QRNA-2.0.1 documentation

Elena Rivas and Sean R. Eddy

{elena,eddy}@genetics.wustl.edu.

Howard Hughes Medical Institute
Department of Genetics
Washington University, St. Louis, MO 63110, USA

Abstract

Basic how–to guide to install and run QRNA-2.0.1.

QRNA-2.0.1 is a *very early* prototype of a new QRNA generation that uses probabilistic evolutionary models to be able to tune all the parameters (transition and emission probabilities) of the three QRNA models to any possible degree of sequence divergence.

Install QRNA

```
Source distribution (qrna-2.0.1.tar.Z):

    setenv QRNADB  (HOME)/qrna-2.0.1/lib  assign a location for the libraries.

    tar zxf qrna-2.0.1.tar.Zr      Unpacks the archive. (makes a new directory, qrna-2.0.1).
    cd qrna-2.0.1                  Moves into the distribution toplevel directory.
    cd squid                       Moves into the directory squid.
    make                           Builds the binaries for the squid library.
    cd ..                          Goes back to toplevel directory.
    cd src                         Moves into the source directory.
    make                           Builds the binaries.
```

It should build cleanly on just about any UNIX machine.

The directory **qrna-2.0.1** includes the following subdirectories,

Demos A collection of files to demonstrate how to use QRNA.

documentation Contains a userguide manual.

lib Contains all the additional files used as input by QRNA. This directory is accessed by setting the environment variable QRNADB.

Licenses The license for this software.

scripts The wrap around Perl scripts used with QRNA.

squid An old version of Sean Eddy's squid library for sequence handling.

src The source files, and headers. It also includes a Makefile. After building QRNA the executable is located in this directory.

Create a QRNA input file from a blast output

- Start with a WUBLASTN output file **foo.blast**.
  Example, the file in the 'qrna-2.0.1/Demos/'' directory **HG_13_RNAs_gene.fa.MGSCv3.fragchrom.blast**, which is the result of blasting a collection of 850 annotated Human RNA genes against the mouse genome.

- Then you need to run the Perl scripts **qrna-2.0.1/scripts/blastn2qrnadepth.pl**.

```
usage: blastn2qrnadepth.pl [options] <blastfile>
options:
sage: blastn2qrnadepth.pl [options] <blastfile>
options:
-d <depth>      :  max number of alignment coverage per position  [ default depth = 1       ]
-e <max_eval>   :  maximum evalue   of blast hits allowed          [ default max_eval = 0.01 ]
-g <org_name>   :  name of the blasting organism                   [ default 'org'           ]
-i <min_id>     :  minimum identity of blast hits allowed          [ default min_id = 0      ]
-j <max_id>     :  maximum identity of blast hits allowed          [ default max_id = 100    ]
-l <min_len>    :  minimum length  of blast hits allowed           [ default min_len = 1     ]
-o <outfile>    :  output qfile                                    [ default = 'blastfile.q' ]
-r              :  calculate depth respect to the database instead of the query
-s <shift>      :  position shift when calculating depth           [ default shift = 1       ]
-w <which>      :  criteria to pick alignments                     [ default which = 'SC'    ]
                                ID -- best % identity
                                SC -- best score
```

- "depth" is a parameter to limit the number of blast alignments that include a given position. If you do not want to constrain the number of blast alignments this way, set "depth" to a very large value. Limiting the depth is important for eukariotic genomes with large amount of repetitive sequences. The default depth = 1 will give the best alignment per position. If you want to allow some redundancy a depth of 5 or 10 is recommended.
- "max_eval" removes blastn alignments with E values above this cutoff.
- "org_name" is the name of the organism as it will appear in the gff file
- "min_id" is the minimum identity of blastn alignments allowed.
- "max_id" is the maximum identity of blastn alignments allowed.
- "min_len" is the minimum length of blastn alignments allowed.
- "shift" is a parameter related to "depth". It indicates how many nucleotides to skip after every check for depth. I recommend not to change it from its default value.
- "which" has to values "SC" or "ID". It determines the criteria to pick the best alignments when applying the depth constraint.

Notice: Even if you do not want to prune **foo.blast**, you have to run this script.

- The result of running
  qrna-2.0.1/scripts/blastn2qrnadepth.pl   foo.blast
  is the three files: **foo.blast.q**, **foo.blast.q.gff**, and **foo.blast.q.rep**

  - **foo.blast.q** is the input file taken by QRNA.
    It is a collection of sequences in fasta format, where two consecutive sequences are the two component of an alignment with the gaps left in place.
  - **foo.blast.q.rep** is a report of the BLASTN alignment that have been pruned in the process of creating **foo.blast.q** according to the options used in **blastn2qrnadepth.pl**.
  - **foo.blast.q.gff** is the gff format version of the surviving aligned regions of the query genome.

- The result of running
  qrna-2.0.1/scripts/blastn2qrnadepth.pl -g human   HG_13_RNAs_gene.fa.MGSCv3.fragchrom.blast
  is the three files in the directory "qrna-2.0.1/Demos/":
  HG_13_RNAs_gene.fa.MGSCv3.fragchrom.blast.q
  HG_13_RNAs_gene.fa.MGSCv3.fragchrom.blast.q.rep
  HG_13_RNAs_gene.fa.MGSCv3.fragchrom.blast.q.gff
  Looking at the .rep file, we observe that initially the .blast file had $69,553$ alignments. And that after using the default filters, we end up with 867 alignments approximately one alignment per RNA gene which is reasonable since by default we have assumed depth = 1.

Run QRNA

- QRNA is a collection of programs, written in C.
  Source code is in "qrna-2.0.1/src/". The executable is called "qrna".

- The "makefile" at "qrna-2.0.1/src/" can be modified to change the C compiler or to change compilation flags.

- To run qrna :
  qrna-2.0.1/src/qrna   [options]   foo.blast.q

- Options are:

```
        qrna -- scores an alignment with the 3 models
                2.0.1 (Tue Aug 19 11:30:55 CDT 2003) using squid 1.5m (Sept 1997)
Usage: qrna [-options] <input_file.q>
where options are:
    -A                : do an all-to-all comparison between the two input files
    -a                : print alignment
    -B                : sre_shuffle the alignment keeping the gap structure of the window
    -b                : shuffle the alignment
    -c <cfgfile>      : <cfgfile> to use to train the rna model (default = tRNA+rRNA)
    -D <codfile>      : include a file of coding-coding frequencies for the coding model
    -d                : log2 form (default = log2-odds space )
    -e <num>          : number of sequen skipped in second file (for multiple comparisons). default 0.
    -F                : change the overall base composition of the 3 models, based on nts frequencies in the
    -f                : use full dp for the probabilistic models (do not conserve the aligment-default is di
    -G                : change the overall base composition of the 3 models, based on nts frequencies for ea
    -g                : do forward (default is viterbi)
    -H <Hexfile>      : include a file of Hexamer frequencies for the coding model
    -h                : print short help and usage info
    -i                : evolutionary time factor (default i=1)
    -j                : use semi-full dp for the probabilistic model (use the alignment created by OTH)
    -k                : allow pseudoknots (not implemented)
    -l <minlenhit>    : change the minlenhit parameter (default 0)
    -L <maxlenhit>    : change the maxlenhit parameter (default provided by longuest sequence)
    -m                : do Forward and Viterbi Diagonal dp
    -n                : do Forward and Viterbi Full    dp
    -o <outfile>      : direct structure-annotated sequence to <outfile>
    -p <pamfile>      : <pamfault> to use (default = BLOSUM62)
    -P                : pedantic, check your evolutionary models for inconsistencies
    -q                : do Forward and Viterbi semi-full dp
    -r                : do Nussinov rna model (default is a 3-state model)
    -R <ribofile>     : <ribofile> to use to train the rna model (default = RIBOPROB85-60)
    -s                : do global (not dp)
    -S                : sweep a collection of motifs(seqfile1) across another bunch of sequences(seqfile2)
    -t                : print traceback
    -v                : verbose debugging output
    -w <num>          : scanning window (default is full length)
    -x <num>          : slide positions (default is 50)
    -y <num>          : grab n sequences at random from the second data file to compare to each one of the


    --cyk     : use CYK algorithm to calculate RNA score (default is Inside).
    --latte   : I just called starbucks with your order...
    --ones    : score with the three models only the given strand.
    --parse   : input is a selex file. Por a given ss, it calculates the probablity of either the best or
    --rnass   : print the alignment with the predicted RNA secondary structure.
    --noends  : do not evaluate ends. Default now: calculate the actual boundaries of the model calls
    --shtoo   : qrna the alignment and also give one shuffled score.

 Debugging, experimentation:
    --regress <f> : save regression test information to file <f>
```

Some useful flags:

- For alignments that are too long, you can score them in chunks using -w ⟨windowsize⟩. The option -x ⟨slide⟩ to decide how many nucleotides to move before you score another chunk of the given alignment. Every window analyzed starts with "length alignment:".

- Option -L determines the maximum length allowed for an alignment to be scored. The default is 1,000. I use this variable to control the memory usage of the program. If you are using the "window" version, memory is determined by the window so you can set -L as large as you want. If not using a window, I would not use a max length larger than 1,500.

- There are three different scoring algorithms:

  global (-s)

  local Viterbi (default)

  local forward (-g).

  You can report any of them or any desired combination of them. I would recommend to use the default.

Looking at a QRNA output:

For file "qrna-2.0.1/Demos/5s_rRNA.q" which contains an alignment of two 5S rRNAs we have,

qrna-2.0.1/src/qrna -a qrna-2.0.1/Demos/5s_rRNA.q > qrna-2.0.1/Demos/5s_rRNA.q.qrna2

The output looks like:

```
#----------------------------------------------------------------------------
#       qrna 2.0.1 (Tue Aug 19 11:30:55 CDT 2003) using squid 1.5m (Sept 1997)
#----------------------------------------------------------------------------
#       PAM model =  BLOSUM62
#----------------------------------------------------------------------------
#       RNA model       =  /mix_tied_linux.cfg
#       RIBOPROB matrix =  /RIBOPROB85-60.mat
#----------------------------------------------------------------------------
#       seq file  =  ../Demos/5s_rRNA.q
#                    #seqs: 2 (max_len = 143)
#----------------------------------------------------------------------------
#       scanning version: full length -- length range = [0,1000]
#----------------------------------------------------------------------------
# 1  [both strands]
>RF00001.5S_rRNA.U26684/296-411 (143)
>RF00001.5S_rRNA.M35563/5-120 (143)

Divergence time (variable): 0.109
[alignment ID = 82.05 MUT = 16.24 GAP = 1.71]

length alignment: 117 (id=82.05) (mut=16.24) (gap=1.71)
posX: 0-116 [0-115](116) -- (0.27 0.24 0.28 0.22)
posY: 0-116 [0-115](116) -- (0.23 0.29 0.29 0.18)


        RF00001.5S_rRNA CCUGAUACCCAUAGAGCUGUGGUACCACCUGAAUCCAUGCCGAACUCAGA
        RF00001.5S_rRNA CCUGGUGUCCAUAGAGCACUGGAACCACCUGAUCCCAUCCCGAACUCAGA

        RF00001.5S_rRNA AGUGAAACGCAGCAUCGCCGAUGGUAGUGUGAG.GUCUCCUCAUGUGAGA
        RF00001.5S_rRNA AGUGAAACGGUGCAUCGCCGAUGGUAGUGUG.GGGCCUCCCCAUGUGAGA

        RF00001.5S_rRNA GUAGGACAGUAUCAGGU
        RF00001.5S_rRNA GUAGGUCAACGCCAGGC


LOCAL_DIAG_VITERBI -- [Inside SCFG]
OTH ends  (+) =  (0..[117]..116)
OTH ends *(-) =  (0..[117]..116)
COD ends  (+) =  (80..[4]..83)
```

4

```
COD ends *(-) =  (60..[25]..84)
RNA ends  (+) =  (0..[117]..116)
RNA ends *(-) =  (0..[117]..116)
winner = RNA
            OTH =     245.255           COD =     244.061           RNA =     252.457
   logoddspostOTH =       0.000  logoddspostCOD =      -1.193  logoddspostRNA =       7.202
     sigmoidalOTH =      -7.206    sigmoidalCOD =      -8.405    sigmoidalRNA =       6.679
```

- Every new blast alignment starts with two lines:
  ">Query_name"
  ">Subject_name"

- "Divergence time" indicates the particular time parameterization of QRNA used. By default QRNA decides on the divergence time (in this case $t = 0.109$) given the percentage identity of the alignment (82.05%) [Divergence time (variable)]. You can set a particular divergence time using the option "-i" [Divergence time (fixed)].

- For each model, and for each strand, you are given the actual local regions (they could be more than one per model and strand) that score according to the model. The notation is (from..[length]..to). Coordinates for *both* strands are given relative to the positive strand. The "*" indicates the strand with the strongest signal for a given model.

- For a given scoring algorithm you get three rows of numbers:

  row 1 The scores of the alignment under each of the three models. The "null" model is a forth model which assumes that the two sequences in the alignment are independent from each other.

  $$\log \frac{P(data|OTH)}{P(data|null)} \qquad \log \frac{P(data|COD)}{P(data|null)} \qquad \log \frac{P(data|RNA)}{P(data|null)}$$

  row 2 The two (COD and RNA) log–odds posterior probabilities respect to the OTH model.

  $$\log \frac{P(OTH|data)}{P(OTH|data)} \qquad \log \frac{P(COD|data)}{P(OTH|data)} \qquad \log \frac{P(RNA|data)}{P(OTH|data)}$$

  row 3 The row you should be paying the most attention to. The three sigmoidal scores calculated using the other two models as null models. The model with the highest sigmoidal score is the winner.

  $$\sigma_{\text{OTH}} = \log \frac{P(data|OTH)P(OTH)}{P(data|COD) + P(data|RNA)} \quad \sigma_{\text{COD}} = \log \frac{P(data|COD)P(COD)}{P(data|OTH) + P(data|RNA)} \quad \sigma_{\text{RNA}} = \log \frac{P(data|RNA)P(RNA)}{P(data|OTH) + P(data|COD)}$$

  These scores are called sigmoidal because for each model M_i

  $$P(M_i|data) = \frac{e^{\sigma_{M_i}}}{1 + e^{\sigma_{M_i}}},$$

  which is a sigmoidal function. This function has the nice property that

  $$P(M_i|data) > 1/2 \Longleftrightarrow \sigma_{M_i} > 0.$$

  We assume a flat prior for the three models.

- Option -a prints the scored alignment.

- Option -B scrambles the columns of the pairwise alignment. The result is an alignment with the same percentage identity, but without any column correlations. Compare the results of using -B for the alignment of two 5S rRNAs given before:

  qrna-2.0.1/src/qrna -a -B qrna-2.0.1/Demos/5s_rRNA.q > qrna2-.0.0/Demos/5s_rRNA.q.shuffle.qrna2

```
#-------------------------------------------------------------------------------------
#      qrna 2.0.1 (Tue Aug 19 11:30:55 CDT 2003) using squid 1.5m (Sept 1997)
#-------------------------------------------------------------------------------------
#      PAM model =  BLOSUM62
#-------------------------------------------------------------------------------------
#      RNA model       =  /mix_tied_linux.cfg
#      RIBOPROB matrix =  /RIBOPROB85-60.mat
#-------------------------------------------------------------------------------------
#      seq file  =  ../Demos/5s_rRNA.q
#                   #seqs: 2 (max_len = 143)
#-------------------------------------------------------------------------------------
#      scanning version: full length -- length range = [0,1000]
#-------------------------------------------------------------------------------------
# 1   [both strands] (sre_shuffled)
>RF00001.5S_rRNA.U26684/296-411 (143)
>RF00001.5S_rRNA.M35563/5-120 (143)


Divergence time (variable): 0.109
[alignment ID = 82.05 MUT = 16.24 GAP = 1.71]

length alignment: 117 (id=82.05) (mut=16.24) (gap=1.71)
posX: 0-116 [0-115](116) -- (0.27 0.24 0.28 0.22)
posY: 0-116 [0-115](116) -- (0.23 0.29 0.29 0.18)



        RF00001.5S_rRNA UCCGGUUAGGGCUAAAUAGUUUCGAAACUACUGGCUCGGGGGCGCUUAAU
        RF00001.5S_rRNA UCCGGUCAGCGCUGAAUUGUUUCGAAACUACUGGCUCGGGGGCCCUCGAU

        RF00001.5S_rRNA CCCACCCAUUCAAAGAACGGACAGGGACUUGAC.ACAGAUUGGCUAGGGA
        RF00001.5S_rRNA CCUACCCAUUCAAGAAUGGGACAGGGACAUG.CGACAGAUCGGCCAGGGA

        RF00001.5S_rRNA CACGCUCGAUAUGUACG
        RF00001.5S_rRNA CACGCCCGUUACGAACG

LOCAL_DIAG_VITERBI -- [Inside SCFG]
OTH ends  (+) =  (0..[117]..116)
OTH ends *(-) =  (0..[117]..116)
COD ends  (+) =  (100..[12]..111)
COD ends *(-) =  (1..[11]..11)
RNA ends  (+) =  (106..[11]..116)
RNA ends *(-) =  (106..[11]..116)
winner = OTH
            OTH =      245.255         COD =    234.758        RNA =     240.633
  logoddspostOTH =        0.000  logoddspostCOD =   -10.497  logoddspostRNA =    -4.621
    sigmoidalOTH =        4.597    sigmoidalCOD =   -10.555    sigmoidalRNA =    -4.622
```

The alignment still has 82.05% identity as the original one, but now it gets classified as "other" since the RNA structure has been destroyed. We use the -B option to estimate false positives.

Example using the scanning version with a window:

Consider file "qrna-2.0.1/Demos/Scerevisiae_orf_v_other_yeasts.q" which contains an alignment of a *S. cerevisiae* ORF The alignment has 514 nucleotides, and we would like to score it with QRNA using a window of 150 nucleotides, and moving the window 50 nucleotides each time.

qrna-2.0.1/src/qrna -w 150 -x 50 qrna-2.0.1/Demos/Scerevisiae_orf_v_other_yeasts.q >
Scerevisiae_orf_v_other_yeasts.q.qrna2

The output for the first two windows looks like:

```
#------------------------------------------------------------------------------
#      qrna 2.0.1 (Tue Aug 19 11:30:55 CDT 2003) using squid 1.5m (Sept 1997)
#------------------------------------------------------------------------------
#      PAM model =  BLOSUM62
#------------------------------------------------------------------------------
#      RNA model       =  /mix_tied_linux.cfg
#      RIBOPROB matrix =  /RIBOPROB85-60.mat
#------------------------------------------------------------------------------
#      seq file  =  ../Demos/Scerevisiae_orf_v_other_yeasts.q
#                   #seqs: 2 (max_len = 545)
#------------------------------------------------------------------------------
#      window version: window = 150   slide = 50 -- length range = [0,9999999]
#------------------------------------------------------------------------------
# 1   [both strands]
>EFB1_I-142172-143158-80<621- (545)
>Contig1207-5-454>998- (545)

length of whole alignment after removing common gaps: 545
Divergence time (variable): 0.023
[alignment ID = 92.66 MUT = 6.79 GAP = 0.55]

length alignment: 150 (id=96.67) (mut=3.33) (gap=0.00)
posX: 0-149 [0-149](150) -- (0.29 0.20 0.18 0.33)
posY: 0-149 [0-149](150) -- (0.27 0.21 0.20 0.32)
LOCAL_DIAG_VITERBI -- [Inside SCFG]
OTH ends *(+) =  (0..[150]..149)
OTH ends  (-) =  (0..[150]..149)
COD ends  (+) =  (78..[72]..149)
COD ends *(-) =  (6..[144]..149)
RNA ends *(+) =  (0..[150]..149)
RNA ends  (-) =  (0..[150]..149)
winner = COD
              OTH =      437.721          COD =      445.056         RNA =      433.639
   logoddspostOTH =        0.000 logoddspostCOD =        7.335 logoddspostRNA =       -4.082
     sigmoidalOTH =       -7.336   sigmoidalCOD =        7.252   sigmoidalRNA =      -11.426

length alignment: 150 (id=94.67) (mut=5.33) (gap=0.00)
posX: 50-199 [50-199](150) -- (0.27 0.25 0.18 0.30)
posY: 50-199 [50-199](150) -- (0.27 0.24 0.21 0.28)
LOCAL_DIAG_VITERBI -- [Inside SCFG]
OTH ends *(+) =  (50..[150]..199)
OTH ends  (-) =  (50..[150]..199)
COD ends  (+) =  (78..[72]..149)
COD ends *(-) =  (54..[144]..197)
RNA ends  (+) =
RNA ends *(-) =  (50..[150]..199)
winner = COD
              OTH =      416.785          COD =      422.932         RNA =      411.383
   logoddspostOTH =        0.000 logoddspostCOD =        6.148 logoddspostRNA =       -5.402
     sigmoidalOTH =       -6.148   sigmoidalCOD =        6.114   sigmoidalRNA =      -11.570
```

How to read the information for each analyzed window:

- Each new analyzed window starts with the line:

  ```
  length alignment:
  ```

- For each window and for each sequence in the alignment we have a line of the form:

  ```
  posX: 50-199 [50-199](150) -- (0.27 0.25 0.18 0.30)
  ```

  The first pair of number represent the first and last coordinates of the window respect to the beginning of the alignment. The pair of numbers in brackets represent the mapping of the window into the coordinate system of sequence X (after removing gaps). The adjacent number in parenthesis is the length of that segment in sequence X. Finally the four decimal numbers in parenthesis are the fraction of A, C, G, and T's in the segment of sequence X involved in that particular window.

Example starting with a blastn output:

File "**qrna-2.0.1/Demos/HG_13_RNAs_gene.fa.MGSCv3.fragchrom.blast**" is a typical WUBLASTN output file. This file is the result of blasting a collection of 850 annotated Human ncRNAs to the mouse genome.

Typing:
    qrna-2.0.1/scripts/blastn2qrnadepth.pl -g human
                   qrna-2.0.1/Demos/HG_13_RNAs_gene.fa.MGSCv3.fragchrom.blast
creates the file
    "**qrna-2.0.1/Demos/HG_13_RNAs_gene.fa.MGSCv3.fragchrom.blast.q**".
This file has selected those WUBLASTN alignment of any length, and any % identity, and E-value ≤ 0.01, and with depth = 1.
The two aligned sequence are now ready to be sent to QRNA.

Typing:
    qrna-2.0.1/src/qrna -w 150 -x 50 /qrna-2.0.1/Demo/HG_13_RNAs_gene.fa.MGSCv3.fragchrom.blast.q >
                                    qrna-2.0.1/Demo/HG_13_RNAs_gene.fa.MGSCv3.fragchrom.blast.q.W150.X50.qrna2

produces an output that starts with:

```
#----------------------------------------------------------------------------------
#      qrna 2.0.1 (Tue Aug 19 11:30:55 CDT 2003) using squid 1.5m (Sept 1997)
#----------------------------------------------------------------------------------
#      PAM model =  BLOSUM62
#----------------------------------------------------------------------------------
#      RNA model       =  /mix_tied_linux.cfg
#      RIBOPROB matrix =  /RIBOPROB85-60.mat
#----------------------------------------------------------------------------------
#      seq file  =  ../Demos/HG_13_RNAs_gene.fa.MGSCv3.fragchrom.blast.q
#                   #seqs: 1734 (max_len = 358)
#----------------------------------------------------------------------------------
#      window version: window = 150   slide = 50 -- length range = [0,9999999]
#----------------------------------------------------------------------------------
# 1  [both strands]
>NT_034563-2555387-2555837-71>406-Telomerase_RNA (344)
>chr3-frag968-96700001-96801000-15935>16265- (344)

length of whole alignment after removing common gaps: 344
Divergence time (variable): 0.272
[alignment ID = 69.48 MUT = 24.42 GAP = 6.10]

length alignment: 150 (id=72.00) (mut=20.00) (gap=8.00)
```

```
posX: 0-149 [0-146](147) -- (0.13 0.39 0.34 0.14)
posY: 0-149 [0-140](141) -- (0.11 0.40 0.35 0.13)
LOCAL_DIAG_VITERBI -- [Inside SCFG]
OTH ends *(+) =  (0..[150]..149)
OTH ends  (-) =  (0..[150]..149)
COD ends *(+) =  (87..[12]..98)
COD ends  (-) =  (76..[9]..84)
RNA ends  (+) =  (84..[45]..128)
RNA ends *(-) =  (48..[102]..149)
winner = RNA
              OTH =     224.740              COD =     207.635              RNA =     229.595
   logoddspostOTH =       0.000  logoddspostCOD =     -17.106  logoddspostRNA =       4.854
     sigmoidalOTH =      -4.854    sigmoidalCOD =     -22.009    sigmoidalRNA =       4.854
```

When alignments have been created using BLASTN, the name of the sequences involved in the alignment contains information about the alignment. That information is extracted from the BLASTN file, and processed by blastn2qrnadepth.pl in the following format:

`>name-from[><]to-more_info (length_alignment)`

where $>$ $(<)$ indicates the alignment involves the positive (negative) strand.
For instance in the previous example,

`>NT_034563-2555387-2555837-71>406-Telomerase_RNA (344)`

name = NT_034563-2555387-2555837
alignment involves positions 71 to 406 of that clone fragment, in the positive strand.
more_info = telomerase_RNA

`>chr3-frag968-96700001-96801000-15935>16265- (344)`

name = chr3-frag968-96700001-96801000
alignment involves positions 15935 to 16265 of the chromosome fragment, in the positive strand.
more_info = none

Analyze the results

After a genome–to–genome QRNA screen, one needs to post process the QRNA output. There are several Perl scripts to parse through a QRNA output file, and obtain different types of information.

phase_count_fast.pl

The first kind of processing one would like to do with a QRNA output is to extract the actual independent genomic regions (loci) that are identified as RNAs or coding by QRNA. A locus can be composed of a single window, a collection of continuous windows from a given alignment, or a collection of windows from different alignments that include that particular region, and are identified with the same function by QRNA. To obtain loci we have the script **phase_count_fast.pl**.

Command line:
    qrna-2.0.1/scripts/phase_count_fast.pl [options] foo.qrna2 query_organism database_organism

```
usage:  phase_count_fast.pl [options] file.qrna org1 org2


options:
-c <case>          :  cases (default is case = 1)
                      possible cases are:
                      0=GLOBAL
                      1=LOCAL_DIAG_VITERBI 2=LOCAL_DIAG_FORWARD
```

```
                             3=LOCAL_SEMI_VITERBI 4=LOCAL_SEMI_FORWARD
                             5=LOCAL_FULL_VITERBI 6=LOCAL_FULL_FORWARD
-i <id_max>        : max identity of alignments analysed (default is id_max = 100)
-l <loci_overlap>  : minimun overlap required to build loci (default is loci_overlap = -1)
-g <typetarget>    : which type of loci you want to analyze (default is typetarget = RNA)
                       possible types of loci are:
                       OTH | COD | RNA | RNA/OTH | RNA/COD | RNA/OTH | COD/OTH | RNA/COD/OTH
-q <file.q>        : include qfile to check if all the alignments were analysed
-s <type_of_score> : type of score (sigmoidal | simple)       [default = sigmoidal]
-t                 : towhomness -- obtains corresponding loci in the other organism
-u <cutoff>        : default is cutoff = 5
-w <whichorg>      : default is whichorg = 1  (use 1-for-org1 2-for-org2 12-for-both)
```

To exemplify some of those possibilities in directory "qrna-2.0.1/Demos/" I have included the following files

- "m52nc.fas-salmonella typhi.q"
  which is an already QRNA–processed collection of $12,037$ WUBLASTN alignments of intergenic *E. coli* to *Salmonella typhi*.

- "m52nc.fas-salmonella typhi.q.W150.X50.qrna2"
  which is the result of running the command line:
  qrna-2.0.1/src/qrna -w 150 -x 50   qrna-2.0.1/Demos/m52nc.fas-salmonella typhi.q

We can extract the RNA loci (with RNA score larger than a cutoff set with option -u to 0 bits) from the file "m52nc.fas-salmonella typhi.q.W150.X50.qrna2" using the following command,
qrna-2.0.1/scripts/phase count fast.pl -u 0.0 qrna-2.0.1/Demos/m52nc.fas-salmonella typhi.q.W150.X50.qrna2
                    Ecoli Styphi > qrna-2.0.1/Demos/m52nc.fas-salmonella typhi.q.W150.X50.qrna2.rnaloci

This scripts identifies 430 independent RNA loci in *E. coli* that score as RNA above 0 bits out of the $14,466$ windows. The script then proceeds to list the coordinates of each of those loci. The result has the following form,

```
--------------Some General Statistics------------------
FILE:                   ../Demos/m52nc.fas-salmonella_typhi.q.W150.X50.qrna2
method:                 LOCAL_DIAG_VITERBI
Cutoff:                 0.0

max id:                 100

# blastn hits:          12037
# windows:              14466
---------------------------------------------------------


--------------Statistics by Windows--------------------
# windows:              14466

RNA>0:                  5016/14466
RNA>cutoff:             5016/14466

COD>0:                  222/14466
COD>cutoff:             222/14466

 in phases:             14461/14466
       RNA:                    5016/14461
       COD:                    222/14461
       OTH:                    9223/14461

 in transitions:        0/14466
```

```
        RNA/COD:                    0/0
        RNA/OTH:                    0/0
        COD/OTH:                    0/0
        RNA/COD/OTH:                0/0
---------------------------------------------------------


--------------Statistics for RNA loci (Ecoli):------------------
# loci: 430
ave_length:      150.83


1-loci ecoli.m52-5531-5682(152) 19 96 (78) 5 RNA -11.20 2.49
2-loci ecoli.m52-9192-9305(114) 1 114 (114) 1 RNA -14.55 2.59
3-loci ecoli.m52-14080-14167(88) 1 88 (88) 1 RNA -9.55 2.23
4-loci ecoli.m52-20509-20814(306) 1 85 (85) 1 RNA -11.41 1.54
5-loci ecoli.m52-20509-20814(306) 224 305 (82) 2 RNA -8.18 1.90
```

The information given for each locus has the following form

```
num-loci name_seq(seq_length) loc_from loc_to (loc_length) number_wind type_loc COD_sc RNA_sc
```

Therefore,

```
1-loci ecoli.m52-5531-5682(152) 19 96 (78) 5 RNA -11.20 2.49
```

means that the first *E. coli* locus corresponds to intergenic sequence named "ecoli.m52-5531-5682" which has a length of 152 nucleotides. The locus has a length of 78 nucleotides and covers the region from nucleotide 19 to 96. Three different windows have contributed to this RNA locus, and the average sigmoidal score for the coding model is -11.20 bits, while the average sigmoidal score for the RNA model is 2.49 bits.

plot_qrna2.pl

Another kind of information that one might want to extract from a QRNA output is a distribution of scores according to percentage identity of the alignments involved. For this purpose we have the script **plot_qrna2.pl**.

Command line:
    qrna-2.0.1/scripts/plot_qrna2.pl [options] foo.qrna2 (foo.shuffle.qrna2)

```
usage: plot_qrna2.pl [options] qrnafile qrnafilesh
options:
-c <case>              :  cases (default is case = 1)
                          possible cases are:
                          0=GLOBAL
                          1=LOCAL_DIAG_VITERBI 2=LOCAL_DIAG_FORWARD
                          3=LOCAL_SEMI_VITERBI 4=LOCAL_SEMI_FORWARD
                          5=LOCAL_FULL_VITERBI 6=LOCAL_FULL_FORWARD
-a <max_n_ali>       : max number of alignments              [default all]
-e <evalue>          : fit to an e-value
-i <min_id_plot>     : for ID plots                          [default min_id_plot = 50]
-j <max_id_plot>     : for ID plots                          [default max_id_plot = 100]
-I <min_id>          : min ID for analysis                   [default min_id = 0]
-J <max_id>          : max ID for analysis                   [default max_id = 100]
-g <max_gap_plot>    : for GAP plots                         [default max_gap_plot = 50]
-l <max_num_win>     : max number of windows to be considered [default all]
-m <max_mut_plot>    : for MUT plots                         [default max_mut_plot = 50]
-n <time increments> : for TIME histo                        [default inc_t = 0.001]
-s <type_of_score>   : type of score (sigmoidal | simple)    [default = sigmoidal]
-t <max_time_plot>   : for plots                             [default max_time_plot = 0.8]
```

```
-T <max_time>          : max time for analysis              [default max_time = 0.8]
-S <min_time>          : min time for analysis              [default min_time = 0]
-u <cutoff>            : plot all logodds, but add stats for those above cutoff [default cutoff = 0]
-w <displaycutoff>     : plot logodds larger than dislplaycutoff         [default display_cutoff = 0]
-y <min_ord>           : min ordinate for plotting scores
-Y <max_ord>           : max ordinate for plotting scores
```

To exemplify some of those possibilities in directory "qrna-2.0.1/Demos/" I have included the following files

- "Rfam.seed.RF00017.SRP.fa.blast.D10.q"
  which is an already QRNA–processed collection of 813 WUBLASTN alignments between the 95 SRP RNAs in the seed of the database of noncoding RNAs Rfam 3.0.

- "Rfam.seed.RF00017.SRP.fa.blast.D10.q.qrna2"
  which is the result of running the command line:
  qrna-2.0.1/src/qrna   qrna-2.0.1/Demos/Rfam.seed.RF00017.SRP.fa.blast.D10.q

- "Rfam.seed.RF00017.SRP.fa.blast.D10.q.shuffle.qrna2"
  which is the result of running the command line:
  qrna-2.0.1/src/qrna   -B   qrna-2.0.1/Demos/Rfam.seed.RF00017.SRP.fa.blast.D10.q

The command line
qrna-2.0.1/scripts/plot_qrna2.pl   Rfam.seed.RF00017.SRP.fa.blast.D10.q.qrna2
                                          Rfam.seed.RF00017.SRP.fa.blast.D10.q.shuffle.qrna2

produces four output files in postscript:

- for COD scores

    - qrna-2.0.1/Demos/Rfam.seed.RF00017.SRP.fa.blast.D10.q.qrna2.id.spreadcod_histo_with_sh.ps

    - qrna-2.0.1/Demos/Rfam.seed.RF00017.SRP.fa.blast.D10.q.qrna2.id.spreadcod_scores_with_sh.ps

- for RNA scores

    - qrna-2.0.1/Demos/Rfam.seed.RF00017.SRP.fa.blast.D10.q.qrna2.id.spreadrna_histo_with_sh.ps

    - qrna-2.0.1/Demos/Rfam.seed.RF00017.SRP.fa.blast.D10.q.qrna2.id.spreadrna_scores_with_sh.ps

For each kind of score (coding or RNA), the two postscript files include the following information,

- The first postscript file is a histogram of the number of coding (RNA) scoring windows above a chosen cutoff (zero bits by default). See Figure 1 for an example.

- The second files produces a distribution of the coding (RNA) scores (average with standard deviation) with the percentage identity of the alignments. See Figure 2 for an example.

Figure 1: Analysis of the collection of 913 BLASTN alignments between SRP RNAs contained in file "Rfam.seed.RF00017.SRP.fa.blast.D10.q". Alignments have been grouped by percentage identity. This figure represent the histogram of the number of alignments bined in each percentage identity interval. In green we find the histogram for the total number of windows analyzed. In blue we have the histogram for those windows that score as RNA above a cutoff of 0 bits. In red we find the histogram for the windows that score as RNA above a cutoff of 0 bits after shuffling the alignments (false positives).

Figure 2: Analysis of the collection of 913 BLASTN alignments between SRP RNAs contained in file "Rfam.seed.RF00017.SRP.fa.blast.D10.q". Alignments have been grouped by percentage identity. This figure represent the RNA scores of all the alignments as a function of the percentage identity of the alignments. "∗" represents the average of the RNA scores. "□" represents the average of the RNA scores of the corresponding shuffled alignments (false positives). The error bars correspond to one standard deviation.