

rnaSeqMap: RNASeq analyses using xmap database back-end

Anna Lesniewska, Michal Okoniewski

October 18, 2010

Vignette for v.0.1.24 - created without XMAP database access!

Contents

1	Introduction	2
2	Database preparation	3
3	Analysis	4
4	ReadSet and NucleotideDistr classes	4
5	Region Mining algorithm	6
6	Interesting gene criteria	7
6.1	Coverage	7
6.2	Splicing	7
6.3	Fold change in regions	8
6.4	Region-based coverage	8
6.5	Utilities for transforming NucleotideDistr objects	8
7	Analyzing whole chromosomes	9
8	Producing output for DESeq and edgeR	10
9	Visualization	10

1 Introduction

rnaSeqMap is a "middleware" library for RNAseq secondary analyses. It constitutes an API for such operations as:

- access to any the reads of the experiment in possibly fastest time, according to any chromosome coordinates
- accessing sets of reads according to genomic annotation in Ensembl
- calculation of coverage and number of reads and transformations of those values
- creating input for significance analysis algorithms - from edgeR and DESeq
- precisely finding significant and consistent regions of expression
- splicing analyses
- visualizations of genes and expression regions

The library is independent from the sequencing technology and reads mapping software. It needs either reads described as genome coordinates in the extended xmap database, or can alternatively read data as big as they can fit in the operational memory. The use with modified xmap database is recommended, as it overcomes memory limitations - thus the library can be efficiently run on not very powerful machines.

The internal features of rnaSeqMap distinctive for this piece of software are:

- sequencing reads and annotations in one common database - extended XMAP [2]
- algorithm for finding irreducible regions of genomic expression - according to Aumann and Lindell [1]
- nucleotide-level splicing analysis
- connectors for further gene- and region-level expression processing to DESeq [5] and edgeR [6]
- the routines for coverage, splicing index and region mining algorithm have been implemented in C for speed

The area of application of the library may be adjusted to the type of experiment performed. Mainly, it is intended to analyze the data from RNAseq with depth-of-coverage estimation of expression, while it can be also used for SAGE, Chip-Seq, copy number analysis, etc. The functions in an appropriate context may be used not only for gene-level and region-level expression estimation, but also eg. for quality control of sequencing data.

The library follows the paradigms of the exonmap package (see [4]), continued currently as xmapcore. The point is: with the use of the annotation database, it gives a framework to analyze genes and intergenic regions gene by gene, filter the genes according to their qualities and visualize the internals of region expression. Due to the qualities of sequencing data, rnaSeqMap may be used not only to genes, but also to all the regions on the genome. The approach does not do any summarization in the initial phases so it is possible to have the fine-grained access by nucleotide (eg. as in NucleotideDistr class).

2 Database preparation

The database which is used in the analysis is xmapcore database, extended by additional relational tables (defined with indexing and partitioning) and stored procedures. First, an instance of xmapcore database has to be installed, then the script for adding the tables seq_reads and seq_experiments have to be run (see rnaSeqMap.sql script in the inst scripts) The table seq_reads is filled with the reads from the sequencers, which may be for example data from .SAM files transformed by sam2xmap.awk script

```
awk -f sam2xmap.awk input.sam > seq_reads.txt
```

The content of the awk script is as follows :

```
# AWK script to convert the SAM file to seq_read table input for X:MAP
# For mouse 19 chromosomes + 20=X,21=Y,22=MT
{
  OFS = "\t"
  strand = rshift($2,4) % 2
  if (strand==0) strand = -1
  cc = substr($3,4)
  if (cc!=cc+0)
  {
    if (cc=="X") cc=20
    else if (cc=="Y") cc=21
    else if (cc=="M") cc=22
    else cc=30;
  }

  print(6, cc, $4, $4+ length($10),strand)
  #remember to change first param to the number of sample!!
}
```

As the chromosome numbers are simplified in the database into numbers, the X, Y and mitochondrial chromosomes are coded as consecutive number after the last numbered chromosome - as in the example for mouse above. In the rnaSeqMap in R, they are recoded according to the species information set by selecting xmapcore database or by setSpecies()

After such preparation, all the data files must be loaded to the database, using mysqlimport, eg:

```
mysqlimport -u mysqlusername -p xmapcorei_homo_sapiens_58 -local seq_reads
```

The seq_reads table has then to be indexed as in the script inst scripts indices.sql.

The table bio_sample has to be created. It contains experiment description: all the samples in seq_reads are identified by an integer number, while in bio_sample the numbers have the treatments and other experimental attributes assigned. The content of this table is assigned to the phenoData slot of the NucleotideDistr objects that are processed.

3 Analysis

After loading the `rnaSeqMap` library, connecting to `xmap` database is done with `xmap.connect`.

```
> library(rnaSeqMap)
> xmap.connect("human-24")
```

Sample objects, used for analysis described in this vignette, may be found in the dataset:

```
> data(sample_data_rnaSeqMap)
> ls()

[1] "rs.list"

> class(rs.list)

[1] "list"

> length(rs.list)

[1] 17

> names(rs.list)

 [1] "ENSG00000164287" "ENSG00000118515" "ENSG00000146859" "ENSG00000174469"
 [5] "ENSG00000002726" "ENSG00000146540" "ENSG00000164877" "ENSG00000214783"
 [9] "ENSG00000165071" "ENSG00000185920" "ENSG00000133816" "ENSG00000129151"
[13] "ENSG00000137494" "ENSG00000220703" "ENSG00000249935" "ENSG00000177757"
[17] "ENSG00000223972"
```

The list contains 17 example genes with their coverage from a real experiment that contains 6 samples.

4 ReadSet and NucleotideDistr classes

`ReadSet` is an initial container for sequencing reads. Using the database filled with reads, the `SeqReads` objects are created by specifying their genome coordinates and filling them with reads information:

```
> rs <- newSeqReads(g.ch, st, en, str)
> rs <- addExperimentsToReadset(rs, 1:6)
```

There are functions that create `SeqReads` objects for genes:

```
> rs <- newSeqReadsFromGene(g)
> rs <- addExperimentsToReadset(rs, 1:6)
```

The table `bio_sample` has to be created. It contains experiment description: all the samples in `seq_reads` are identified by an integer number, while in `bio_sample` the numbers have the treatments and other experimental attributes assigned. The content of this table is assigned to the `phenoData` slot of the `NucleotideDistr` objects that are processed.

```
> rs <- rs.list[[1]]
> nd.cov <- getCoverageFromRS(rs, 1:6)
> distribs(nd.cov)[52000:52010, ]
```

```
      1  2 3 4 5 6
52000 4 10 3 1 1 4
52001 4  9 3 1 1 4
52002 4 10 3 1 1 2
52003 4 10 3 0 1 2
52004 4  9 3 0 1 2
52005 4  9 3 0 1 2
52006 3 10 3 0 2 2
52007 3 10 3 0 2 2
52008 3  9 3 0 2 2
52009 3  9 3 0 2 2
52010 3  9 3 0 2 2
```

The distributions in an object may be normalized by getting the sum of reads for each sample from the database or by any other scaling vector:

```
> nsums <- getSumsExp()
```

```
> nsums
```

```
[1] 1067676 25033137 85415567 21236253 18152006 25132926
```

```
> nd.cov <- normalizeBySum(nd.cov, nsums)
```

```
> distribs(nd.cov)[52000:52010, ]
```

```
      1  2 3 4 5 6
52000 4 0 0 0 0 0
52001 4 0 0 0 0 0
52002 4 0 0 0 0 0
52003 4 0 0 0 0 0
52004 4 0 0 0 0 0
52005 4 0 0 0 0 0
52006 3 0 0 0 0 0
52007 3 0 0 0 0 0
52008 3 0 0 0 0 0
52009 3 0 0 0 0 0
52010 3 0 0 0 0 0
```

5 Region Mining algorithm

The region mining algorithm is an adaptation of the sliding-and-combing two window algorithm by Lindell and Aumann [1]. Originally, it was used to find quantitative association rules from datasets with two numerical attributes. In the case of genomic data, especially RNAseq, the X dimension is the genome coordinate in bp, the Y dimension (decisive attribute) is the level of coverage for a given nucleotide. By specifying the μ and *minsup* parameters, we may tune it to find regions of a given magnitude of coverage and minimal length.

The regions found by this algorithm are proven to be irreducible. This means here, that if the region has mean coverage over μ , then when divided into any two sub-regions, those sub-regions will both have also mean coverage higher than μ . This follows the intuition of a region with a biologically consistent expression - namely - an exon. So if one applies properly tuned Lindell algorithm to the RNAseq data, should be able to find precise boundaries of real transcription.

From a `NucleotideDistr` object, another object may be created that includes the distributions equal 0 when the part of the distribution is not within irreducible region, or appropriate μ -level when the nucleotide in a given sample has a coverage that makes it a part of an irreducible region:

```
> nd.reg <- findRegionsAsND(nd.cov, 6, 10)
> distribs(nd.reg)[52000:52010, ]
```

```
      1 2 3 4 5 6
52000 6 0 0 0 0 0
52001 6 0 0 0 0 0
52002 6 0 0 0 0 0
52003 6 0 0 0 0 0
52004 6 0 0 0 0 0
52005 6 0 0 0 0 0
52006 6 0 0 0 0 0
52007 6 0 0 0 0 0
52008 6 0 0 0 0 0
52009 6 0 0 0 0 0
52010 6 0 0 0 0 0
```

Alternatively, regions from a single distribution may be also presented as an `RangedData` object:

```
> findRegionsAsIR(nd.cov, 6, 10, 1)
```

```
IRanges of length 9
      start      end width
[1] 54412169 54412201    33
[2] 54452556 54452582    27
[3] 54456010 54456244   235
[4] 54456846 54457070   225
[5] 54459903 54462265  2363
[6] 54462356 54462406    51
```

```
[7] 54462469 54462485    17
[8] 54462842 54462871    30
[9] 54463010 54463060    51
```

6 Interesting gene criteria

The genes may be regarded as interesting or significant according to many features, that can be calculated using their `NucleotideDistr` values.

6.1 Coverage

Coverage is included in the object (access with `distr()`), however in the heuristics for finding genes, we various statistical measures of the coverage for each sample:

```
> apply(distribs(nd.cov), 2, max)
  1  2  3  4  5  6
56  1  1  1  1  2

> apply(distribs(nd.cov), 2, mean)
      1          2          3          4          5          6
0.660823183 0.011963909 0.001661654 0.000797594 0.009105864 0.010784135

> apply(distribs(nd.cov), 2, sum)
  1    2    3    4    5    6
39769  720  100  48  548  649

> apply(distribs(nd.cov), 2, sd)
      1          2          3          4          5          6
3.29843494 0.10872429 0.04072985 0.02823068 0.09498998 0.10953246
```

6.2 Splicing

We use a modified version of splicing index from [3] paper. The value of index is calculated for each nucleotide in the investigated region. Gene level proportion is calculated using sum of coverage functions in the region for both treatments.

However, many nucleotides have no coverage in one or both conditions. Often the coverage in one condition is much higher than another. Thus the value of splicing index is set to 1 or -1 when second or first conditions have zero coverage. In the case when the \log_2 value exceeds the range (-1,1) it is limited to -1 or 1 respectively.

```
> nd.si <- getSIFromND(nd.cov, c(2, 3))
> distribs(nd.si)[52000:52010, ]
52000 52001 52002 52003 52004 52005 52006 52007 52008 52009 52010
      0      0      0      0      0      0      0      0      0      0      0
```

Splicing index may be plotted for both genes and any genomic regions:

```
> distrSIPlot(nd.cov, 1, 2, mi = 5, minsup = 5)
```

6.3 Fold change in regions

For the genes that have enough coverage, it is possible to calculate nucleotide-level fold change using:

```
> nd.fc <- getFCFromND(nd.cov, c(2, 3))
> distribs(nd.fc)[52000:52010, ]

52000 52001 52002 52003 52004 52005 52006 52007 52008 52009 52010
      0      0      0      0      0      0      0      0      0      0      0
```

6.4 Region-based coverage

Region-based coverage is a function of discretizing the `NucleotideDistr` object according to a sequence of Lindell algorithm searches. It assigns to the resulting `NucleotideDistr` the value of maximal μ from a sequence for which there is a region covering it

```
> nd.rbcov <- regionBasedCoverage(nd.cov, seqq = 1:10, maxexp = 40,
+   minsup = 10)
> distribs(nd.rbcov)[52000:52010, ]

      1  2  3  4  5  6
52000  8  0  0  0  0  0
52001  8  0  0  0  0  0
52002  8  0  0  0  0  0
52003  8  0  0  0  0  0
52004  8  0  0  0  0  0
52005  8  0  0  0  0  0
52006  8  0  0  0  0  0
52007  8  0  0  0  0  0
52008  8  0  0  0  0  0
52009  8  0  0  0  0  0
52010  8  0  0  0  0  0
```

Such a distribution is useful for removing artifacts from the distributions: high peaks and sudden drops of signal.

The plots are depicting the coverage before and after the region-based transformation.

```
> distrCOVPlot(nd.cov, c(3, 5))
> distrCOVPlot(nd.rbcov, c(3, 5))
```

6.5 Utilities for transforming `NucleotideDistr` objects

Additional utilities to operate on `NucleotideDistr` objects are functions for combining two, averaging, getting the sum or \log_2 for all the elements.

7 Analyzing whole chromosomes

The functions `spaceInChromosome` and `geneInChromosome` are finding all the genes and intergenic spaces in given regions. Such a table of genes or spaces may be used to build a genome-wide pipeline for finding interesting gene expression or novel areas of intergenic expression (possibly novel genes).

```
> my.genes <- geneInChromosome(22, 2e+07, 20400000, 1)
> my.genes
> my.spaces <- spaceInChromosome(22, 2e+07, 20400000, 1)
> my.spaces
```

The example of function used to evaluate genes or spaces:

```
> test.gene <- function(g, exps, nsums, mi = 5, minsup = 15) {
+   rs <- newSeqReadsFromGene(g)
+   rs <- addExperimentsToReadset(rs, exps)
+   cat("added reads..\n")
+   nd.cov <- getCoverageFromRS(rs, exps)
+   cat("calculated coverage...\n")
+   nd.cov <- normalizeBySum(nd.cov, nsums)
+   cat("normalized...\n")
+   nd.reg <- findRegionsAsND(nd.cov, as.integer(mi), minsup = minsup)
+   ir.reg <- findRegionsAsIR(nd.cov, as.integer(mi), minsup = minsup)
+   cat("ran region search algorithm...\n")
+   print(ir.reg)
+   out <- g
+   out <- c(out, apply(distribns(nd.cov), 2, max))
+   out <- c(out, apply(distribns(nd.cov), 2, mean))
+   out <- c(out, apply(distribns(nd.reg), 2, max))
+   out
+ }
> test.space <- function(exps, ch, st, en, str, nsums, mi = 5,
+   minsup = 15) {
+   g.ch <- rnaSeqMap:::.chromosome.number(ch)
+   rs <- newSeqReads(g.ch, st, en, str)
+   rs <- addExperimentsToReadset(rs, exps)
+   nd.cov <- getCoverageFromRS(rs, exps)
+   nd.cov <- normalizeBySum(nd.cov, nsums)
+   cat("Running ch ", g.ch, " start", st, "\n")
+   nd.reg <- findRegionsAsND(nd.cov, as.integer(mi), minsup = minsup)
+   out <- c(ch, st, en, str)
+   out <- c(out, apply(distribns(nd.cov), 2, max))
+   out <- c(out, apply(distribns(nd.cov), 2, mean))
+   out <- c(out, apply(distribns(nd.reg), 2, max))
+   out
+ }
> interesting.genes <- NULL
```

```

> for (i in 1:length(my.genes)) {
+   cat("Running gene ", i, "-----\n")
+   interesting.genes <- rbind(interesting.genes, test.gene(my.genes[i],
+     1:6, nsums))
+ }
> interesting.spaces <- NULL
> for (i in 104:(dim(my.spaces))[1]) {
+   cat("Running space ", i, "-----\n")
+   interesting.spaces <- rbind(interesting.spaces, test.space(1:2,
+     22, my.spaces[i, 1], my.spaces[i, 2], my.spaces[i, 3]))
+ }

```

8 Producing output for DESeq and edgeR

The functions `buildDESeq` and `buildDGEList` are the connectors between `rnaSeqMap` and the analysis of expression on the gene level as implemented in `DESeq` [5] and `edgeR` [6].

9 Visualization

`rnaSeqMap` has a set of chromosome-based plots that may be used for distributions contained in a `NucleotideDistr` object or just a region defined by chromosome coordinates - taking data from the database and not creating any objects.

Examples of the plots:

Histogram-based distribution of coverage:

```

> plotCoverageHistogram(1, 11000, 13500, 1, 1, skip = 50)

```

Coverage plots for a gene, with official exons of the gene plotted below:

```

> distrCOVPlotg(names(rs.list)[16], 1:3)

```

References

- [1] Aumann Y and Lindell Y, A Statistical Theory for Quantitative Association Rules, Journal of Intelligent Information Systems, 2003
- [2] Yates T, Okoniewski M and Miller C, X:Map: annotation and visualization of genome structure for Affymetrix exon array analysis, Nucleic Acids Research, 2007
- [3] Gardina P et al, Alternative splicing and differential gene expression in colon cancer detected by a whole genome exon array, BMC Genomics, 2006
- [4] Okoniewski M et al, An annotation infrastructure for the analysis and interpretation of Affymetrix exon array data, Genome Biology, 2007
- [5] Anders S and Huber W, Differential expression analysis for sequence count data, Nature Preceedings, 2010
- [6] Robinson M et al, edgeR: a Bioconductor package for differential expression analysis of digital gene expression data, Bioinformatics, Bioinformatics, 2009