HTML processing for a minute and talk about how Python handles variables. Python has two built-in functions

8.6. Dictionary-based string formatting

There is an alternative form of string formatting that uses dictionaries instead of tuples of values.

8.7. Quoting attribute values

This is generally a good idea. It's a simple (and silly) descendant of `HTMLParser`. It runs blocks of text through a series of substitutions, but it makes sure that anything

8.8. Introducing dialect.py

It's a simple (and silly) descendant of `HTMLParser`. It runs blocks of text through a series of substitutions, but it makes sure that anything

8.9. Putting it all together

It's time to put everything you've learned so far to good use. I hope you were paying attention.

8.10. Summary

Python provides you with a powerful tool, `lxml`, to manipulate HTML by turning its structure into an object model. You can use

Chapter 9. XML Processing

9.1. Diving in

There are two basic ways to work with XML. One is called SAX ("Simple API for XML"), and it works by reading the XML

9.2. Packages

Actually parsing an XML document is very simple: one line of code. However, before you get to that line of code, you need

9.3. Parsing XML

As I was saying, actually parsing an XML document is very simple: one line of code. Where you go from there is up to you

9.4. Unicode

Unicode is a system to represent characters from all the world's different languages. When Python parses an XML document

9.5. Searching for elements

Traversing XML documents by stepping through each node can be tedious. If you're looking for something in particular, be

9.6. Accessing element attributes

XML elements can have one or more attributes, and it is incredibly simple to access them once you have parsed an XML document

9.7. Segue

OK, that's it for the hard-core XML stuff. The next chapter will continue to use these same example programs, but focus on

Chapter 10. Scripts and Streams

10.1. Abstracting input sources

One of Python's greatest strengths is its dynamic binding, and one powerful use of dynamic binding is the *file-like object*.

10.2. Standard input, output, and error

UNIX users are already familiar with the concept of standard input, standard output, and standard error. This section is for

10.3. Caching node lookups

`lxml` employs several tricks which may or may not be useful to you in your XML processing. The first one takes advantage of the

10.4. Finding direct children of a node

Another useful technique when parsing XML documents is finding all the direct child elements of a particular element. For i

10.5. Creating separate handlers by node type

The third useful XML processing tip involves separating your code into logical functions, based on node types and element

10.6. Handling command-line arguments

Python fully supports creating programs that can be run on the command line, complete with command-line arguments and

10.7. Putting it all together

You've covered a lot of ground. Let's step back and see how all the pieces fit together.

10.8. Summary

Python comes with powerful libraries for parsing and manipulating XML documents. The `xml.etree.ElementTree` module takes an XML file and parses it

Chapter 11. HTTP Web Services

11.1. Diving in

You've learned about HTML processing and XML processing, and along the way you saw how to download a web page and

11.2. How not to fetch data over HTTP

Let's say you want to download a resource over HTTP, such as a syndicated Atom feed. But you don't just want to download

11.3. Features of HTTP

There are five important features of HTTP which you should support.

11.4. Debugging HTTP web services

First, let's turn on the debugging features of Python's HTTP library and see what's being sent over the wire. This will be useful

11.5. Setting the User-Agent

The first step to improving your HTTP web services client is to identify yourself properly with a `User-Agent` header. To do that, you need to modify

11.6. Handling Last-Modified and ETag

Now that you know how to add custom HTTP headers to your web service requests, let's look at adding support for `Last-Modified` and `ETag` headers.

11.7. Handling redirects

You can support permanent and temporary redirects using a different kind of custom URL handler.

11.8. Handling compressed data

The last important HTTP feature you want to support is compression. Many web services have the ability to send data compressed

11.9. Putting it all together

You've seen all the pieces for building an intelligent HTTP web services client. Now let's see how they all fit together.

11.10. Summary

The `find` and its functions should now make perfect sense.

Chapter 12. SOAP Web Services

12.1. Diving In

You use Google, right? It's a popular search engine. Have you ever wished you could programmatically access Google search?

12.2. Installing the SOAP Libraries

Unlike the other code in this book, this chapter relies on libraries that do not come pre-installed with Python.

12.3. First Steps with SOAP

The heart of SOAP is the ability to call remote functions. There are a number of public access SOAP servers that provide search services.

12.4. Debugging SOAP Web Services

The SOAP libraries provide an easy way to see what's going on behind the scenes.

12.5. Introducing WSDL

The `SOAPProxy` class proxies local method calls and transparently turns them into invocations of remote SOAP methods. As you've seen, this is done by using the `__getattr__` method.

12.6. Introspecting SOAP Web Services with WSDL

Like many things in the web services arena, WSDL has a long and checkered history, full of political strife and intrigue. I won't go into the details here.

12.7. Searching Google

Let's finally turn to the sample code that you saw at the beginning of this chapter, which does something more useful and interesting than just searching Google.

12.8. Troubleshooting SOAP Web Services

Of course, the world of SOAP web services is not all happiness and light. Sometimes things go wrong.

12.9. Summary

SOAP web services are very complicated. The specification is very ambitious and tries to cover many different use cases for web services.

Chapter 13. Unit Testing

13.1. Introduction to Roman numerals

In previous chapters, you "dived in" by immediately looking at code and trying to understand it as quickly as possible. Now we'll take a different approach.

13.2. Diving in

Now that you've completely defined the behavior you expect from your conversion functions, you're going to do something interesting with them.

13.3. Introducing `romantest.py`

This is the complete test suite for your Roman numeral conversion functions, which are yet to be written but will eventually

13.4. Testing for success

The most fundamental part of unit testing is constructing individual test cases. A test case answers a single question about t

13.5. Testing for failure

It is not enough to test that functions succeed when given good input; you must also test that they fail when given bad input

13.6. Testing for sanity

Often, you will find that a unit of code contains a set of reciprocal functions, usually in the form of conversion functions w

Chapter 14. Test-First Programming

14.1. roman.py, stage 1

Now that the unit tests are complete, it's time to start writing the code that the test cases are attempting to test. You're going

14.2. roman.py, stage 2

Now that you have the framework of the module laid out, it's time to start writing code and passing test cases.

14.3. roman.py, stage 3

Now that behaves correctly with good input (integers from 1 to 9) it's time to make it behave correctly with bad input (every

14.4. roman.py, stage 4

Now that is done, it's time to start coding. Thanks to the rich data structure that maps individual Roman numerals to integ

14.5. roman.py, stage 5

Now that works properly with good input, it's time to fit in the last piece of the puzzle: making it work properly with bad i

Chapter 15. Refactoring

15.1. Handling bugs

Despite your best efforts to write comprehensive unit tests, bugs happen. What do I mean by "bug"? A bug is a test case yo

15.2. Handling changing requirements

Despite your best efforts to pin your customers to the ground and extract exact requirements from them on pain of horrible

15.3. Refactoring

The best thing about comprehensive unit testing is not the feeling you get when all your test cases finally pass, or even the

15.4. Postscript

A clever reader read the previous section and took it to the next level. The biggest headache (and performance drain) in the

15.5. Summary

Unit testing is a powerful concept which, if properly implemented, can both reduce maintenance costs and increase flexibility.

Chapter 16. Functional Programming

16.1. Diving in

In Chapter 13, *Unit Testing*, you learned about the philosophy of unit testing. In Chapter 14, *Test-First Programming*, you

16.2. Finding the path

When running Python scripts from the command line, it is sometimes useful to know where the currently running script is located.

16.3. Filtering lists revisited

You're already familiar with using list comprehensions to filter lists. There is another way to accomplish this same thing, with

16.4. Mapping lists revisited

You're already familiar with using list comprehensions to map one list into another. There is another way to accomplish this

16.5. Data-centric programming

By now you're probably scratching your head wondering why this is better than using `for` loops and straight function calls. And

16.6. Dynamically importing modules

OK, enough philosophizing. Let's talk about dynamically importing modules.

16.7. Putting it all together

You've learned enough now to deconstruct the first seven lines of this chapter's code sample: reading a directory and importing

16.8. Summary

The program and its output should now make perfect sense.

Chapter 17. Dynamic functions

17.1. Diving in

I want to talk about plural nouns. Also, functions that return other functions, advanced regular expressions, and generators.

17.2. plural.py, stage 1

So you're looking at words, which at least in English are strings of characters. And you have rules that say you need to find

17.3. plural.py, stage 2

Now you're going to add a level of abstraction. You started by defining a list of rules: if this, then do that, otherwise go to

17.4. plural.py, stage 3

Defining separate named functions for each match and apply rule isn't really necessary. You never call them directly; you

17.5. plural.py, stage 4

Let's factor out the duplication in the code so that defining new rules can be easier.

17.6. plural.py, stage 5

You've factored out all the duplicate code and added enough abstractions so that the pluralization rules are defined in a list

17.7. plural.py, stage 6

Now you're ready to talk about generators.

17.8. Summary

You talked about several different advanced techniques in this chapter. Not all of them are appropriate for every situation.

Chapter 18. Performance Tuning

18.1. Diving in

There are so many pitfalls involved in optimizing your code, it's hard to know where to start.

18.2. Using the timeit Module

The most important thing you need to know about optimizing Python code is that you shouldn't write your own timing function

18.3. Optimizing Regular Expressions

The first thing the Soundex function checks is whether the input is a non-empty string of letters. What's the best way to do

18.4. Optimizing Dictionary Lookups

The second step of the Soundex algorithm is to convert characters to digits in a specific pattern. What's the best way to do

18.5. Optimizing List Operations

The third step in the Soundex algorithm is eliminating consecutive duplicate digits. What's the best way to do this?

18.6. Optimizing String Manipulation

The final step of the Soundex algorithm is padding short results with zeros, and truncating long results. What is the best wa

18.7. Summary

This chapter has illustrated several important aspects of performance tuning in Python, and performance tuning in general.

Appendix C. Tips and tricks

Chapter 1. Installing Python

Chapter 2. Your First Python Program

2.1. Diving in

In the ActivePython IDE on Windows, you can run the Python program you're editing by choosing File->Run... (**Ctrl-R**). Output is

In the Python IDE on Mac OS, you can run a Python program with Python->Run window... (**Cmd-R**), but there is an important option

On UNIX-compatible systems (including Mac OS X), you can run a Python program from the command line:

2.2. Declaring Functions

In Visual Basic, functions (that return a value) start with **Function** and subroutines (that do not return a value) start with **Sub**. There are no subro

In Java, C++, and other statically-typed languages, you must specify the datatype of the function return value and each function argu

2.3. Documenting Functions

Triple quotes are also an easy way to define a string with both single and double quotes, like **"""** in Perl.

Many Python IDEs use the **help** function to provide context-sensitive documentation, so that when you type a function name, its **help** appears as a tool

2.4. Everything Is an Object

In Python is like in Perl. Once you **import** a Python module, you access its functions with **module.function**. Once you **use** a Perl module, you access its functi

2.5. Indenting Code

Python uses carriage returns to separate statements and a colon and indentation to separate code blocks. C++ and Java use semicol

2.6. Testing Modules

Like C, Python uses **==** for comparison and **=** for assignment. Unlike C, Python does not support in-line assignment, so there's no chan

On MacPython, there is an additional step to make the **__import__** trick work. Pop up the module's options menu by clicking the black triangle

Chapter 3. Native Datatypes

3.1. Introducing Dictionaries

A dictionary in Python is like a hash in Perl. In Perl, variables that store hashes always start with a **%** character. In Python, variables c

A dictionary in Python is like an instance of the `dict` class in Java.

A dictionary in Python is like an instance of the `Object` in Visual Basic.

3.1.2. Modifying Dictionaries

Dictionaries have no concept of order among elements. It is incorrect to say that the elements are "out of order"; they are simply unordered.

3.2. Introducing Lists

A list in Python is like an array in Perl. In Perl, variables that store arrays always start with the `@` character; in Python, variables can be named anything.

A list in Python is much more than an array in Java (although it can be used as one if that's really all you want out of life). A better analogy is a `Vector` in Java.

3.2.3. Searching Lists

Before version 2.2.1, Python had no separate boolean datatype. To compensate for this, Python accepted almost anything in a boolean context:

- 0 is false; all other numbers are true.

- An empty string (`''`) is false, all other strings are true.

- An empty list (`[]`) is false; all other lists are true.

- An empty tuple (`()`) is false; all other tuples are true.

- An empty dictionary (`{}`) is false; all other dictionaries are true.

These rules still apply in Python 2.2.1 and beyond, but now you can also use an actual boolean, which has a value of `True` or `False`. Note the capitalization.

3.3. Introducing Tuples

Tuples can be converted into lists, and vice-versa. The built-in `tuple()` function takes a list and returns a tuple with the same elements, and the built-in `list()` function takes a tuple and returns a list with the same elements.

3.4. Declaring variables

When a command is split among several lines with the line-continuation marker (`"\"`), the continued lines can be indented in any manner.

3.5. Formatting Strings

String formatting in Python uses the same syntax as the `printf` function in C.

3.7. Joining Lists and Splitting Strings

`join()` works only on lists of strings; `join()` does not do any type coercion. Joining a list that has one or more non-string elements will raise an error.

`split()` is a useful technique when you want to search a string for a substring and then work with everything before the substring (which is often the case).

Chapter 4. The Power Of Introspection

4.2. Using Optional and Named Arguments

The only thing you need to do to call a function is specify a value (somehow) for each required argument; the manner and order in which you specify the arguments is up to you.

4.3.3. Built-In Functions

Python comes with excellent reference manuals, which you should peruse thoroughly to learn all the modules Python has to offer. Be sure to read the `help()` function.

4.7. Using lambda Functions

Functions are a matter of style. Using them is never required; anywhere you could use them, you could define a separate normal function.

4.8. Putting It All Together

In SQL, you must use `IS NULL` to compare a null value. In Python, you can use either `is` or `is None`, but `is` is faster.

Chapter 5. Objects and Object-Oriented Programming

5.2. Importing Modules Using `from module import`

In Python is like in Perl; in Python is like in Perl.

In Python is like in Java; in Python is like in Java.

Use sparingly, because it makes it difficult to determine where a particular function or attribute came from, and that makes debugging harder.

5.3. Defining Classes

The `class` statement in Python is like an empty set of braces (`{` in Java or C).

In Python, the ancestor of a class is simply listed in parentheses immediately after the class name. There is no special keyword like `base`.

5.3.1. Initializing and Coding Classes

By convention, the first argument of any Python class method (the reference to the current instance) is called `self`. This argument fills the role of `this` in other languages.

5.3.2. Knowing When to Use `self` and `__init__`

Methods are optional, but when you define one, you must remember to explicitly call the ancestor's method (if it defines one). This is done by using `super()`.

5.4. Instantiating Classes

In Python, simply call a class as if it were a function to create a new instance of the class. There is no explicit operator like C++ or Java.

5.5. Exploring `UserDict`: A Wrapper Class

In the ActivePython IDE on Windows, you can quickly open any module in your library path by selecting File->Locate... (**Ctrl-L**).

Java and Powerbuilder support function overloading by argument list, *i.e.* one class can have multiple methods with the same name.

Guido, the original author of Python, explains method overriding this way: "Derived classes may override methods of their base class."

Always assign an initial value to all of an instance's data attributes in the `__init__` method. It will save you hours of debugging later, tracking down uninitialized attributes.

In versions of Python prior to 2.2, you could not directly subclass built-in datatypes like strings, lists, and dictionaries. To compensate, you could use a wrapper class.

5.6.1. Getting and Setting Items

When accessing data attributes within a class, you need to qualify the attribute name: `obj.attr`. When calling other methods within a class, you use `obj.method()`.

5.7. Advanced Special Class Methods

In Java, you determine whether two string variables reference the same physical memory location by using `obj == other_obj`. This is called *object identity*.

While other object-oriented languages only let you define the physical model of an object ("this object has a `method()`"), Python's `__str__` and `__repr__` methods let you define the logical model of an object.

5.8. Introducing Class Attributes

In Java, both static variables (called class attributes in Python) and instance variables (called data attributes in Python) are defined in the class definition.

There are no constants in Python. Everything can be changed if you try hard enough. This fits with one of the core principles of Python: *everything is mutable*.

5.9. Private Functions

In Python, all special methods (like `__init__`) and built-in attributes (like `__name__`) follow a standard naming convention: they both start with and end with two underscores.

Chapter 6. Exceptions and File Handling

6.1. Handling Exceptions

Python uses `try/except` to handle exceptions and `raise` to generate them. Java and C++ use `try/catch` to handle exceptions, and `throw` to generate them.

6.5. Working with Directories

Whenever possible, you should use the functions in `os` and `shutil` for file, directory, and path manipulations. These modules are wrappers for the `os` and `shell` commands.

Chapter 7. Regular Expressions

7.4. Using the `{n,m}` Syntax

There is no way to programmatically determine that two regular expressions are equivalent. The best you can do is write a lot of tests.

Chapter 8. HTML Processing

8.2. Introducing `sgmlib.py`

Python 2.0 had a bug where `SGMLParser` would not recognize declarations at all (`parser.parseDeclaration()` would never be called), which meant that `SGMLParser` were silently ignored.

In the ActivePython IDE on Windows, you can specify command line arguments in the "Run script" dialog. Separate multiple arguments with spaces.

8.4. Introducing `BaseHTMLProcessor.py`

The HTML specification requires that all non-HTML (like client-side JavaScript) must be enclosed in HTML comments, but not all HTML processors support this.

8.5. locals and globals

Python 2.2 introduced a subtle but important change that affects the namespace search order: nested scopes. In versions of Python prior to 2.2, the search order was global, local, then global. In Python 2.2, the search order is local, then global, then global.

Using the `getattr` and `setattr` functions, you can get the value of arbitrary variables dynamically, providing the variable name as a string. This mirrors the behavior of the `getattr` and `setattr` functions in Python 2.2.

8.6. Dictionary-based string formatting

Using dictionary-based string formatting with `dict` is a convenient way of making complex string formatting expressions more readable.

Chapter 9. XML Processing

9.2. Packages

A package is a directory with the special `__init__.py` file in it. The `__init__.py` file defines the attributes and methods of the package. It doesn't need to define anything, but it must exist.

9.6. Accessing element attributes

This section may be a little confusing, because of some overlapping terminology. Elements in an XML document have attributes, and attributes have values.

Like a dictionary, attributes of an XML element have no ordering. Attributes may *happen to be* listed in a certain order in the original document, but that order is not guaranteed.

Chapter 10. Scripts and Streams

Chapter 11. HTTP Web Services

11.6. Handling Last-Modified and ETag

In these examples, the HTTP server has supported both `Last-Modified` and `ETag` headers, but not all servers do. As a web services client, you should be able to handle either.

Chapter 12. SOAP Web Services

Chapter 13. Unit Testing

13.2. Diving in

`unittest` is included with Python 2.1 and later. Python 2.0 users can download it from <http://pyunit.sourceforge.net/>.

Chapter 14. Test-First Programming

14.3. roman.py, stage 3

The most important thing that comprehensive unit testing can tell you is when to stop coding. When all the unit tests for a function pass, you can stop coding.

14.5. roman.py, stage 5

When all of your tests pass, stop coding.

Chapter 15. Refactoring

15.3. Refactoring

Whenever you are going to use a regular expression more than once, you should compile it to get a pattern object, then call the method

Chapter 16. Functional Programming

16.2. Finding the path

The pathnames and filenames you pass to `os.path` do not need to exist.

`os.path.abspath()` not only constructs full path names, it also normalizes them. That means that if you are in the `current` directory, `os.path.abspath()` will return `os.path.abspath()` It normalizes

Like the other functions in the `os.path` module, `os.path.abspath()` is cross-platform. Your results will look slightly different than my examples if you're

Chapter 17. Dynamic functions

Chapter 18. Performance Tuning

18.2. Using the timeit Module

You can use the `timeit` module on the command line to test an existing Python program, without modifying the code. See <http://docs.python.org/2/library/timeit.html>

The `timeit` module only works if you already know what piece of code you need to optimize. If you have a larger Python program and don't

Appendix D. List of examples

Chapter 1. Installing Python

- 1.3. Python on Mac OS X
 - Example 1.1. Two versions of Python
- 1.5. Python on RedHat Linux
 - Example 1.2. Installing on RedHat Linux 9
- 1.6. Python on Debian GNU/Linux
 - Example 1.3. Installing on Debian GNU/Linux
- 1.7. Python Installation from Source
 - Example 1.4. Installing from source
- 1.8. The Interactive Shell
 - Example 1.5. First Steps in the Interactive Shell

Chapter 2. Your First Python Program

- 2.1. Diving in
 - Example 2.1. odbchelper.py
- 2.3. Documenting Functions
 - Example 2.2. Defining the buildConnectionString Function's doc string
- 2.4. Everything Is an Object
 - Example 2.3. Accessing the buildConnectionString Function's doc string
- 2.4.1. The Import Search Path
 - Example 2.4. Import Search Path
- 2.5. Indenting Code
 - Example 2.5. Indenting the buildConnectionString Function
 - Example 2.6. if Statements

Chapter 3. Native Datatypes

- 3.1.1. Defining Dictionaries
 - Example 3.1. Defining a Dictionary
- 3.1.2. Modifying Dictionaries
 - Example 3.2. Modifying a Dictionary
 - Example 3.3. Dictionary Keys Are Case-Sensitive
 - Example 3.4. Mixing Datatypes in a Dictionary
- 3.1.3. Deleting Items From Dictionaries
 - Example 3.5. Deleting Items from a Dictionary
- 3.2.1. Defining Lists
 - Example 3.6. Defining a List
 - Example 3.7. Negative List Indices
 - Example 3.8. Slicing a List
 - Example 3.9. Slicing Shorthand
- 3.2.2. Adding Elements to Lists
 - Example 3.10. Adding Elements to a List
 - Example 3.11. The Difference between extend and append
- 3.2.3. Searching Lists
 - Example 3.12. Searching a List
- 3.2.4. Deleting List Elements
 - Example 3.13. Removing Elements from a List
- 3.2.5. Using List Operators
 - Example 3.14. List Operators

- 3.3. Introducing Tuples
 - Example 3.15. Defining a tuple
 - Example 3.16. Tuples Have No Methods
- 3.4. Declaring variables
 - Example 3.17. Defining the myParams Variable
- 3.4.1. Referencing Variables
 - Example 3.18. Referencing an Unbound Variable
- 3.4.2. Assigning Multiple Values at Once
 - Example 3.19. Assigning multiple values at once
 - Example 3.20. Assigning Consecutive Values
- 3.5. Formatting Strings
 - Example 3.21. Introducing String Formatting
 - Example 3.22. String Formatting vs. Concatenating
 - Example 3.23. Formatting Numbers
- 3.6. Mapping Lists
 - Example 3.24. Introducing List Comprehensions
 - Example 3.25. The keys, values, and items Functions
 - Example 3.26. List Comprehensions in buildConnectionString, Step by Step
- 3.7. Joining Lists and Splitting Strings
 - Example 3.27. Output of odbchelper.py
 - Example 3.28. Splitting a String

Chapter 4. The Power Of Introspection

- 4.1. Diving In
 - Example 4.1. apihelper.py
 - Example 4.2. Sample Usage of apihelper.py
 - Example 4.3. Advanced Usage of apihelper.py
- 4.2. Using Optional and Named Arguments
 - Example 4.4. Valid Calls of info
- 4.3.1. The type Function
 - Example 4.5. Introducing type
- 4.3.2. The str Function
 - Example 4.6. Introducing str
 - Example 4.7. Introducing dir
 - Example 4.8. Introducing callable
- 4.3.3. Built-In Functions
 - Example 4.9. Built-in Attributes and Functions
- 4.4. Getting Object References With getattr
 - Example 4.10. Introducing getattr
- 4.4.1. getattr with Modules
 - Example 4.11. The getattr Function in apihelper.py
- 4.4.2. getattr As a Dispatcher
 - Example 4.12. Creating a Dispatcher with getattr
 - Example 4.13. getattr Default Values
- 4.5. Filtering Lists
 - Example 4.14. Introducing List Filtering
- 4.6. The Peculiar Nature of and and or
 - Example 4.15. Introducing and
 - Example 4.16. Introducing or
- 4.6.1. Using the and-or Trick
 - Example 4.17. Introducing the and-or Trick
 - Example 4.18. When the and-or Trick Fails

- Example 4.19. Using the and-or Trick Safely
- 4.7. Using lambda Functions
 - Example 4.20. Introducing lambda Functions
- 4.7.1. Real-World lambda Functions
 - Example 4.21. split With No Arguments
- 4.8. Putting It All Together
 - Example 4.22. Getting a doc string Dynamically
 - Example 4.23. Why Use str on a doc string?
 - Example 4.24. Introducing ljust
 - Example 4.25. Printing a List

Chapter 5. Objects and Object-Orientation

- 5.1. Diving In
 - Example 5.1. fileinfo.py
- 5.2. Importing Modules Using from module import
 - Example 5.2. import module vs. from module import
- 5.3. Defining Classes
 - Example 5.3. The Simplest Python Class
 - Example 5.4. Defining the FileInfo Class
- 5.3.1. Initializing and Coding Classes
 - Example 5.5. Initializing the FileInfo Class
 - Example 5.6. Coding the FileInfo Class
- 5.4. Instantiating Classes
 - Example 5.7. Creating a FileInfo Instance
- 5.4.1. Garbage Collection
 - Example 5.8. Trying to Implement a Memory Leak
- 5.5. Exploring UserDict: A Wrapper Class
 - Example 5.9. Defining the UserDict Class
 - Example 5.10. UserDict Normal Methods
 - Example 5.11. Inheriting Directly from Built-In Datatype dict
- 5.6.1. Getting and Setting Items
 - Example 5.12. The __getitem__ Special Method
 - Example 5.13. The __setitem__ Special Method
 - Example 5.14. Overriding __setitem__ in MP3FileInfo
 - Example 5.15. Setting an MP3FileInfo's name
- 5.7. Advanced Special Class Methods
 - Example 5.16. More Special Methods in UserDict
- 5.8. Introducing Class Attributes
 - Example 5.17. Introducing Class Attributes
 - Example 5.18. Modifying Class Attributes
- 5.9. Private Functions
 - Example 5.19. Trying to Call a Private Method

Chapter 6. Exceptions and File Handling

- 6.1. Handling Exceptions
 - Example 6.1. Opening a Non-Existent File
- 6.1.1. Using Exceptions For Other Purposes
 - Example 6.2. Supporting Platform-Specific Functionality
- 6.2. Working with File Objects
 - Example 6.3. Opening a File
- 6.2.1. Reading Files

- Example 6.4. Reading a File
- 6.2.2. Closing Files
 - Example 6.5. Closing a File
- 6.2.3. Handling I/O Errors
 - Example 6.6. File Objects in MP3FileInfo
- 6.2.4. Writing to Files
 - Example 6.7. Writing to Files
- 6.3. Iterating with for Loops
 - Example 6.8. Introducing the for Loop
 - Example 6.9. Simple Counters
 - Example 6.10. Iterating Through a Dictionary
 - Example 6.11. for Loop in MP3FileInfo
- 6.4. Using sys.modules
 - Example 6.12. Introducing sys.modules
 - Example 6.13. Using sys.modules
 - Example 6.14. The `__module__` Class Attribute
 - Example 6.15. sys.modules in fileinfo.py
- 6.5. Working with Directories
 - Example 6.16. Constructing Pathnames
 - Example 6.17. Splitting Pathnames
 - Example 6.18. Listing Directories
 - Example 6.19. Listing Directories in fileinfo.py
 - Example 6.20. Listing Directories with glob
- 6.6. Putting It All Together
 - Example 6.21. listDirectory

Chapter 7. Regular Expressions

- 7.2. Case Study: Street Addresses
 - Example 7.1. Matching at the End of a String
 - Example 7.2. Matching Whole Words
- 7.3.1. Checking for Thousands
 - Example 7.3. Checking for Thousands
- 7.3.2. Checking for Hundreds
 - Example 7.4. Checking for Hundreds
- 7.4. Using the {n,m} Syntax
 - Example 7.5. The Old Way: Every Character Optional
 - Example 7.6. The New Way: From n o m
- 7.4.1. Checking for Tens and Ones
 - Example 7.7. Checking for Tens
 - Example 7.8. Validating Roman Numerals with {n,m}
- 7.5. Verbose Regular Expressions
 - Example 7.9. Regular Expressions with Inline Comments
- 7.6. Case study: Parsing Phone Numbers
 - Example 7.10. Finding Numbers
 - Example 7.11. Finding the Extension
 - Example 7.12. Handling Different Separators
 - Example 7.13. Handling Numbers Without Separators
 - Example 7.14. Handling Leading Characters
 - Example 7.15. Phone Number, Wherever I May Find Ye
 - Example 7.16. Parsing Phone Numbers (Final Version)

Chapter 8. HTML Processing

- 8.1. Diving in
 - Example 8.1. BaseHTMLProcessor.py
 - Example 8.2. dialect.py
 - Example 8.3. Output of dialect.py
- 8.2. Introducing sgmlib.py
 - Example 8.4. Sample test of sgmlib.py
- 8.3. Extracting data from HTML documents
 - Example 8.5. Introducing urllib
 - Example 8.6. Introducing urlister.py
 - Example 8.7. Using urlister.py
- 8.4. Introducing BaseHTMLProcessor.py
 - Example 8.8. Introducing BaseHTMLProcessor
 - Example 8.9. BaseHTMLProcessor output
- 8.5. locals and globals
 - Example 8.10. Introducing locals
 - Example 8.11. Introducing globals
 - Example 8.12. locals is read-only, globals is not
- 8.6. Dictionary-based string formatting
 - Example 8.13. Introducing dictionary-based string formatting
 - Example 8.14. Dictionary-based string formatting in BaseHTMLProcessor.py
 - Example 8.15. More dictionary-based string formatting
- 8.7. Quoting attribute values
 - Example 8.16. Quoting attribute values
- 8.8. Introducing dialect.py
 - Example 8.17. Handling specific tags
 - Example 8.18. SGMLParser
 - Example 8.19. Overriding the handle_data method
- 8.9. Putting it all together
 - Example 8.20. The translate function, part 1
 - Example 8.21. The translate function, part 2: curiouser and curiouser
 - Example 8.22. The translate function, part 3

Chapter 9. XML Processing

- 9.1. Diving in
 - Example 9.1. kgp.py
 - Example 9.2. toolbox.py
 - Example 9.3. Sample output of kgp.py
 - Example 9.4. Simpler output from kgp.py
- 9.2. Packages
 - Example 9.5. Loading an XML document (a sneak peek)
 - Example 9.6. File layout of a package
 - Example 9.7. Packages are modules, too
- 9.3. Parsing XML
 - Example 9.8. Loading an XML document (for real this time)
 - Example 9.9. Getting child nodes
 - Example 9.10. toxml works on any node
 - Example 9.11. Child nodes can be text
 - Example 9.12. Drilling down all the way to text
- 9.4. Unicode
 - Example 9.13. Introducing unicode
 - Example 9.14. Storing non-ASCII characters
 - Example 9.15. sitecustomize.py

Example 9.16. Effects of setting the default encoding

Example 9.17. Specifying encoding in .py files

Example 9.18. russiansample.xml

Example 9.19. Parsing russiansample.xml

9.5. Searching for elements

Example 9.20. binary.xml

Example 9.21. Introducing getElementByTagName

Example 9.22. Every element is searchable

Example 9.23. Searching is actually recursive

9.6. Accessing element attributes

Example 9.24. Accessing element attributes

Example 9.25. Accessing individual attributes

Chapter 10. Scripts and Streams

10.1. Abstracting input sources

Example 10.1. Parsing XML from a file

Example 10.2. Parsing XML from a URL

Example 10.3. Parsing XML from a string (the easy but inflexible way)

Example 10.4. Introducing StringIO

Example 10.5. Parsing XML from a string (the file-like object way)

Example 10.6. openAnything

Example 10.7. Using openAnything in kgp.py

10.2. Standard input, output, and error

Example 10.8. Introducing stdout and stderr

Example 10.9. Redirecting output

Example 10.10. Redirecting error information

Example 10.11. Printing to stderr

Example 10.12. Chaining commands

Example 10.13. Reading from standard input in kgp.py

10.3. Caching node lookups

Example 10.14. loadGrammar

Example 10.15. Using the ref element cache

10.4. Finding direct children of a node

Example 10.16. Finding direct child elements

10.5. Creating separate handlers by node type

Example 10.17. Class names of parsed XML objects

Example 10.18. parse, a generic XML node dispatcher

Example 10.19. Functions called by the parse dispatcher

10.6. Handling command-line arguments

Example 10.20. Introducing sys.argv

Example 10.21. The contents of sys.argv

Example 10.22. Introducing getopt

Example 10.23. Handling command-line arguments in kgp.py

Chapter 11. HTTP Web Services

11.1. Diving in

Example 11.1. openanything.py

11.2. How not to fetch data over HTTP

Example 11.2. Downloading a feed the quick-and-dirty way

11.4. Debugging HTTP web services

Example 11.3. Debugging HTTP

- 11.5. Setting the User-Agent
 - Example 11.4. Introducing urllib2
 - Example 11.5. Adding headers with the Request
- 11.6. Handling Last-Modified and ETag
 - Example 11.6. Testing Last-Modified
 - Example 11.7. Defining URL handlers
 - Example 11.8. Using custom URL handlers
 - Example 11.9. Supporting ETag/If-None-Match
- 11.7. Handling redirects
 - Example 11.10. Accessing web services without a redirect handler
 - Example 11.11. Defining the redirect handler
 - Example 11.12. Using the redirect handler to detect permanent redirects
 - Example 11.13. Using the redirect handler to detect temporary redirects
- 11.8. Handling compressed data
 - Example 11.14. Telling the server you would like compressed data
 - Example 11.15. Decompressing the data
 - Example 11.16. Decompressing the data directly from the server
- 11.9. Putting it all together
 - Example 11.17. The openanything function
 - Example 11.18. The fetch function
 - Example 11.19. Using openanything.py

Chapter 12. SOAP Web Services

- 12.1. Diving In
 - Example 12.1. search.py
 - Example 12.2. Sample Usage of search.py
- 12.2.1. Installing PyXML
 - Example 12.3. Verifying PyXML Installation
- 12.2.2. Installing fpconst
 - Example 12.4. Verifying fpconst Installation
- 12.2.3. Installing SOAPpy
 - Example 12.5. Verifying SOAPpy Installation
- 12.3. First Steps with SOAP
 - Example 12.6. Getting the Current Temperature
- 12.4. Debugging SOAP Web Services
 - Example 12.7. Debugging SOAP Web Services
- 12.6. Introspecting SOAP Web Services with WSDL
 - Example 12.8. Discovering The Available Methods
 - Example 12.9. Discovering A Method's Arguments
 - Example 12.10. Discovering A Method's Return Values
 - Example 12.11. Calling A Web Service Through A WSDL Proxy
- 12.7. Searching Google
 - Example 12.12. Introspecting Google Web Services
 - Example 12.13. Searching Google
 - Example 12.14. Accessing Secondary Information From Google
- 12.8. Troubleshooting SOAP Web Services
 - Example 12.15. Calling a Method With an Incorrectly Configured Proxy
 - Example 12.16. Calling a Method With the Wrong Arguments
 - Example 12.17. Calling a Method and Expecting the Wrong Number of Return Values
 - Example 12.18. Calling a Method With An Application-Specific Error

Chapter 13. Unit Testing

- 13.3. Introducing romantest.py
 - Example 13.1. romantest.py
- 13.4. Testing for success
 - Example 13.2. testToRomanKnownValues
- 13.5. Testing for failure
 - Example 13.3. Testing bad input to toRoman
 - Example 13.4. Testing bad input to fromRoman
- 13.6. Testing for sanity
 - Example 13.5. Testing toRoman against fromRoman
 - Example 13.6. Testing for case

Chapter 14. Test-First Programming

- 14.1. roman.py, stage 1
 - Example 14.1. roman1.py
 - Example 14.2. Output of romantest1.py against roman1.py
- 14.2. roman.py, stage 2
 - Example 14.3. roman2.py
 - Example 14.4. How toRoman works
 - Example 14.5. Output of romantest2.py against roman2.py
- 14.3. roman.py, stage 3
 - Example 14.6. roman3.py
 - Example 14.7. Watching toRoman handle bad input
 - Example 14.8. Output of romantest3.py against roman3.py
- 14.4. roman.py, stage 4
 - Example 14.9. roman4.py
 - Example 14.10. How fromRoman works
 - Example 14.11. Output of romantest4.py against roman4.py
- 14.5. roman.py, stage 5
 - Example 14.12. roman5.py
 - Example 14.13. Output of romantest5.py against roman5.py

Chapter 15. Refactoring

- 15.1. Handling bugs
 - Example 15.1. The bug
 - Example 15.2. Testing for the bug (romantest61.py)
 - Example 15.3. Output of romantest61.py against roman61.py
 - Example 15.4. Fixing the bug (roman62.py)
 - Example 15.5. Output of romantest62.py against roman62.py
- 15.2. Handling changing requirements
 - Example 15.6. Modifying test cases for new requirements (romantest71.py)
 - Example 15.7. Output of romantest71.py against roman71.py
 - Example 15.8. Coding the new requirements (roman72.py)
 - Example 15.9. Output of romantest72.py against roman72.py
- 15.3. Refactoring
 - Example 15.10. Compiling regular expressions
 - Example 15.11. Compiled regular expressions in roman81.py
 - Example 15.12. Output of romantest81.py against roman81.py
 - Example 15.13. roman82.py
 - Example 15.14. Output of romantest82.py against roman82.py
 - Example 15.15. roman83.py
 - Example 15.16. Output of romantest83.py against roman83.py

15.4. Postscript

Example 15.17. roman9.py

Example 15.18. Output of romantest9.py against roman9.py

Chapter 16. Functional Programming

16.1. Diving in

Example 16.1. regression.py

Example 16.2. Sample output of regression.py

16.2. Finding the path

Example 16.3. fullpath.py

Example 16.4. Further explanation of os.path.abspath

Example 16.5. Sample output from fullpath.py

Example 16.6. Running scripts in the current directory

16.3. Filtering lists revisited

Example 16.7. Introducing filter

Example 16.8. filter in regression.py

Example 16.9. Filtering using list comprehensions instead

16.4. Mapping lists revisited

Example 16.10. Introducing map

Example 16.11. map with lists of mixed datatypes

Example 16.12. map in regression.py

16.6. Dynamically importing modules

Example 16.13. Importing multiple modules at once

Example 16.14. Importing modules dynamically

Example 16.15. Importing a list of modules dynamically

16.7. Putting it all together

Example 16.16. The regressionTest function

Example 16.17. Step 1: Get all the files

Example 16.18. Step 2: Filter to find the files you care about

Example 16.19. Step 3: Map filenames to module names

Example 16.20. Step 4: Mapping module names to modules

Example 16.21. Step 5: Loading the modules into a test suite

Example 16.22. Step 6: Telling unittest to use your test suite

Chapter 17. Dynamic functions

17.2. plural.py, stage 1

Example 17.1. plural1.py

Example 17.2. Introducing re.sub

Example 17.3. Back to plural1.py

Example 17.4. More on negation regular expressions

Example 17.5. More on re.sub

17.3. plural.py, stage 2

Example 17.6. plural2.py

Example 17.7. Unrolling the plural function

17.4. plural.py, stage 3

Example 17.8. plural3.py

17.5. plural.py, stage 4

Example 17.9. plural4.py

Example 17.10. plural4.py continued

Example 17.11. Unrolling the rules definition

Example 17.12. plural4.py, finishing up

- Example 17.13. Another look at buildMatchAndApplyFunctions
- Example 17.14. Expanding tuples when calling functions
- 17.6. plural.py, stage 5
 - Example 17.15. rules.en
 - Example 17.16. plural5.py
- 17.7. plural.py, stage 6
 - Example 17.17. plural6.py
 - Example 17.18. Introducing generators
 - Example 17.19. Using generators instead of recursion
 - Example 17.20. Generators in for loops
 - Example 17.21. Generators that generate dynamic functions

Chapter 18. Performance Tuning

- 18.1. Diving in
 - Example 18.1. soundex/stage1/soundex1a.py
- 18.2. Using the timeit Module
 - Example 18.2. Introducing timeit
- 18.3. Optimizing Regular Expressions
 - Example 18.3. Best Result So Far: soundex/stage1/soundex1e.py
- 18.4. Optimizing Dictionary Lookups
 - Example 18.4. Best Result So Far: soundex/stage2/soundex2c.py
- 18.5. Optimizing List Operations
 - Example 18.5. Best Result So Far: soundex/stage2/soundex2c.py

Appendix E. Revision history

Revision History	
Revision 5.4	2004-05-20
<p>Added Section 12.1, “Diving In”.</p> <p>Added Section 12.2, “Installing the SOAP Libraries”.</p> <p>Added Section 12.3, “First Steps with SOAP”.</p> <p>Added Section 12.4, “Debugging SOAP Web Services”.</p> <p>Added Section 12.5, “Introducing WSDL”.</p> <p>Added Section 12.6, “Introspecting SOAP Web Services with WSDL”.</p> <p>Added Section 12.7, “Searching Google”.</p> <p>Added Section 12.8, “Troubleshooting SOAP Web Services”.</p> <p>Added Section 12.9, “Summary”.</p> <p>Incorporated technical reviewer revisions in Chapter 16, <i>Functional Programming</i> and Chapter 18, <i>Performance Tuning</i>.</p>	
Revision 5.3	2004-05-12
<p>Added example to Section 18.3, “Optimizing Regular Expressions”. Thanks, Paul.</p> <p>Incorporated copyediting revisions into Chapter 5, <i>Objects and Object-Oriented</i> and Chapter 6, <i>Exceptions and File Handling</i>.</p> <p>Fixed URL of Section 9.7, “Segue”.</p>	
Revision 5.2	2004-05-09
<p>Fixed URL of Section 14.1, “roman.py, stage 1”.</p> <p>Added Section 18.1, “Diving in”.</p> <p>Added Section 18.2, “Using the timeit Module”.</p> <p>Added Section 18.3, “Optimizing Regular Expressions”.</p> <p>Added Section 18.4, “Optimizing Dictionary Lookups”.</p> <p>Added Section 18.5, “Optimizing List Operations”.</p> <p>Added Section 18.6, “Optimizing String Manipulation”.</p> <p>Added Section 18.7, “Summary”.</p>	
Revision 5.1	2004-05-05
<p>Clarified Example 7.7, “Checking for Tens” and Example 7.8, “Validating Roman Numerals with {n,m}”.</p> <p>Clarified Example 7.10, “Finding Numbers”.</p> <p>Fixed typo in Example 11.6, “Testing Last-Modified”. Thanks, Jesir.</p> <p>Fixed typo in Example 3.11, “The Difference between extend and append”. Thanks, Daniel.</p> <p>Incorporated technical reviewer revisions.</p> <p>Incorporated copy editor revisions in Chapter 1, <i>Installing Python</i>, Chapter 2, <i>Your First Python Program</i>, Chapter 3, <i>Native Data</i></p>	
Revision 5.0	2004-04-16
<p>Added Section 11.1, “Diving in”.</p> <p>Added Section 11.2, “How not to fetch data over HTTP”.</p> <p>Added Section 11.3, “Features of HTTP”.</p> <p>Added Section 11.4, “Debugging HTTP web services”.</p> <p>Added Section 11.5, “Setting the User-Agent”.</p> <p>Added Section 11.6, “Handling Last-Modified and ETag”.</p> <p>Added Section 11.7, “Handling redirects”.</p> <p>Added Section 11.8, “Handling compressed data”.</p> <p>Added Section 11.9, “Putting it all together”.</p> <p>Added Section 11.10, “Summary”.</p> <p>Added Example 3.11, “The Difference between extend and append”.</p> <p>Changed descriptions of how to download Python throughout Chapter 1, <i>Installing Python</i> to be more generic and less version-specific.</p> <p>Changed references of “module” to “program” in Section 2.1, “Diving in” and Section 2.4, “Everything Is an Object” since we have</p> <p>Added explicit instructions in Section 2.4, “Everything Is an Object” for the reader to open their Python IDE and follow along with</p>	

Changed all examples and descriptions that referred to truth values `1` and `0` to refer to `True` and `False`.
 Updated Example 3.22, “String Formatting vs. Concatenating” to show new Python 2.3 `message`.
 Fixed typo in Example 17.19, “Using generators instead of recursion”.
 Fixed typo in Section 7.7, “Summary”.
 Fixed typo in Example 17.9, “plural4.py”.

Revision 4.9	2004-03-25
<p>Finished Section 16.7, “Putting it all together”. Added Section 16.8, “Summary”. Split unit testing introduction into two chapters, Chapter 13, <i>Unit Testing</i> and Chapter 14, <i>Test-First Programming</i>. Fixed typo in Example 17.12, “plural4.py, finishing up”. Fixed typo in Example 17.18, “Introducing generators”.</p>	
Revision 4.8	2004-03-25
<p>Finished Section 17.7, “plural.py, stage 6”. Finished Section 17.8, “Summary”. Fixed broken links in Appendix A, <i>Further reading</i>, Appendix B, <i>A 5-minute review</i>, Appendix C, <i>Tips and tricks</i>, Appendix D, <i>List of...</i></p>	
Revision 4.7	2004-03-21
<p>Added Section 17.1, “Diving in”. Added Section 17.2, “plural.py, stage 1”. Added Section 17.3, “plural.py, stage 2”. Added Section 17.4, “plural.py, stage 3”. Added Section 17.5, “plural.py, stage 4”. Added Section 17.6, “plural.py, stage 5”. Added Section 17.7, “plural.py, stage 6” (unfinished). Added Section 17.8, “Summary” (unfinished).</p>	
Revision 4.6	2004-03-14
<p>Finished Section 7.4, “Using the {n,m} Syntax”. Finished Section 7.5, “Verbose Regular Expressions”. Finished Section 7.6, “Case study: Parsing Phone Numbers”. Expanded Section 7.7, “Summary”.</p>	
Revision 4.5	2004-03-07
<p>Added Section 7.1, “Diving In”. Added Section 7.4, “Using the {n,m} Syntax” (incomplete). Added Section 7.5, “Verbose Regular Expressions” (incomplete). Added Section 7.6, “Case study: Parsing Phone Numbers” (incomplete). Added Section 7.7, “Summary”. Moved Section 7.2, “Case Study: Street Addresses” and Section 7.3, “Case Study: Roman Numerals” to regular expressions chapter. Added Example 6.20, “Listing Directories with glob”. Added Example 6.7, “Writing to Files”. Added Example 5.11, “Inheriting Directly from Built-In Datatype dict”. Added Example 10.11, “Printing to stderr”. Added Example 4.12, “Creating a Dispatcher with getattr” and Example 4.13, “getattr Default Values”. Added Example 2.6, “if Statements”. Added Example 3.23, “Formatting Numbers”. Split Chapter 5, <i>Objects and Object-Oriented</i> into 2 chapters: Chapter 5, <i>Objects and Object-Oriented</i> and Chapter 6, <i>Exceptions</i>. Split Chapter 9, <i>XML Processing</i> into 2 chapters: Chapter 9, <i>XML Processing</i> and Chapter 10, <i>Scripts and Streams</i>. Split Chapter 13, <i>Unit Testing</i> into 2 chapters: Chapter 13, <i>Unit Testing</i> and Chapter 15, <i>Refactoring</i>. Renamed <code>no</code> in Chapter 4, <i>The Power Of Introspection</i>. Fixed incorrect back-reference in Section 8.5, “locals and globals”. Fixed broken example links in Section 8.1, “Diving in”. Fixed missing line in example in Section 9.1, “Diving in”.</p>	

Fixed typo in Section 8.2, “Introducing <code>sgmlib.py</code> ”.	
Revision 4.4	2003-10-08
<p>Added Section 1.1, “Which Python is right for you?”.</p> <p>Added Section 1.2, “Python on Windows”.</p> <p>Added Section 1.3, “Python on Mac OS X”.</p> <p>Added Section 1.4, “Python on Mac OS 9”.</p> <p>Added Section 1.5, “Python on RedHat Linux”.</p> <p>Added Section 1.6, “Python on Debian GNU/Linux”.</p> <p>Added Section 1.7, “Python Installation from Source”.</p> <p>Added Section 1.9, “Summary”.</p> <p>Removed preface.</p> <p>Fixed typo in Example 3.27, “Output of <code>odbc helper.py</code>”.</p> <p>Added link to PEP 257 in Section 2.3, “Documenting Functions”.</p> <p>Fixed link to <i>How to Think Like a Computer Scientist</i> (http://www.ibiblio.org/obp/thinkCSpy/) in Section 3.4.2, “Assigning Multiple Values”.</p> <p>Added note about implied assert in Section 3.3, “Introducing Tuples”.</p>	
Revision 4.3	2003-09-28
<p>Added Section 16.6, “Dynamically importing modules”.</p> <p>Added Section 16.7, “Putting it all together” (incomplete).</p> <p>Fixed links in Appendix F, <i>About the book</i>.</p>	
Revision 4.2.1	2003-09-17
<p>Fixed links on index page.</p> <p>Fixed syntax highlighting.</p>	
Revision 4.2	2003-09-12
<p>Fixed typos in Section 16.4, “Mapping lists revisited”, Section 16.3, “Filtering lists revisited”, Section 7.2, “Case Study: Street Address”.</p> <p>Fixed external link in Section 5.3, “Defining Classes”. Thanks, Harry.</p> <p>Changed wording at the end of Section 4.5, “Filtering Lists”. Thanks, Paul.</p> <p>Added sentence in Section 13.5, “Testing for failure” to make it clearer that we’re passing a function to <code>isinstance()</code>, not a function name as a string.</p> <p>Fixed typo in Section 8.8, “Introducing <code>dialect.py</code>”. Thanks, Wellie.</p> <p>Fixed links to dialectized examples.</p> <p>Fixed external link to the history of Roman numerals. Thanks to many concerned Roman numeral fans around the world.</p>	
Revision 4.1	2002-07-28
<p>Added Section 10.3, “Caching node lookups”.</p> <p>Added Section 10.4, “Finding direct children of a node”.</p> <p>Added Section 10.5, “Creating separate handlers by node type”.</p> <p>Added Section 10.6, “Handling command-line arguments”.</p> <p>Added Section 10.7, “Putting it all together”.</p> <p>Added Section 10.8, “Summary”.</p> <p>Fixed typo in Section 6.5, “Working with Directories”. It’s not <code>os.listdir()</code>. Thanks, Abhishek.</p> <p>Fixed typo in Section 3.7, “Joining Lists and Splitting Strings”. When evaluated (instead of printed), the Python IDE will display single quotes around the string.</p> <p>Changed example in Section 4.8, “Putting It All Together” to use a user-defined function, since Python 2.2 obsoleted the old example.</p> <p>Fixed typo in Section 9.4, “Unicode”, “anyway” to “anywhere”. Thanks Frank.</p> <p>Fixed typo in Section 13.6, “Testing for sanity”, doubled word “accept”. Thanks Ralph.</p> <p>Fixed typo in Section 15.3, “Refactoring”, “Matches 0 to 3 Characters, not 4”. Thanks Ralph.</p> <p>Clarified and expanded explanation of implied slice indices in Example 3.9, “Slicing Shorthand”. Thanks Petr.</p> <p>Added historical note in Section 5.5, “Exploring UserDict: A Wrapper Class” now that Python 2.2 supports subclassing built-in data types.</p> <p>Added explanation of <code>dict.items()</code> method in Example 5.9, “Defining the UserDict Class”. Thanks Petr.</p> <p>Clarified Python’s lack of overloading in Section 5.5, “Exploring UserDict: A Wrapper Class”. Thanks Petr.</p> <p>Fixed typo in Example 8.8, “Introducing <code>BaseHTMLProcessor</code>”. HTML comments end with two dashes and a bracket, not one. Thanks Petr.</p> <p>Changed tense of note about nested scopes in Section 8.5, “Locals and Globals” now that Python 2.2 is out. Thanks Petr.</p> <p>Fixed typo in Example 8.14, “Dictionary-based string formatting in <code>BaseHTMLProcessor.py</code>”; a space should have been a non-breaking space.</p>	

	<p>Added title to note on derived classes in Section 5.5, “Exploring UserDict: A Wrapper Class”. Thanks Petr.</p> <p>Added title to note on downloading in Section 15.3, “Refactoring”. Thanks Petr.</p> <p>Fixed typesetting problem in Example 16.6, “Running scripts in the current directory”; tabs should have been spaces, and the line number should be right-aligned.</p> <p>Fixed capitalization typo in the tip on truth values in Section 3.2, “Introducing Lists”. It’s <code>and</code> not <code>AND</code>. Thanks to everyone who pointed it out.</p> <p>Changed section titles of Section 3.1, “Introducing Dictionaries”, Section 3.2, “Introducing Lists”, and Section 3.3, “Introducing Tuples” to be more descriptive.</p> <p>Upgraded to version 1.52 of the DocBook XSL stylesheets.</p> <p>Upgraded to version 6.52 of the SAXON XSLT processor from Michael Kay.</p> <p>Various accessibility-related stylesheet tweaks.</p> <p>Somewhere between this revision and the last one, she said yes. The wedding will be next spring.</p>
Revision 4.0-2	2002-04-26
	<p>Fixed typo in Example 4.15, “Introducing and”.</p> <p>Fixed typo in Example 2.4, “Import Search Path”.</p> <p>Fixed Windows help file (missing table of contents due to base stylesheet changes).</p>
Revision 4.0	2002-04-19
	<p>Expanded Section 2.4, “Everything Is an Object” to include more about import search paths.</p> <p>Fixed typo in Example 3.7, “Negative List Indices”. Thanks to Brian for the correction.</p> <p>Rewrote the tip on truth values in Section 3.2, “Introducing Lists”, now that Python has a separate boolean datatype.</p> <p>Fixed typo in Section 5.2, “Importing Modules Using from module import” when comparing syntax to Java. Thanks to Rick for the correction.</p> <p>Added note in Section 5.5, “Exploring UserDict: A Wrapper Class” about derived classes always overriding ancestor classes.</p> <p>Fixed typo in Example 5.18, “Modifying Class Attributes”. Thanks to Kevin for the correction.</p> <p>Added note in Section 6.1, “Handling Exceptions” that you can define and raise your own exceptions. Thanks to Rony for the suggestion.</p> <p>Fixed typo in Example 8.17, “Handling specific tags”. Thanks for Rick for the correction.</p> <p>Added note in Example 8.18, “SGMLParser” about what the return codes mean. Thanks to Howard for the suggestion.</p> <p>Added <code>__getattr__</code> function when creating <code>__instance</code> in Example 10.6, “openAnything”. Thanks to Ganesan for the idea.</p> <p>Added link in Section 13.3, “Introducing <code>romantest.py</code>” to explanation of why test cases belong in a separate file.</p> <p>Changed Section 16.2, “Finding the path” to use <code>os.path</code> instead of <code>os</code>. Thanks to Marc for the idea.</p> <p>Added code samples (<code>re.compile</code> and <code>re.match</code>) for the upcoming regular expressions chapter.</p> <p>Updated and expanded list of Python distributions on home page.</p>
Revision 3.9	2002-01-01
	<p>Added Section 9.4, “Unicode”.</p> <p>Added Section 9.5, “Searching for elements”.</p> <p>Added Section 9.6, “Accessing element attributes”.</p> <p>Added Section 10.1, “Abstracting input sources”.</p> <p>Added Section 10.2, “Standard input, output, and error”.</p> <p>Added simple counter loop examples (good usage and bad usage) in Section 6.3, “Iterating with for Loops”. Thanks to Kevin for the suggestion.</p> <p>Fixed typo in Example 3.25, “The keys, values, and items Functions” (two elements of <code>iteritems</code> were reversed).</p> <p>Fixed mistake in Section 4.3, “Using type, str, dir, and Other Built-In Functions” with regards to the name of the <code>__main__</code> module. Thanks to Rick for the correction.</p> <p>Added additional example in Section 16.2, “Finding the path” to show how to run unit tests in the current working directory, instead of the current directory.</p> <p>Modified explanation of how to derive a negative list index from a positive list index in Example 3.7, “Negative List Indices”. Thanks to Rick for the suggestion.</p> <p>Updated links on home page for downloading latest version of Python.</p> <p>Added link on home page to Bruce Eckel’s preliminary draft of <i>Thinking in Python</i> (http://www.mindview.net/Books/TIPython), a book about Python.</p>
Revision 3.8	2001-11-18
	<p>Added Section 16.2, “Finding the path”.</p> <p>Added Section 16.3, “Filtering lists revisited”.</p> <p>Added Section 16.4, “Mapping lists revisited”.</p> <p>Added Section 16.5, “Data-centric programming”.</p> <p>Expanded sample output in Section 16.1, “Diving in”.</p> <p>Finished Section 9.3, “Parsing XML”.</p>
Revision 3.7	2001-09-30
	<p>Added Section 9.2, “Packages”.</p>

	<p>Added Section 9.3, “Parsing XML”.</p> <p>Cleaned up introductory paragraph in Section 9.1, “Diving in”. Thanks to Matt for this suggestion.</p> <p>Added Java tip in Section 5.2, “Importing Modules Using from module import”. Thanks to Ori for this suggestion.</p> <p>Fixed mistake in Section 4.8, “Putting It All Together” where I implied that you could not use <code>is</code> to compare to a null value in Python.</p> <p>Clarified in Section 3.2, “Introducing Lists” where I said that <code>is</code> was equivalent to <code>==</code>. The result is the same, but <code>is</code> is faster because it does not create a new object.</p> <p>Fixed mistake in Section 3.2, “Introducing Lists” where I said that <code>is</code> was equivalent to <code>==</code>. In fact, it’s equivalent to <code>==</code> since it doesn’t create a new object.</p> <p>Fixed typographical laziness in Chapter 2, <i>Your First Python Program</i>; when I was writing it, I had not yet standardized on putting <code>is</code> before <code>is</code>.</p> <p>Fixed mistake in Section 2.2, “Declaring Functions” where I said that statically typed languages always use explicit variable declarations.</p> <p>Added link to Spanish translation (http://es.diveintopython.org/).</p>
Revision 3.6.4	2001-09-06
	<p>Added code in <code>is</code> to handle non-HTML entity references, and added a note about it in Section 8.4, “Introducing BaseHTMLProcessor”.</p> <p>Modified Example 8.11, “Introducing globals” to include <code>is</code> in the output.</p>
Revision 3.6.3	2001-09-04
	<p>Fixed typo in Section 9.1, “Diving in”.</p> <p>Added link to Korean translation (http://kr.diveintopython.org/html/index.htm).</p>
Revision 3.6.2	2001-08-31
	<p>Fixed typo in Section 13.6, “Testing for sanity” (the last requirement was listed twice).</p>
Revision 3.6	2001-08-31
	<p>Finished Chapter 8, <i>HTML Processing</i> with Section 8.9, “Putting it all together” and Section 8.10, “Summary”.</p> <p>Added Section 15.4, “Postscript”.</p> <p>Started Chapter 9, <i>XML Processing</i> with Section 9.1, “Diving in”.</p> <p>Started Chapter 16, <i>Functional Programming</i> with Section 16.1, “Diving in”.</p> <p>Fixed long-standing bug in colorizing script that improperly colorized the examples in Chapter 8, <i>HTML Processing</i>.</p> <p>Added link to French translation (http://fr.diveintopython.org/toc.html). They did the right thing and translated the source XML, so I don’t have to.</p> <p>Upgraded to version 1.43 of the DocBook XSL stylesheets.</p> <p>Upgraded to version 6.43 of the SAXON XSLT processor from Michael Kay.</p> <p>Massive stylesheet changes, moving away from a table-based layout and towards more appropriate use of cascading style sheets. Unfortunately, this broke some of the examples.</p> <p>Moved to Ant (http://jakarta.apache.org/ant/) to have better control over the build process, which is especially important now that I’m using CVS.</p> <p>Consolidated the available downloadable archives; previously, I had different files for each platform, because the .zip files that Python 2.1.1 couldn’t handle.</p> <p>Now hosting the complete XML source, XSL stylesheets, and associated scripts and libraries on SourceForge. There’s also a CVS access URL.</p> <p>Re-licensed the example code under the new-and-improved GPL-compatible Python 2.1.1 license (http://www.python.org/2.1.1/license.html).</p>
Revision 3.5	2001-06-26
	<p>Added explanation of strong/weak/static/dynamic datatypes in Section 2.2, “Declaring Functions”.</p> <p>Added case-sensitivity example in Section 3.1, “Introducing Dictionaries”.</p> <p>Use <code>is</code> in Chapter 5, <i>Objects and Object-Oriented Programming</i> to compensate for inferior operating systems whose files aren’t case-sensitive.</p> <p>Fixed indentation problems in code samples in PDF version.</p>
Revision 3.4	2001-05-31
	<p>Added Section 14.5, “roman.py, stage 5”.</p> <p>Added Section 15.1, “Handling bugs”.</p> <p>Added Section 15.2, “Handling changing requirements”.</p> <p>Added Section 15.3, “Refactoring”.</p> <p>Added Section 15.5, “Summary”.</p> <p>Fixed yet another stylesheet bug that was dropping nested <code>is</code> tags.</p>
Revision 3.3	2001-05-24
	<p>Added Section 13.2, “Diving in”.</p> <p>Added Section 13.3, “Introducing romantest.py”.</p> <p>Added Section 13.4, “Testing for success”.</p> <p>Added Section 13.5, “Testing for failure”.</p> <p>Added Section 13.6, “Testing for sanity”.</p>

<p>Added Section 14.1, “roman.py, stage 1”.</p> <p>Added Section 14.2, “roman.py, stage 2”.</p> <p>Added Section 14.3, “roman.py, stage 3”.</p> <p>Added Section 14.4, “roman.py, stage 4”.</p> <p>Tweaked stylesheets in an endless quest for complete Netscape/Mozilla compatibility.</p>	
Revision 3.2	2001-05-03
<p>Added Section 8.8, “Introducing dialect.py”.</p> <p>Added Section 7.2, “Case Study: Street Addresses”.</p> <p>Fixed bug in <code>method</code> that would produce incorrect declarations (adding a space where it couldn’t be).</p> <p>Fixed bug in CSS (introduced in 2.9) where body background color was missing.</p>	
Revision 3.1	2001-04-18
<p>Added code in <code>no</code> to handle declarations, now that Python 2.1 supports them.</p> <p>Added note about nested scopes in Section 8.5, “locals and globals”.</p> <p>Fixed obscure bug in Example 8.1, “BaseHTMLProcessor.py” where attribute values with character entities would not be properly escaped.</p> <p>Now recommending (but not requiring) Python 2.1, due to its support of declarations in <code>no</code>.</p> <p>Updated download links on the home page (http://diveintopython.org/) to point to Python 2.1, where available.</p> <p>Moved to versioned filenames, to help people who redistribute the book.</p>	
Revision 3.0	2001-04-16
<p>Fixed minor bug in code listing in Chapter 8, <i>HTML Processing</i>.</p> <p>Added link to Chinese translation on home page (http://diveintopython.org/).</p>	
Revision 2.9	2001-04-13
<p>Added Section 8.5, “locals and globals”.</p> <p>Added Section 8.6, “Dictionary-based string formatting”.</p> <p>Tightened code in Chapter 8, <i>HTML Processing</i>, specifically <code>no</code> to use fewer and simpler regular expressions.</p> <p>Fixed a stylesheet bug that was inserting blank pages between chapters in the PDF version.</p> <p>Fixed a script bug that was stripping the <code>no</code> from the home page (http://diveintopython.org/).</p> <p>Added link to Python Cookbook (http://www.activestate.com/ASPN/Python/Cookbook/), and added a few links to individual recipes.</p> <p>Switched to Google (http://www.google.com/services/free.html) for searching on <code>no</code>.</p> <p>Upgraded to version 1.36 of the DocBook XSL stylesheets, which was much more difficult than it sounds. There may still be lingering bugs.</p>	
Revision 2.8	2001-03-26
<p>Added Section 8.3, “Extracting data from HTML documents”.</p> <p>Added Section 8.4, “Introducing BaseHTMLProcessor.py”.</p> <p>Added Section 8.7, “Quoting attribute values”.</p> <p>Tightened up code in Chapter 4, <i>The Power Of Introspection</i>, using the built-in function <code>no</code> instead of manually checking types.</p> <p>Moved Section 5.2, “Importing Modules Using from module import” from Chapter 4, <i>The Power Of Introspection</i> to Chapter 5, <i>Object-Oriented Programming</i>.</p> <p>Fixed typo in code example in Section 5.1, “Diving In” (added colon).</p> <p>Added several additional downloadable example scripts.</p> <p>Added Windows Help output format.</p>	
Revision 2.7	2001-03-16
<p>Added Section 8.2, “Introducing sgmlib.py”.</p> <p>Tightened up code in Chapter 8, <i>HTML Processing</i>.</p> <p>Changed code in Chapter 2, <i>Your First Python Program</i> to use <code>no</code> method instead of <code>no</code>.</p> <p>Moved Section 3.4.2, “Assigning Multiple Values at Once” section to Chapter 2, <i>Your First Python Program</i>.</p> <p>Edited note about <code>no</code> string method, and provided a link to the new entry in <i>The Whole Python FAQ</i> (http://www.python.org/doc/FAQ).</p> <p>Rewrote Section 4.6, “The Peculiar Nature of and and or” to emphasize the fundamental nature of <code>and</code> and <code>or</code> and de-emphasize the <code>and</code> trick.</p> <p>Reorganized language comparisons into <code>no</code>.</p>	
Revision 2.6	2001-02-28
<p>The PDF and Word versions now have colorized examples, an improved table of contents, and properly indented <code>no</code> and <code>no</code>.</p> <p>The Word version is now in native Word format, compatible with Word 97.</p>	

The PDF and text versions now have fewer problems with improperly converted special characters (like trademark symbols and currency symbols).

Added link to download Word version for UNIX, in case some twisted soul wants to import it into StarOffice or something.

Fixed several `#` which were missing titles.

Fixed stylesheets to work around bug in Internet Explorer 5 for Mac OS which caused colorized words in the examples to be displayed in a different color.

Fixed archive corruption in Mac OS downloads.

In first section of each chapter, added link to download examples. (My access logs show that people skim or skip the two pages where the examples are.)

Tightened the home page (<http://diveintopython.org/>) and preface even more, in the hopes that someday someone will read them.

Soon I hope to get back to actually writing this book instead of debugging it.

Revision 2.5	2001-02-23
<p>Added Section 6.4, "Using sys.modules".</p> <p>Added Section 6.5, "Working with Directories".</p> <p>Moved Example 6.17, "Splitting Pathnames" from Section 3.4.2, "Assigning Multiple Values at Once" to Section 6.5, "Working with Directories".</p> <p>Added Section 6.6, "Putting It All Together".</p> <p>Added Section 5.10, "Summary".</p> <p>Added Section 8.1, "Diving in".</p> <p>Fixed program listing in Example 6.10, "Iterating Through a Dictionary" which was missing a colon.</p>	
Revision 2.4.1	2001-02-12
<p>Changed newsgroup links to use "news:" protocol, now that <code>#</code> is defunct.</p> <p>Added file sizes to download links.</p>	
Revision 2.4	2001-02-12
<p>Added "further reading" links in most sections, and collated them in Appendix A, <i>Further reading</i>.</p> <p>Added URLs in parentheses next to external links in text version.</p>	
Revision 2.3	2001-02-09
<p>Rewrote some of the code in Chapter 5, <i>Objects and Object-Oriented Programming</i> to use class attributes and a better example of multi-variables.</p> <p>Reorganized Chapter 5, <i>Objects and Object-Oriented Programming</i> to put the class sections first.</p> <p>Added Section 5.8, "Introducing Class Attributes".</p> <p>Added Section 6.1, "Handling Exceptions".</p> <p>Added Section 6.2, "Working with File Objects".</p> <p>Merged the "review" section in Chapter 5, <i>Objects and Object-Oriented Programming</i> into Section 5.1, "Diving In".</p> <p>Colorized all program listings and examples.</p> <p>Fixed important error in Section 2.2, "Declaring Functions": functions that do not explicitly return a value return <code>None</code> so you can assign them to variables.</p> <p>Added minor clarifications to Section 2.3, "Documenting Functions", Section 2.4, "Everything Is an Object", and Section 3.4, "Declaring Functions".</p>	
Revision 2.2	2001-02-02
<p>Edited Section 4.4, "Getting Object References With getattr".</p> <p>Added titles to <code>#</code> tags, so they can have their cute little tooltips too.</p> <p>Changed the look of the revision history page.</p> <p>Fixed problem I introduced yesterday in my HTML post-processing script that was causing invalid HTML character references and broken links.</p> <p>Upgraded to version 1.29 of the DocBook XSL stylesheets.</p>	
Revision 2.1	2001-02-01
<p>Rewrote the example code of Chapter 4, <i>The Power Of Introspection</i> to use <code>__dict__</code> instead of <code>__dict__</code> and <code>__dict__</code> and rewrote explanatory text to match.</p> <p>Added example of list operators in Section 3.2, "Introducing Lists".</p> <p>Added links to relevant sections in the summary lists at the end of each chapter (Section 3.8, "Summary" and Section 4.9, "Summary").</p>	
Revision 2.0	2001-01-31
<p>Split Section 5.6, "Special Class Methods" into three sections, Section 5.5, "Exploring UserDict: A Wrapper Class", Section 5.6, "Special Class Methods", and Section 5.7, "Special Class Methods".</p> <p>Changed notes on garbage collection to point out that Python 2.0 and later can handle circular references without additional coding.</p> <p>Fixed UNIX downloads to include all relevant files.</p>	
Revision 1.9	2001-01-15
<p>Removed introduction to Chapter 2, <i>Your First Python Program</i>.</p> <p>Removed introduction to Chapter 4, <i>The Power Of Introspection</i>.</p>	

	Removed introduction to Chapter 5, <i>Objects and Object-Orientation</i> . Edited text ruthlessly. I tend to ramble.
Revision 1.8	2001-01-12
	Added more examples to Section 3.4.2, "Assigning Multiple Values at Once". Added Section 5.3, "Defining Classes". Added Section 5.4, "Instantiating Classes". Added Section 5.6, "Special Class Methods". More minor stylesheet tweaks, including adding titles to <code>h</code> ags, which, if your browser is cool enough, will display a description of the
Revision 1.71	2001-01-03
	Made several modifications to stylesheets to improve browser compatibility.
Revision 1.7	2001-01-02
	Added introduction to Chapter 2, <i>Your First Python Program</i> . Added introduction to Chapter 4, <i>The Power Of Introspection</i> . Added review section to Chapter 5, <i>Objects and Object-Orientation</i> [later removed] Added Section 5.9, "Private Functions". Added Section 6.3, "Iterating with for Loops". Added Section 3.4.2, "Assigning Multiple Values at Once". Wrote scripts to convert book to new output formats: one single HTML file, PDF, Microsoft Word 97, and plain text. Registered the domain and moved the book there, along with links to download the book in all available output formats for offline use. Modified the XSL stylesheets to change the header and footer navigation that displays on each page. The top of each page is branded
Revision 1.6	2000-12-11
	Added Section 4.8, "Putting It All Together". Finished Chapter 4, <i>The Power Of Introspection</i> with Section 4.9, "Summary". Started Chapter 5, <i>Objects and Object-Orientation</i> with Section 5.1, "Diving In".
Revision 1.5	2000-11-22
	Added Section 4.6, "The Peculiar Nature of and and or". Added Section 4.7, "Using lambda Functions". Added appendix that lists section abstracts. Added appendix that lists tips. Added appendix that lists examples. Added appendix that lists revision history. Expanded example of mapping lists in Section 3.6, "Mapping Lists". Encapsulated several more common phrases into entities. Upgraded to version 1.25 of the DocBook XSL stylesheets.
Revision 1.4	2000-11-14
	Added Section 4.5, "Filtering Lists". Added documentation to Section 4.3, "Using type, str, dir, and Other Built-In Functions". Added example in Section 3.3, "Introducing Tuples". Added additional note about trick under MacPython. Switched to the SAXON XSLT processor from Michael Kay. Upgraded to version 1.24 of the DocBook XSL stylesheets. Added db-html processing instructions with explicit filenames of each chapter and section, to allow deep links to content even if I a Made several common phrases into entities for easier reuse. Changed several <code>h</code> ags to <code>h</code>
Revision 1.3	2000-11-09
	Added section on dynamic code execution. Added links to relevant section/example wherever I refer to previously covered concepts. Expanded introduction of chapter 2 to explain what the function actually does. Explicitly placed example code under the GNU General Public License and added appendix to display license. [Note 8/16/2001: co

Changed links to licenses to use flags, now that I know how to use them.	
Revision 1.2	2000-11-06
<p>Added first four sections of chapter 2.</p> <p>Tightened up preface even more, and added link to Mac OS version of Python.</p> <p>Filled out examples in "Mapping lists" and "Joining strings" to show logical progression.</p> <p>Added output in chapter 1 summary.</p>	
Revision 1.1	2000-10-31
<p>Finished chapter 1 with sections on mapping and joining, and a chapter summary.</p> <p>Toned down the preface, added links to introductions for non-programmers.</p> <p>Fixed several typos.</p>	
Revision 1.0	2000-10-30
Initial publication.	

Appendix F. About the book

This book was written in DocBook XML (<http://www.oasis-open.org/docbook/>) using Emacs (<http://www.gnu.org/software/emacs/>), and con

If you're interested in learning more about DocBook for technical writing, you can download the XML source (<http://diveintopython.org/dow>

Appendix G. GNU Free Documentation License

Version 1.1, March 2000

Copyright (C) 2000 Free Software Foundation, Inc. 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA Everyone is permitted to copy and distribute verbatim copies of this license document, provided the copyright notice and this notice are preserved on all copies.

G.0. Preamble

The purpose of this License is to make a manual, textbook, or other written document "free" in the sense of freedom: to assure everyone the right to copy and distribute the document, unaltered and in quantity, to use the document in any medium, and to revise and update the document.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals written in Free Documentation License format.

G.1. Applicability and definitions

This License applies to any manual or other work that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers of the Document to the public and is separated from the rest of the Document by a title page.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for automatic processing.

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML input format, and any other format designed to be processed by machines.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the Title Page.

G.2. Verbatim copying

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notice, and the Title Page are preserved on all copies.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

G.3. Copying in quantity

If you publish printed copies of the Document numbering more than 100, and the Document's license notice requires Cover Texts, you must include the full Texts on all covers of each copy.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and the remainder on adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in the Title Page that you have prepared Transparent copies for those who can be contacted through the address listed in the Title Page.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to correct any inaccuracies in the Document.

G.4. Modifications

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release it under this License, that you add a new title page, that you change the old title page to read "This Document is derived from GNU Free Documentation License, Version 1.1, March 2000," and that you delete the notice in the Title Page referring to the old version of the License.

Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should be given in parentheses), and from the title of any other document intended to be distributed together with the Modified Version.

List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, including those responsible for releasing it under this License.

State on the Title page the name of the publisher of the Modified Version, as the publisher.3.

Preserve all the copyright notices of the Document.4.

Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.5.

Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the

Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.7.

Include an unaltered copy of this License.8.

Preserve the section entitled "History", and its title, and add to it an item stating at least the title, year, new authors, and publisher of

Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise

In any section entitled "Acknowledgements" or "Dedications", preserve the section's title, and preserve in the section all the substan

Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are n

Delete any section entitled "Endorsements". Such a section may not be included in the Modified Version.13.

Do not retitl e any existing section as "Endorsements" or to conflict in title with any Invariant Section.14.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied f

You may add a section entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties-

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the li

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imp

G.5. Combining documents

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy

In the combination, you must combine any sections entitled "History" in the various original documents, forming one section entitled "Histor

G.6. Collections of documents

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies o

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of thi

G.7. Aggregation with independent works

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one quarter of the

G.8. Translation

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing

G.9. Termination

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to

G.10. Future revisions of this license

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new ver

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this Lic

G.11. How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notice in the document.

Copyright (c) YEAR YOUR NAME. Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with the Invariant Sections listed below.

If you have no Invariant Sections, write "with no Invariant Sections" instead of saying which ones are invariant. If you have no Front-Cover Text, write "no Front-Cover Text" instead of "no Front-Cover Text required". If you have no Back-Cover Text, write "no Back-Cover Text" instead of "no Back-Cover Text required".

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license (e.g., GNU GPL or GNU LGPL) in order to allow their use in free software. For more information on this topic, see the GNU Free Documentation License, section 7.

Appendix H. Python license

H.A. History of the software

Python was created in the early 1990s by Guido van Rossum at Stichting Mathematisch Centrum (CWI) in the Netherlands as a successor of

Following the release of Python 1.6, and after Guido van Rossum left CNRI to work with commercial software developers, it became clear th

After Python 2.0 was released by BeOpen.com, Guido van Rossum and the other PythonLabs developers joined Digital Creations. All intelle

Thanks to the many outside volunteers who have worked under Guido's direction to make these releases possible.

H.B. Terms and conditions for accessing or otherwise using Python

H.B.1. PSF license agreement

This LICENSE AGREEMENT is between the Python Software Foundation ("PSF"), and the Individual or Organization ("Licensee")
Subject to the terms and conditions of this License Agreement, PSF hereby grants Licensee a nonexclusive, royalty-free, world-wide
In the event Licensee prepares a derivative work that is based on or incorporates Python 2.1.1 or any part thereof, and wants to make
PSF is making Python 2.1.1 available to Licensee on an "AS IS" basis. PSF MAKES NO REPRESENTATIONS OR WARRANTIES
PSF SHALL NOT BE LIABLE TO LICENSEE OR ANY OTHER USERS OF PYTHON 2.1.1 FOR ANY INCIDENTAL, SPECIAL, OR CONSEQUENTIAL DAMAGES OR LOSS.
This License Agreement will automatically terminate upon a material breach of its terms and conditions.
Nothing in this License Agreement shall be deemed to create any relationship of agency, partnership, or joint venture between PSF and
By copying, installing or otherwise using Python 2.1.1, Licensee agrees to be bound by the terms and conditions of this License Agreement.

H.B.2. BeOpen Python open source license agreement version 1

This LICENSE AGREEMENT is between BeOpen.com ("BeOpen"), having an office at 160 Saratoga Avenue, Santa Clara, CA 95051
Subject to the terms and conditions of this BeOpen Python License Agreement, BeOpen hereby grants Licensee a non-exclusive, royalty-free, world-wide
BeOpen is making the Software available to Licensee on an "AS IS" basis. BEOPEN MAKES NO REPRESENTATIONS OR WARRANTIES
BEOPEN SHALL NOT BE LIABLE TO LICENSEE OR ANY OTHER USERS OF THE SOFTWARE FOR ANY INCIDENTAL, SPECIAL, OR CONSEQUENTIAL DAMAGES OR LOSS.
This License Agreement will automatically terminate upon a material breach of its terms and conditions.
This License Agreement shall be governed by and interpreted in all respects by the law of the State of California, excluding conflict of laws provisions.
By copying, installing or otherwise using the software, Licensee agrees to be bound by the terms and conditions of this License Agreement.

H.B.3. CNRI open source GPL-compatible license agreement

This LICENSE AGREEMENT is between the Corporation for National Research Initiatives, having an office at 1895 Preston White Drive
Subject to the terms and conditions of this License Agreement, CNRI hereby grants Licensee a nonexclusive, royalty-free, world-wide
In the event Licensee prepares a derivative work that is based on or incorporates Python 1.6.1 or any part thereof, and wants to make
CNRI is making Python 1.6.1 available to Licensee on an "AS IS" basis. CNRI MAKES NO REPRESENTATIONS OR WARRANTIES
CNRI SHALL NOT BE LIABLE TO LICENSEE OR ANY OTHER USERS OF PYTHON 1.6.1 FOR ANY INCIDENTAL, SPECIAL, OR CONSEQUENTIAL DAMAGES OR LOSS.
This License Agreement will automatically terminate upon a material breach of its terms and conditions.
This License Agreement shall be governed by the federal intellectual property law of the United States, including without limitation the Copyright Act of 1976.
By clicking on the "ACCEPT" button where indicated, or by copying, installing or otherwise using Python 1.6.1, Licensee agrees to be bound by the terms and conditions of this License Agreement.

H.B.4. CWI permissions statement and disclaimer

Copyright (c) 1991 - 1995, Stichting Mathematisch Centrum Amsterdam, The Netherlands. All rights reserved.

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that

STICHTING MATHEMATISCH CENTRUM DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.