

BitSeq User Guide

Peter Glaus, Antti Honkela and Magnus Rattray

October 13, 2015

1 Abstract

The *BitSeq* package is targeted for transcript expression analysis and differential expression analysis of RNA-seq data in two stage process. In the first stage it uses Bayesian inference methodology to infer expression of individual transcripts from individual RNA-seq experiments. The second stage of *BitSeq* embraces the differential expression analysis of transcript expression. Providing expression estimates from replicates of multiple conditions, Log-Normal model of the estimates is used for inferring the condition mean transcript expression and ranking the transcripts based on the likelihood of differential expression.

2 Citing BitSeq

The *BitSeq* package is based on probabilistic models and inference methods described in the manuscript [1]. For citing *BitSeq* in publications please refer to the manuscript above and to the source of the software.

3 Installing the BitSeq package

The recommended way to install *BitSeq* is to use the `biocLite` function available from the bioconductor website.

```
> source("http://www.bioconductor.org/biocLite.R")
> biocLite("BitSeq")
```

To load the package start R and run

```
> library(BitSeq)
```

4 Preparing data

The size of data normally analysed by *BitSeq* and results represented by samples from posterior distribution usually ranges in Gigabytes. Keeping this kind of datasets in memory within R environment would be rather inefficient and in most cases unnecessary, thus most of the data used within *BitSeq* is loaded and saved directly to the local hard drive.

The input for *BitSeq* package can be either SAM or BAM file containing aligned reads, as well as reference Fasta file. These files do not need to be loaded into the environment as *BitSeq* will read the data from disk space. The transcriptome Fasta file can be either downloaded from Ensembl website ¹ or constructed by UCSC genome browser at <http://genome.ucsc.edu/cgi-bin/hgTables>. In this example we will use file *ensSelect1.fasta* which contains sequence of all transcripts for five genes from the Ensembl homo sapiens, release 65 annotation.

For the alignment of reads to transcriptome we recommend using software *bowtie*, which is able to report all valid alignments for all the reads. The following *bash* commands can be used to create the SAM file with alignments:

¹link for homo sapiens, release 66: ftp://ftp.ensembl.org/pub/release-66/fasta/homo_sapiens/cdna/Homo_sapiens.GRCh37.66.cdna.all.fa.gz

```
# create bowtie reference index for the annotation
$ bowtie-build -f --ntoa ensSelect1.fasta ensSelect1-index
# align reads in data-c0b0.fastq against index
$ bowtie -q -v 3 -3 0 -p 4 -a -m 100 --sam ensSelect1-index \
  data-c0b0.fastq data-c0b0.sam
```

In the following examples we will be using the *data-c0b0.sam* file provided with the package. To make the life easier we set our current directory to the *extdata* directory.

```
> # save the current directory
> # (we move back to old_directory at the end of vignette)
> old_directory = getwd();
> on.exit(setwd(old_directory))
> # move to directory with the data
> setwd(system.file("extdata", package="BitSeq"));
```

5 Basic use

5.1 Estimating expression

To estimate expression we use the function `getExpression`, which takes as an input the SAM file with alignments as well as reference Fasta file that was used for the alignment. The function returns a `list` in which the first item `exp` is a `DataFrame` with expression mean and standard deviation of each transcript. The second item `fn` is a file name of a file containing all the expression samples, which are used in the later DE analysis. The last two items are counts, vector containing estimated read counts per transcript, and `trInfo`, `DataFrame` with information about transcripts.

The `log` option tells the function to return mean and standard deviation of logged samples and the last three options, which are parameters for the sampling algorithm, are passed to the `estimateExpression` function used for expression inference.

```
> res1 <- getExpression("data-c0b0.sam",
+   "ensSelect1.fasta",
+   log = TRUE,
+   MCMC_burnIn=200,
+   MCMC_samplesN=200,
+   MCMC_samplesSave=50,
+   seed=47)

[time: +0.00 m]
Reads: all(Ntotal): 4817  mapped(Nmap): 4663
[time: +0.00 m]
[time: +0.00 m]
Alignments: 21259.
[time: +0.00 m]
N mapped: 4663
N total: 4817
All alignments: 25910
Isoforms: 56
Burn in: 200 DONE. [time: +0.00 m]

Sampling DONE. [time: +0.00 m]
rHat (for 200 samples)
  rHat  (rH theta|  tid | mean theta)
1.2451 ( 1.2527 |   15 |  0.08371)
1.2394 ( 1.2936 |   14 |  0.02992)
1.1827 ( 1.1530 |   31 |  0.00660)
Mean rHat of worst 10 transcripts: 1.137175
Mean C0: (0 0 0 0 ). Nunmap: 154
```

Producing 915 final samples from each chain.

```
Sampling DONE. [time: +0.03 m]
rHat (for 915 samples)
  rHat   (rH theta|   tid | mean theta)
  1.0312 ( 1.0183 |    2 |  0.01813)
  1.0220 ( 1.0216 |   32 |  0.00756)
  1.0211 ( 1.0119 |   43 |  0.01082)
Mean rHat of worst 10 transcripts: 1.013715
Mean C0: (0 0 0 0 ). Nunmap: 154
```

```
Total samples: 5260
DONE. [time: +0.03 m]
```

The data being processed in this vignette is a small, example dataset, thus it is safe to use lower values for `MCMC_burnIn`, `MCMC_samplesN`, `MCMC_samplesSave`. For normal sized dataset, you can use default values of these parameters (all three parameters have default value 1000).

To view histogram of log RPKM expression use:

```
> hist(res1$exp$mean)
```

We can load the expression samples using function `loadSamples`, which returns `DataFrame` containing all expression samples.

```
> samples1 <- loadSamples(res1$fn)
```

Following command produces plot showing correlation between two transcript expression estimates:

```
> plot( unlist(s2["ENST00000436661",]), unlist(s2["ENST00000373501",]))
```

The `getExpression` function first computes probabilities for every alignment and then uses this data in Markov chain Monte Carlo algorithm which samples from the posterior distribution of transcript expression. Both these steps are computationally very expensive and can take several hours to finish. In case of MCMC sampling, which has to converge to the correct posterior distribution this can take more than day for extensive data.

5.2 Identifying differentially expressed transcripts

In the differential expression analysis we will use the expression samples produced in first step as well as expression samples from other experiments provided with the package. It is essential to use biological replicates in both conditions in order to account for biological variation which could otherwise cause false positive DE calls.

```
> cond1Files = c(res1$fn,"data-c0b1.rpkm")
> cond2Files = c("data-c1b1.rpkm","data-c1b1.rpkm")
```

We use the function `getDE` to infer the Probability of Positive Log Ratio for each transcript. The function again returns list with first item `pplr` containing `DataFrame` with PPLR and other information. The second item `fn` contains list of filenames with the names of produced files. Using the options `samples=TRUE`, the function creates also files containing the condition mean expression samples.

```
> de1 <- getDE(list(cond1Files, cond2Files), samples=TRUE)
```

```
[time: +0.00 s]
# 5 done. [time: +0.00 s]
# 10 done. [time: +0.00 s]
# 15 done. [time: +0.00 s]
# 20 done. [time: +0.00 s]
# 25 done. [time: +0.00 s]
# 30 done. [time: +0.00 s]
# 35 done. [time: +0.00 s]
# 40 done. [time: +0.00 s]
```

```
# 45 done. [time: +0.00 s]
# 50 done. [time: +0.00 s]

> print(de1$fn)

$pplr
[1] "/tmp/Rtmp5cYbZQ/dataBS-DE-1735554756dec.pplr"

$samplesFiles
[1] "/tmp/Rtmp5cYbZQ/dataBS-DE-1735554756dec-C0.est"
[2] "/tmp/Rtmp5cYbZQ/dataBS-DE-1735554756dec-C1.est"
```

Now we can rank the transcripts based on the PPLR value to identify the ones with the highest probability of being differentially expressed:

```
> head( de1$pplr[ order(abs(0.5-de1$pplr$pplr), decreasing=TRUE ), ], 5)

DataFrame with 5 rows and 6 columns
      pplr 1~2 log2FC 1~2      ciLow      ciHigh
      <numeric> <numeric> <numeric> <numeric>
ENST00000370994 0.0192308 -0.936489 -1.6936100 -0.110386
ENST00000465588 0.9230770  0.808699 -0.0793691  1.927240
ENST00000473383 0.8461540  1.174720 -0.5146190  3.702100
ENST00000476493 0.8076920  0.622623 -0.8252730  2.027940
ENST00000493607 0.2307690 -0.312473 -1.3938300  0.558528
      mean 1      mean 2
      <numeric> <numeric>
ENST00000370994 9.48878  8.83966
ENST00000465588 8.83961  9.40016
ENST00000473383 6.65256  7.46682
ENST00000476493 8.17439  8.60596
ENST00000493607 10.68160 10.46500
```

6 Advanced use

Both expression estimation and identification of differentially expressed transcripts involves multiple steps which are independent. Computing these steps independently might be useful for keeping intermediate results in case of crash or error. As BitSeq makes extensive use of local files, it is essential to set path to working directory containing alignment files and which will be used for storing results of individual steps. In this example we use the *extdata* directory provided with the package:

```
> setwd(system.file("extdata", package="BitSeq"))
```

6.1 Stage 1 - Transcript expression analysis

6.1.1 Pre-processing alignments

In the pre-processing step, the `parseAlignment` function reads the SAM file and assigns a probability to every valid alignment. These probabilities are saved into *.prob* file and are the direct input for the expression estimation. We have to specify the reference file which is used for identifying base mismatches and we use uniform model for the read distribution along transcript:

```
> parseAlignment( "data-c0b0.sam",
+   outFile = "data-c0b0.pro",
+   trSeqFile = "ensSelect1.fasta",
+   trInfoFile = "data.tr",
+   verbose = TRUE )
```

```

Assuming alignment file in 'sam' format.
Using alignments' header for transcript information.
Initializing fasta sequence reader.
Found gene names in sequence file, updating transcript information.
[time: +0.00 m]
Using uniform read distribution.
Reads: all(Ntotal): 4817 mapped(Nmap): 4663
    491 reads were used to estimate empirical distributions.
[time: +0.02 m]
Writing alignment probabilities.
# 481 done. [time: +0.00 m]
# 962 done. [time: +0.00 m]
# 1443 done. [time: +0.00 m]
# 1924 done. [time: +0.00 m]
# 2405 done. [time: +0.00 m]
# 2886 done. [time: +0.00 m]
# 3367 done. [time: +0.00 m]
# 3848 done. [time: +0.00 m]
# 4329 done. [time: +0.00 m]
# 4810 done. [time: +0.00 m]
[time: +0.00 m]
Analyzed 4817 reads:
    154 had no alignments
The rest had 21259 alignments:
    21259 single-read alignments
Computing effective lengths.
Transcript information saved into data.tr.
[time: +0.00 m]
DONE. [time: +0.02 m]

```

The program passes the SAM file twice and produces the *data-c0b0.prob* file with the alignment probabilities as well as transcript information file *data.tr* which contains transcript names and lengths extracted from the SAM file.

6.1.2 Estimating transcript expression

The `estimateExpression` function implements a generative model of RNA-seq data and infers the transcript expression using Markov chain Monte Carlo algorithm. The default MCMC algorithm is the Collapsed Gibbs sampling which converges faster than regular Gibbs sampling (selectable by option `gibbs=TRUE`). It is the most time consuming part of the *BitSeq* analysis process as it uses multiple independent chains to sample the expression values and it iterates until the chains converge to the same distribution.

The following example runs the sampler using the *.prob* file from previous step, produces expression in RPKM measure and produces files with the prefix *data-c0b0*. It will produce two files, file *data-c0b0.rpkm* will contain a row for each transcript with MCMC_samplesSave RPKM expression samples. The second file *data-c0b0.thetaMeans* will contain the mean relative expression values for every transcript.

```

> estimateExpression("data-c0b0.prob", outFile = "data-c0b0",
+   outputType = "RPKM", trInfoFile = "data.tr",
+   MCMC_burnIn = 200, MCMC_samplesN = 200, MCMC_samplesSave = 100,
+   MCMC_chainsN = 2)

```

```

N mapped: 4663
N total: 4817
All alignments: 25910
Isoforms: 56
Burn in: 200 DONE. [time: +0.00 m]

```

```

Sampling DONE. [time: +0.00 m]
rHat (for 200 samples)
    rHat    (rH theta|    tid | mean theta)

```

```

1.0992 ( 1.0523 |      2 | 0.01449)
1.0774 ( 1.0867 |     13 | 0.00274)
1.0476 ( 1.0484 |     43 | 0.01256)
Mean rHat of worst 10 transcripts: 1.039252
Mean C0: (0 0 ). Nunmap: 154

```

Producing 1429 final samples from each chain.

```

Sampling DONE. [time: +0.03 m]
rHat (for 1429 samples)
  rHat   (rH theta|      tid | mean theta)
1.0192 ( 1.0188 |      54 | 0.00787)
1.0084 ( 1.0082 |       1 | 0.16534)
1.0052 ( 1.0047 |      16 | 0.00495)
Mean rHat of worst 10 transcripts: 1.005553
Mean C0: (0 0 ). Nunmap: 154

```

```

Total samples: 3658
DONE. [time: +0.03 m]

```

The behavior of the sampling algorithm can be adjusted by optional arguments, such as `MCMC_chainsN` which selects the number of chains. After producing `MCMC_burnIn` burn-in samples, the algorithm produces first `MCMC_samplesN` samples from each chain in the first iteration. These are used to estimate the number of samples needed for recording `MCMC_samplesSave` effective samples, in the second, final, iteration.

6.1.3 Convergence checking via possible scale reduction estimation

The `estimateExpressionLegacy` uses different convergence checking mechanism which mostly results in multiple iterations, producing more samples in total. After each iteration, the possible scale reduction of marginal posterior variance \hat{R} is computed for each transcript expression and the ten highest values are reported. If the average of ten highest possible scale reductions is less than the `MCMC_scaleReduction` parameter, then the sampler produces one last iteration during which subset of `MCMC_samplesSave` samples is recorded. Otherwise the program continues with next iteration in which it produces twice as many samples. The program terminates either after reaching the target scale reduction or after iteration which produces `MCMC_samplesNmax` samples. All these parameters can be set also within a parameters file specified by the option `parFile` with the advantage that the parameters such as `MCMC_scaleReduction` or `MCMC_samplesNmax` can be adjusted while the sampler is running, example of the parameters file *parameters1.txt* is provided in the *extdata* directory.

```

> estimateExpressionLegacy("data-c0b0.pro", outFile = "data-c0b0",
+   outputType = "RPKM", trInfoFile = "data.tr",
+   MCMC_burnIn = 200, MCMC_samplesN = 200, MCMC_samplesSave = 100,
+   MCMC_scaleReduction = 1.1, MCMC_chainsN = 2)

```

```

N mapped: 4663
N total: 4817
All alignments: 25910
Isoforms: 56
Burn in: 200 DONE. [time: +0.00 m]

```

```

Sampling DONE. [time: +0.00 m]
rHat (for 200 samples)
  rHat   (rH theta|      tid | mean theta)
1.1870 ( 1.1659 |       2 | 0.01799)
1.0577 ( 1.0729 |      20 | 0.00129)
1.0466 ( 1.0478 |      31 | 0.00741)
Mean rHat of worst 10 transcripts: 1.050925
(target: 1.100)
Mean C0: (0 0 ). Nunmap: 154

```

Producing 200 final samples from each chain.

```
Sampling DONE. [time: +0.00 m]
rHat (for 200 samples)
  rHat   (rH theta|   tid | mean theta)
  1.1393 ( 1.1277 |   43 |  0.01162)
  1.1369 ( 1.1248 |   14 |  0.02514)
  1.0979 ( 1.0808 |   32 |  0.00761)
Mean rHat of worst 10 transcripts: 1.084790
(target: 1.100)
Mean CO: (0 0 ). Nunmap: 154
```

```
WARNING: Following transcripts failed to converge entirely
(however the estimates might still be usable):
  2 transcripts (full list is in the output file)
Total samples: 1200
DONE. [time: +0.00 m]
```

6.1.4 Examining the samples

Again, we can view the resulting file (*data-c0b0.rpkm* in this case) using the function `loadSamples`.

```
> loadSamples("data-c0b0.rpkm")
```

6.1.5 Estimating transcript expression with Variational Bayes algorithm

Variational Bayes is an alternative inference algorithm for estimating transcript expression. While it accurately estimates mean expression, it underestimates the variance. The algorithm is implemented in the function `estimateVBExpression` which takes the same input as `estimateExpression`, the *.prob* file and optionally the *.tr* file.

The default output is the mean θ and parameters of Dirichlet distribution reported in *data-c0b0-vb.m_alphas*. The θ can be multiplied by total number of reads to produce the estimated read counts per transcript. The first row of the data corresponds to noise transcript and should be discarded.

```
> estimateVBExpression("data-c0b0.pro", outFile = "data-c0b0-vb",
+   outputType = "RPKM", trInfoFile = "data.tr")
```

```
N mapped: 4663
N total: 4817
All alignments: 25910
Isoforms: 56
```

```
End: bound decrease
iter(s):   38  bound: -58548.611 grad: 0.0001778 beta: 0.8942972
```

The function can also produce also samples from the posterior distribution in form of RPKM, θ or counts when option `samples` is used.

6.2 Stage 2 - Differential expression analysis

6.2.1 Preparing for Differential Expression analysis

In the differential expression analysis we are comparing samples from two different conditions. Also in order to estimate biological variance of transcript expression, we have to use data from at least one extra biological replicates. We first specify the files containing expression samples from the Stage 1, using file *data-c0b0.rpkm* computed in previous example and three other files provided with the package:

```
> allConditions = list(c("data-c0b0.rpkm", "data-c0b1.rpkm"),
+   c("data-c1b1.rpkm", "data-c1b1.rpkm"))
```

The estimation of expression specific hyperparameters for the DE model requires pre computing joint mean expression over all experiments using the `getMeanVariance` function. As the DE model uses logged expression samples, we have to compute the mean and variance of logged expression samples:

```
> getMeanVariance(allConditions, outFile = "data.means", log = TRUE )
```

6.2.2 Estimating model hyperparameters

The hyperparameters for the model are estimated from the entire data using Metropolis-Hastings MCMC algorithm. The values are smoothed afterwards using the non-parametric Lowess smoothing algorithm:

```
> estimateHyperPar( outFile = "data.par",
+   conditions = allConditions,
+   meanFile = "data.means",
+   verbose = TRUE )
```

```
Transcripts in expression file: 55
Samples are not logged. (will log for you)
Using -100 as minimum instead of log(0).
Files with samples loaded.
Number of all replicates: 4
seed: 1444791434
Expression step: 0.0433709
[time: +0.00 s]
Running sampler.
.....
Sampling done.
Have 10 parameters to smooth.
# alphaSmooth f: 0.2 nSteps: 5
# betaSmooth f: 0.2 nSteps: 5
DONE.
[time: +0.10 m]
```

6.2.3 Inferring condition mean expression and calculating Probability of Positive Log Ratio

The model for Differential Expression analysis uses the posterior samples from expression analysis to infer samples of the mean expression for each transcript in every condition. Function `estimateDE` computes the samples and uses them to compute the Probability of Positive Log Ratio, which is the probability of a transcript being up-regulated in the first condition as well as inverse probability of transcript being down-regulated in the first condition. The PPLR, mean \log_2 fold change with confidence intervals and mean condition mean expression are saved into the final output file with extension `.pplr` and prefix specified by the option `outFile`:

```
> estimateDE(allConditions, outFile = "data", parFile = "data.par" )

# 5 done. [time: +0.00 s]
# 10 done. [time: +0.00 s]
# 15 done. [time: +0.00 s]
# 20 done. [time: +0.00 s]
# 25 done. [time: +0.00 s]
# 30 done. [time: +0.00 s]
# 35 done. [time: +0.00 s]
# 40 done. [time: +0.00 s]
# 45 done. [time: +0.00 s]
# 50 done. [time: +0.00 s]

> ##
> ## pretend run with three conditions and normalization constants
> ##
> cond3Files = c("data-c2b0.rpkm", "data-c2b1.rpkm", "data-c2b2.rpkm")
> estimateDE(list( allConditions[[1]], allConditions[[2]], cond3Files),
```



```
+      outFile="mydata",
+      parFile="mydata.par",
+      norm=c(1.0, 0.999, 1.0017, 0.9998, 1.0, 0.99, 0.97),
+      pretend=TRUE)

estimateDE data-c0b0.rpkm data-c0b1.rpkm C data-c1b1.rpkm data-c1b1.rpkm C\
  data-c2b0.rpkm data-c2b1.rpkm data-c2b2.rpkm --outPrefix mydata\
  --parameters mydata.par --norm 1,0.999,1.0017,0.9998,1,0.99,0.97
```

In case one is interested in the condition mean expression samples as well, they can be obtained by using the `samples` flag:

```
> estimateDE(allConditions, outFile = "data", parFile = "data.par",
+   samples = TRUE )

# 5 done. [time: +0.00 s]
# 10 done. [time: +0.00 s]
# 15 done. [time: +0.00 s]
# 20 done. [time: +0.00 s]
# 25 done. [time: +0.00 s]
# 30 done. [time: +0.00 s]
# 35 done. [time: +0.00 s]
# 40 done. [time: +0.00 s]
# 45 done. [time: +0.00 s]
# 50 done. [time: +0.00 s]
```

This produces three extra files, the first two *data-C0.est*, *data-C1.est* containing the condition means for each condition and the third file *data.estVar* containing samples of inferred variance for the first condition.

7 External use of BitSeq

All major computation in *BitSeq* is executed by running C++ libraries, and there is also a C++ only implementation of *BitSeq* package available at <https://github.com/BitSeq/BitSeq>. The standalone package can be particularly useful for clusters without support for R or Bioconductor.

In order to facilitate the use of C++ version of the package, the relevant functions in R interface provide `pretend` option. Using this option, the computation will not be executed, instead the function will print out one or more command line commands which can be directly used with the C++ implementation.

```
> res1 <- getExpression("data-c0b0.sam",
+   "ensSelect1.fasta",
+   outPrefix="localDir/data-c0b0",
+   log = TRUE,
+   MCMC_burnIn=200,
+   MCMC_samplesN=200,
+   pretend=TRUE)

parseAlignment data-c0b0.sam --outFile localDir/data-c0b0.proba --trSeqFile\
  ensSelect1.fasta --trInfoFile localDir/data-c0b0.tr --uniform
estimateExpression localDir/data-c0b0.proba --outPrefix localDir/data-c0b0\
  --outType rpkm --trInfoFile localDir/data-c0b0.tr --MCMC_burnIn 200\
  --MCMC_samplesN 200
getVariance localDir/data-c0b0.rpkm --outFile localDir/data-c0b0.mean --log
```

8 Session Info

```
> sessionInfo()
```

```
R version 3.2.2 (2015-08-14)
Platform: x86_64-apple-darwin13.4.0 (64-bit)
Running under: OS X 10.9.5 (Mavericks)

locale:
[1] C/en_US.UTF-8/en_US.UTF-8/C/en_US.UTF-8/en_US.UTF-8

attached base packages:
[1] parallel  stats4    stats      graphics  grDevices
[6] utils     datasets  methods    base

other attached packages:
[1] BitSeq_1.14.0      zlibbioc_1.16.0
[3] Rsamtools_1.22.0   Biostrings_2.38.0
[5] XVector_0.10.0     GenomicRanges_1.22.0
[7] GenomeInfoDb_1.6.0 IRanges_2.4.0
[9] S4Vectors_0.8.0    BiocGenerics_0.16.0

loaded via a namespace (and not attached):
[1] bitops_1.0-6        futile.options_1.0.0
[3] futile.logger_1.4.1 BiocStyle_1.8.0
[5] BiocParallel_1.4.0  lambda.r_1.1.7
[7] tools_3.2.2
```

References

- [1] Glaus, P., Honkela, A. and Rattray, M. (2012). Identifying differentially expressed transcripts from RNA-seq data with biological variation. *Bioinformatics*, **28**(13), 1721-1728.