

# baySeq: Empirical Bayesian analysis of patterns of differential expression in count data

Thomas J. Hardcastle

April 22, 2010

## 1 Introduction

This vignette is intended to give a rapid introduction to the commands used in implementing two methods of evaluating differential expression in Solexa-type, or *count* data by means of the `baySeq` R package. For fuller details on the methods being used, consult Hardcastle & Kelly [1]. The major improvement made in this release is the option to include region length in evaluating differential expression between genomic regions (e.g. genes). See Section 7 for more details.

We assume that we have discrete data from a set of sequencing or other high-throughput experiments, arranged in a matrix such that each column describes a sample and each row describes some entity for which counts exist. For example, the rows may correspond to the different sequences observed in a sequencing experiment. The data then consists of the number of times each sequence is observed in each sample. We wish to determine which, if any, rows of the data correspond to some patterns of differential expression across the samples. This problem has been addressed for pairwise differential expression by the `edgeR` [2] package.

However, `baySeq` takes an alternative approach to analysis that allows more complicated patterns of differential expression than simple pairwise comparison, and thus is able to cope with more complex experimental designs. We also observe that the methods implemented in `baySeq` perform at least as well, and in some circumstances considerably better than those implemented in `edgeR` [1].

`baySeq` uses empirical Bayesian methods to estimate the posterior likelihoods of each of a set of models that define patterns of differential expression for each row. This approach begins by considering a distribution for the row defined by a set of underlying parameters for which some prior distribution exists. By estimating this prior distribution from the data, we are able to assess, for a given model about the relatedness of our underlying parameters for multiple libraries, the posterior likelihood of the model.

In forming a set of models upon the data, we consider which patterns are biologically likely to occur in the data. For example, suppose we have count data from some organism in condition *A* and condition *B*. Suppose further that we have two biological replicates for each condition, and hence four libraries  $A_1, A_2, B_1, B_2$ , where  $A_1, A_2$  and  $B_1, B_2$  are the replicates. It is reasonable to suppose that at least some of the rows may be unaffected by our experimental conditions *A* and *B*, and the count data for each sample in these rows will be *equivalent*. These data need not in general be identical across each sample

due to random effects and different library sizes, but they will share the same underlying parameters. However, some of the rows may be influenced by the different experimental conditions  $A$  and  $B$ . The count data for the samples  $A_1$  and  $A_2$  will then be equivalent, as will the count data for the samples  $B_1$  and  $B_2$ . However, the count data between samples  $A_1, A_2, B_1, B_2$  will not be equivalent. For such a row, the data from samples  $A_1$  and  $A_2$  will then share the same set of underlying parameters, the data from samples  $B_1$  and  $B_2$  will share the same set of underlying parameters, but, crucially, the two sets will not be identical.

Our task is thus to determine the posterior likelihood of each model for each row of the data. We can do this by considering either a Poisson or negative-binomial distribution upon the sequencing count data. The Poisson method is considerably faster as a closed form conjugate prior exists for this distribution. The negative-binomial solution is slower as it requires a numerical solution for the prior, but is probably a better fit for most data. In experimental data, we have found that the Poisson method is likely to give poor results if true biological replicates are not available; in most human studies, for example. In general, therefore, the use of the negative-binomial methods is recommended.

## 2 Preparation

We begin by loading the `baySeq` package.

```
> library(baySeq)
```

Note that because the experiments that `baySeq` is designed to analyse are usually massive, we should use (if possible) parallel processing as implemented by the `snow` package. We therefore need to load the `snow` package (if it exists), define a *cluster* and load the `baySeq` library onto each member of the cluster. If `snow` is not present, we can proceed anyway with a `NULL` cluster. Results may be slightly different depending on whether or not a cluster is used owing to the non-deterministic elements of the method.

```
> if ("snow" %in% installed.packages()[, 1]) {  
+   library(snow)  
+   cl <- makeCluster(4, "SOCK")  
+ } else cl <- NULL
```

Here we have (if the `snow` package is installed) defined a cluster of four processors on sockets; that is to say, on the local machine. If the local machine has multiple processors this may be a valid method of accelerating `baySeq`, but if very large data sets are being analysed we may wish to consider some other form of parallelisation (e.g. LAM/MPI) that allows processors on multiple nodes to be used. See the `snow` documentation for details on how to achieve this.

We load a simulated data set consisting of count data on one thousand counts and library sizes for ten libraries.

```
> data(simCount)  
> data(libsizes)  
> simCount[1:10, ]
```

	[,1]	[,2]	[,3]	[,4]	[,5]	[,6]	[,7]	[,8]	[,9]	[,10]
[1,]	4	1	5	2	3	0	1	1	1	0
[2,]	1	0	9	6	5	0	1	0	0	1
[3,]	9	2	5	5	14	2	3	1	0	4
[4,]	7	3	8	2	2	0	1	0	1	0
[5,]	2	2	4	7	0	0	0	0	0	1
[6,]	2	1	0	1	0	4	3	5	5	3
[7,]	9	8	8	8	9	1	2	1	0	0
[8,]	9	5	7	8	7	1	2	0	1	2
[9,]	6	2	2	3	0	0	0	0	0	0
[10,]	1	0	2	0	1	3	17	2	2	10

```
> libsizes
```

```
[1] 75373 40153 75403 34285 55975 53287 80477 37655 41171 77510
```

The data are simulated such that the first hundred counts show differential expression between the first five libraries and the second five libraries. Our replicate structure, used to estimate the prior distributions on the data, can thus be defined as

```
> replicates <- c(1, 1, 1, 1, 1, 2, 2, 2, 2, 2)
```

We can also establish two group structures for the data.

Each member (vector) contained within the 'groups' list corresponds to one model upon the data. In this setting, a model describes the patterns of data we expect to see at least some of the tags correspond to. In this simple example, we expect that some of the tags will be equivalently expressed between all ten libraries. This corresponds to the 'NDE' model, or vector  $c(1,1,1,1,1,1,1,1,1,1)$  - all libraries belong to the same group for these tags.

We also expect that some tags will show differential expression between the first five libraries and the second five libraries. For these tags, the two sets of libraries belong to different groups, and so we have the model 'DE', or vector  $c(1,1,1,1,1,2,2,2,2,2)$  - the first five libraries belong to group 1 and the second five libraries to group 2. We thus have the following group structure

```
> groups <- list(NDE = c(1, 1, 1, 1, 1, 1, 1, 1, 1, 1), DE = c(1,
+ 1, 1, 1, 1, 2, 2, 2, 2, 2))
```

In a more complex experimental design (Section 1) we might have several additional models. The key to constructing vectors corresponding to a model is to see for which groups of libraries we expect equivalent expression of tags.

We note that the group for DE corresponds to the replicate structure. This will often be the case, but need not be in more complex experimental designs.

The ultimate aim of the `baySeq` package is to evaluate posterior likelihoods of each model for each row of the data.

We begin by combining the count data, library sizes and user-defined groups into a `countData` object.

```
> CD <- new("countData", data = simCount, replicates = replicates,
+ libsizes = libsizes, groups = groups)
```

We can also optionally add annotation details into the `@annotation` slot of the `countData` object.

```
> CD@annotation <- data.frame(name = paste("count", 1:1000, sep = "_"))
```

### 3 Poisson-Gamma Approach

We first try to identify the posterior likelihoods of each model for each tag assuming a Poisson distribution on each tag with a rate that is Gamma distributed. That is, if  $Y_{ij}$  is an element of the data where  $i$  is the row of the data, and  $j$  is the sample, then

$$Y_{ij} \sim Poi(\theta_j l_j)$$

where the  $l_j$  is the library size of sample  $j$  (or some other suitable scaling factor) and

$$\theta_j \sim \Gamma(\alpha_j, \beta_j)$$

The relationships between the  $\alpha_j, \beta_j$  for each  $j$  are determined by the model being investigated such that, if and only if samples  $X$  and  $Y$  belong to the same group, then  $\alpha_X = \alpha_Y$  and  $\beta_X = \beta_Y$ .

We begin by trying to establish the parameters of the Gamma distribution by bootstrapping from the data and applying maximum likelihood methods. We are able to adjust the parameters of the bootstrapping; here we take twenty sets of count data, establish the maximum likelihood Gamma parameters, and iterate over 1000 cases. In general more than 5000 iterations is recommended but is used here for speed of calculation.

We then take the mean of the maximum likelihood estimates to acquire a prior on the rate distribution.

```
> CDP.Poi <- getPriors.Pois(CD, samplesize = 20, takemean = TRUE,
+   cl = cl)
```

The calculated priors are stored in the `@priors` slot of the `countData` object produced.

```
> CDP.Poi@priors

$priors
$priors$NDE
$priors$NDE[[1]]
[1] 1.274190 28204.487818

$priors$DE
$priors$DE[[1]]
[1] 9.936345e-01 2.005718e+04

$priors$DE[[2]]
[1] 1.071952 24986.724537
```

For each model, we get a set of priors. In the Poisson-Gamma approach, we get, for each group in the model, a pair of parameters which define the Gamma distribution that we shall use as a prior distribution for the rates of the Poisson distributions that describe how many counts we see in each row of the data. Thus, in the model of differential expression, there are two groups in the data and we find two sets of parameters.

Having acquired a set of prior distributions on the rate parameter of the Poisson distribution, we can calculate the posterior likelihoods of each model for each tag. We need to either provide an initial prior likelihood on each model via the `prs` parameter, or provide some other means of estimating the prior likelihood on the model. If `pET = "BIC"` then the prior likelihood on each model will be estimated from the Bayesian Information Criterion, and the `prs` parameter is not necessary (and will be ignored).

```
> CDPPost.Poi <- getLikelihoods.Pois(CDP.Poi, pET = "BIC", cl = cl)
> CDPPost.Poi@estProps
```

```
      NDE      DE
0.6835688 0.3164312
```

```
> CDPPost.Poi@posteriors[1:10, ]
```

```
      NDE      DE
[1,] -2.445109 -9.070917e-02
[2,] -6.921181 -9.871512e-04
[3,] -5.273955 -5.136475e-03
[4,] -7.510238 -5.476007e-04
[5,] -4.978730 -6.906592e-03
[6,] -3.250291 -3.953422e-02
[7,] -16.039593 -1.081667e-07
[8,] -9.778648 -5.664996e-05
[9,] -5.714065 -3.304690e-03
[10,] -10.191890 -3.747369e-05
```

```
> CDPPost.Poi@posteriors[101:110, ]
```

```
      NDE      DE
[1,] -0.2004772361 -1.705619123
[2,] 0.0000000000 -35.539787903
[3,] -0.1319280436 -2.090737547
[4,] -0.1018535009 -2.334714296
[5,] -0.0004798705 -7.642234268
[6,] -0.1591869874 -1.916213609
[7,] -0.3361409721 -1.253591610
[8,] -6.0948962900 -0.002256889
[9,] -0.1510126003 -1.964948280
[10,] -0.0683643507 -2.716891226
```

The estimated posterior likelihoods for each model are stored in the natural logarithmic scale in the `@posteriors` slot of the `countDataPosterior`. The  $n$ th column of the posterior likelihoods matrix corresponds to the  $n$ th model as listed in the `group` slot of `CDPPost.Poi`.

Here the assumption of a Poisson distribution gives an estimate of

```
DE
0.3164312
```

as the proportion of differential expressed counts in the simulated data, where in fact the proportion is known to be 0.1.

## 4 Negative-Binomial Approach

We next try the same analysis assuming a Negative Binomial distribution on the data. We first estimate an empirical distribution on the parameters of the Negative Binomial distribution by bootstrapping from the data, taking individual counts and finding the quasi-likelihood parameters for a Negative Binomial distribution. By taking a sufficiently large sample, an empirical distribution on the parameters is estimated. A sample size of around 10000 iterations is suggested, depending on the data being used), but 1000 is used here to rapidly generate the plots and tables.

```
> CDP.NBML <- getPriors.NB(CD, samplesize = 1000, estimation = "QL",
+   c1 = c1)
```

The calculated priors are stored in the `@priors` slot of the `countData` object produced as before. For the negative-binomial method, we are unable to form a conjugate prior distribution. Instead, we build an empirical prior distribution which we record in the list object `$priors` of the slot `@priors`. Each member of this list object corresponds to one of the models defined by the `group` slot of the `countData` object and contains the estimated parameters for each of the individual counts selected under the models. The vector `$sampled` contained in the slot `@priors` describes which rows were sampled to create these sets of parameters.

We then acquire posterior likelihoods as before, estimating the proportions of differentially expressed counts. We can repeatedly bootstrap the prior estimation to improve accuracy; here three bootstraps are used.

```
> CDPPost.NBML <- getLikelihoods.NB(CDP.NBML, pET = "BIC", c1 = c1)
```

```
.
```

```
> CDPPost.NBML@estProps
```

```
[1] 0.7756488 0.2243512
```

```
> CDPPost.NBML@posteriors[1:10, ]
```

	NDE	DE
[1,]	-1.403657	-0.281961040
[2,]	-1.766588	-0.187432831
[3,]	-1.632115	-0.217553741
[4,]	-3.480933	-0.031262284
[5,]	-1.311025	-0.314085629
[6,]	-1.738875	-0.193242439
[7,]	-6.387487	-0.001683896
[8,]	-5.155549	-0.005784006
[9,]	-1.949987	-0.153472891
[10,]	-2.619871	-0.075599181

```
> CDPost.NBML@posteriors[101:110, ]
```

```
          NDE      DE
[1,] -0.1508101458 -1.966191
[2,] -0.0001777517 -8.635212
[3,] -0.1303547941 -2.101965
[4,] -0.0440829473 -3.143643
[5,] -0.0150211078 -4.205800
[6,] -0.1360356679 -2.062085
[7,] -0.2038614743 -1.690514
[8,] -0.1300648059 -2.104050
[9,] -0.1664504938 -1.875128
[10,] -0.0271896482 -3.618483
```

Here the assumption of a Negative Binomial distribution with priors estimated by maximum likelihood gives an estimate of

```
[1] 0.2243512
```

as the proportion of differentially expressed counts in the simulated data, where in fact the proportion is known to be 0.1.

## 5 Results

We can ask for the top differentially expressed tags using the `topCounts` function.

```
> topCounts(CDPost.NBML, group = 2)
```

```
      name X1 X2 X3 X4 X5 X6 X7 X8 X9 X10 Likelihood
1 count_80  1  1  0  1  1 13 21  8  6  20  0.9996228
2 count_78  1  1  0  1  1  8 13  7  9  10  0.9992802
3 count_66  0  0  0  0  0 15 10  4  4  10  0.9989883
4 count_26 13  4 11  5  7  1  1  1  0  0  0.9989096
5 count_21  2  0  1  1  0 15 15  6  5  11  0.9986313
6 count_7   9  8  8  8  9  1  2  1  0  0  0.9983175
7 count_72  0  0  1  0  0  7  6  4  3  8  0.9973797
8 count_83 14  6  9  2  9  1  0  0  1  1  0.9972926
9 count_64  6  6  8 11  9  1  1  0  0  1  0.9966427
10 count_27  5  3  6  4  7  0  0  0  1  0  0.9948027
```

We can compare the accuracy of the methods by using the `getTPs` function to find the number of true positives. We then use this to plot the false positive rate.

```
> NBML.TPs <- getTPs(CDPost.NBML, group = 2, TPs = 1:100)
> Poi.TPs <- getTPs(CDPost.Poi, group = 2, TPs = 1:100)
```

Figure 1 shows that false positive rates for the bootstrapped Negative Binomial approach with maximum likelihood priors (red) are much lower than for the Poisson-Gamma conjugacy approach (blue). This approach is therefore significantly more accurate, although potentially computationally more intensive and thus slower than the Poisson-Gamma conjugacy.

Finally, we shut down the cluster (assuming it was started to begin with).




Figure 1: (Log) false positive rates for differentially expressed tag discovery in a simulated dataset using the Poisson-Gamma approach (blue) and the Negative Binomial approach (red).

```
> if (!is.null(c1)) stopCluster(c1)
```

## 6 More Complex Experimental Designs

To illustrate the way in which a model is specified for more complex experimental designs, we consider a factorial design containing eight libraries. Table 1 describes the experimental design in more detail. Samples 1, 2, 3 & 4 are from condition A while samples 5, 6, 7 & 8 are from condition B. However, samples 1, 2, 5 & 6 have also been subjected to some condition C, while samples 3, 4, 7 & 8 have been subjected to some condition D.

	Condition A	Condition B
Condition C	Samples 1, 2	Samples 5, 6
Condition D	Samples 3, 4	Samples 7, 8

Table 1: An example factorial design experiment in which samples 1 and 2 are subjected to experimental conditions A and C, samples 3 and 4 are subjected to conditions B and C, samples 5 and 6 are subjected to conditions A and C and samples 7 and 8 are subjected to conditions B and D.



We prepare the baySeq library and cluster as before.

```
> library(baySeq)
> if ("snow" %in% installed.packages()[, 1]) {
+   library(snow)
+   cl <- makeCluster(4, "SOCK")
+ } else cl <- NULL
```

We load a simulated data set corresponding to the factorial design described. The first hundred cases show differential expression caused by differences between condition A and condition B, while the second hundred cases show differential expression caused by differences between condition C and condition D.

```
> data(factData)
> data(factlibsizes)
```

We establish our replicate structure, together with three group structures on the data. We observe that in this case, the replicate structure does not correspond to any of the group structures.

```
> replicates <- c(1, 1, 2, 2, 3, 3, 4, 4)
> factgroups <- list(NDE = c(1, 1, 1, 1, 1, 1, 1, 1), DE.A.B = c(1,
+   1, 1, 1, 2, 2, 2, 2), DE.C.D = c(1, 1, 2, 2, 1, 1, 2, 2))
```

The first group assumes no differential expression between samples. The second group assumes differential expression between samples experiencing condition A and samples experiencing condition B. The third group assumes differential expression between samples experiencing condition C and samples experiencing condition D.

We could also consider the possibility of interactions between effects, by considering a group  $c(1,1,2,2,3,3,4,4)$ . However, in this simulated data set, no such data exists and so we need not consider this group. It should be noted, however, that such a group would only find that an interaction effect takes place in some elements of the data. Like an ANOVA test, it is necessary to examine the data to determine what form the effect takes.

Having established a group structure, we proceed as before.

```
> CDfact <- new("countData", data = factCount, replicates = replicates,
+   libsizes = factlibsizes, groups = factgroups)
> CDfact@annotation <- data.frame(name = paste("count", 1:1000,
+   sep = "_"))
> CDfactP.NBML <- getPriors.NB(CDfact, samplesize = 1000, estimation = "QL",
+   cl = cl)
> CDfactPost.NBML <- getLikelihoods.NB(CDfactP.NBML, pET = "BIC",
+   cl = cl)
.
> CDfactPost.NBML@estProps
```

[1] 0.5451340 0.2299319 0.2249341

We can then ask for the tags showing most differential expression caused by the difference between conditions A and B

```
> topCounts(CDfactPost.NBML, group = 2)

      name X1 X2 X3 X4 X5 X6 X7 X8 Likelihood
1 count_94 19 41 19 22  3  2  3  4  0.9998662
2 count_81  6 12  5 10  0  0  0  0  0.9993401
3 count_88  0  0  0  0 10  8  6  8  0.9992040
4 count_41  2  5  4  1 29 11 29 23  0.9990795
5 count_13  2  6  6  6 32 32 59 37  0.9988702
6 count_60  5 14  6  6  0  0  1  0  0.9988108
7 count_49  5 19 14 12  2  1  2  1  0.9979943
8 count_7   0  1  1  0 19 16  9  8  0.9977774
9 count_75  1  4  1  0 30  7 33 14  0.9966735
10 count_95  3 14 13  8  1  0  0  0  0.9966193
```

And for those tags showing most differential expression caused by the difference between conditions C and D

```
> topCounts(CDfactPost.NBML, group = 3)

      name X1 X2 X3 X4 X5 X6 X7 X8 Likelihood
1 count_138  3  1 21 23  2  1 29 26  0.9999568
2 count_126 53 78 10  3 47 39  7  6  0.9997064
3 count_161 15 27  1  1 10 15  1  0  0.9996563
4 count_200 18 16  0  1 11 10  1  2  0.9991862
5 count_180  1  2  6 11  0  1 14 15  0.9991334
6 count_155 10 14 40 56  5 12 81 73  0.9985078
7 count_125  2  1 11  8  2  0 19  8  0.9983540
8 count_105  0  3  8  5  0  0  9  8  0.9964886
9 count_166  6 18  1  0 24 15  0  3  0.9952872
10 count_186  9 58  1  3 24 15  3  4  0.9927474
```

## 7 Application to Genomic Regions

So far, we have assumed that the count data deals with individual tags from a sequencing machine. If we consider count data derived from grouping together tags relating to genomic regions, for example, in looking at the number of tags that match to a gene in mRNA-Seq, we need to include the region length in the calculations of differential expression. The reasons for this are clear; if a region of length 200 bases has 200 tags in sample A and 400 tags in sample B, this may be good evidence for differential expression. If we saw the same difference in a region 2e6 bases long, this would be much poorer evidence of differential expression. We can specify values for the '@seglens' slot of the 'countData' class to explore such data.

We load a simulated data set consisting of count data on one thousand counts and library sizes for ten libraries. We use the same library sizes as before, but now the first column of 'simSeg' contains the length of the segment.

```

> data(simSeg)
> data(libsizes)
> simSeg[1:10, ]

      length
[1,]  6385  150   72   87   61  340  16  28  31   6  45
[2,] 15078  455  515 1670  120  449  34 133  34  32  63
[3,]  5577   20   5   11   12   12 106 256  93  47 152
[4,] 25139  808  529  891 1114 2705 135  89  41  37 312
[5,] 16423   86   22   36   15   13 344 802 120 462 857
[6,] 21752 1675  503 1007  193 1033  28   0  10 177 146
[7,]   434   50   19   29   12   59 196 780  50 214 510
[8,] 15370   55   45   75   27   77 210 426 137 216 260
[9,] 43786 2581 1245 2234  940  982 128 644  53 199 257
[10,] 11835   56   26   40   51   55 263 424 441 231 286

> libsizes

[1] 75373 40153 75403 34285 55975 53287 80477 37655 41171 77510

```

As before, the data are simulated such that the first hundred segments show differential expression between the first five libraries and the second five libraries. We thus establish two groups and a replicate structure as before.

```

> replicates <- c(1, 1, 1, 1, 1, 2, 2, 2, 2, 2)
> groups <- list(NDE = c(1, 1, 1, 1, 1, 1, 1, 1, 1, 1), DE = c(1,
+   1, 1, 1, 1, 2, 2, 2, 2, 2))

```

We now combine the count data, library sizes and user-defined groups into a `countData` object as before, but now we additionally specify segment lengths by setting the `@seglens` slot.

```

> SD <- new("countData", data = simSeg[, -1], replicates = replicates,
+   libsizes = libsizes, groups = groups, seglens = simSeg[,
+   1])

```

We can also optionally add annotation details into the `@annotation` slot of the `countData` object.

```

> SD@annotation <- data.frame(name = paste("gene", 1:1000, sep = "_"))

```

We can calculate the priors for this `countData` object as before, for both the Poisson-Gamma method and the Negative Binomial method.

```

> SDP.NBML <- getPriors.NB(SD, samplesize = 1000, estimation = "QL",
+   cl = cl)
> SDP.Pois <- getPriors.Pois(SD, samplesize = 20, cl = cl)

```

We can then calculate the posterior likelihoods as before.

```

> SDPost.Pois <- getLikelihoods.Pois(SDP.Pois, pET = "BIC", cl = cl)
> SDPost.NBML <- getLikelihoods.NB(SDP.NBML, pET = "BIC", cl = cl)

```

If we ignore segment length, and merely use the count data, we could acquire posteriors as before.

```
> CSD <- new("countData", data = simSeg[, -1], replicates = replicates,
+   libsizes = libsizes, groups = groups)
> CSD@annotation <- data.frame(name = paste("gene", 1:1000, sep = "_"))
> CSDP.NBML <- getPriors.NB(CSD, samplesize = 1000, estimation = "QL",
+   cl = cl)
> CSDPost.NBML <- getLikelihoods.NB(CSDP.NBML, pET = "BIC", cl = cl)
```

```
> CSDP.Pois <- getPriors.Pois(CSD, samplesize = 20, cl = cl)
> CSDPost.Pois <- getLikelihoods.Pois(CSDP.Pois, pET = "BIC", cl = cl)
```

If we look at the false discovery rates (Figure 2), we find that the Negative Binomial method (red) handles data with different segment lengths much better than the Poisson-Gamma method (blue). We also see that, for the Negative Binomial method, ignoring segment lengths can have a small negative effect on the results.




Figure 2: (Log) false positive rates for differentially expressed region discovery in a simulated dataset using both the Negative Binomial and Poisson-Gamma methods, including and ignoring segment length data.

We are also, using the Negative Binomial methods, able to ask if any segments have no true expression by setting `nullData = TRUE`. If the distribution of rates at each segment is bimodal, this approach will be able to differentiate between segments which have no true expression of any kind, and those which have expression (and may be differentially expressed). If we use this option, we implicitly introduce another grouping. The `prs` vector, which defines the prior likelihood of each grouping, should thus sum to less than 1, with the residual going to the implicit group. We can use this method whether or not segment lengths are specified; however, to ignore segment length may have a significant effect on the results.

In the `simSeg` object, the second hundred rows are simulated to show no true expression.

```
> NSDPost.NBML <- getLikelihoods.NB(SDP.NBML, pET = "BIC", nullData = TRUE,
+   bootStraps = 1, cl = cl)
.
> NCSDPost.NBML <- getLikelihoods.NB(CSDP.NBML, pET = "BIC", nullData = TRUE,
+   bootStraps = 1, cl = cl)
.
```

We can see the top segments for which there is no true expression by setting `group = NULL` in the `topCounts` function. Similarly, we can find true positive rates for by setting `group = NULL` in the `getTPs` function. Figure 3 shows that neglecting the segment lengths causes a substantial difference in false positive rates in the analysis of the data neglecting segment lengths (blue) compared to the analysis of the data including segment length information (red) in the identification of segments for which there is no true expression.

```
> topCounts(NSDPost.NBML, group = NULL)

      name V1 V2 V3 V4 V5 V6 V7 V8 V9 V10 Likelihood
1 gene_139 10  6 13  5  9 10 13  5 11  19  1.0000000
2 gene_106 32 14 22  5 14 19 21 15 15  19  1.0000000
3 gene_194 17 11 16 10 13  7 27 14  4  22  0.9999999
4 gene_169 24  7  8  8 17 16 16 15 13  13  0.9999999
5 gene_108 12  8 14 11  8  9 13  6  9  7  0.9999999
6 gene_199 29 14 30 19 25  7 29 20 10  26  0.9999998
7 gene_181 14  8  6  4  7  7 10  4  2  11  0.9999997
8 gene_101 16  5  5  6  5  5 12  2  2  10  0.9999997
9 gene_115  9  5 22  7 20  8 16 10 14  26  0.9999997
10 gene_168 11  5 15  6  4 12 16 10 11  19  0.9999996

> topCounts(NCSDPost.NBML, group = NULL)

      name V1 V2 V3 V4 V5 V6 V7 V8 V9 V10 Likelihood
1 gene_128  5  1  5  1  3  5  4  1  3  3  1
2 gene_234  6  1  5  2  4  2  3  3  2  8  1
3 gene_218  6  2  4  0  3  2  5  3  4  5  1
4 gene_548  5  1  3  3  5  6  2  2  2  5  1
```

5	gene_557	4	4	11	3	5	7	6	2	1	7	1
6	gene_890	2	3	4	4	5	3	11	3	1	7	1
7	gene_735	4	2	6	5	8	3	6	4	7	7	1
8	gene_103	8	3	2	3	4	9	8	2	5	4	1
9	gene_193	2	1	1	0	1	3	3	0	1	3	1
10	gene_449	4	3	9	3	3	2	9	2	0	4	1

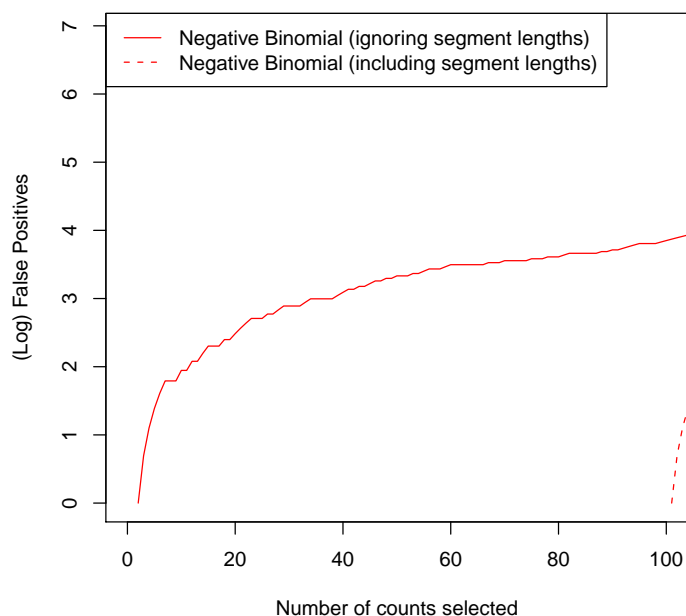


Figure 3: (Log) false positive rates for discovery of regions with no genuine expression in a simulated dataset using the the Negative Binomial approach considering region lengths (solid) and neglecting region lengths (dashed).

In order to look for segments for which there is no true expression, we assume that the rates of production of tags are bimodally distributed. If this is not the case, then this approach may cause serious problems in the analysis of the data. It is possible to gain some idea of whether the rates are bimodally distributed by examining the priors estimated (by Negative Binomial methods) for a `countData` object by using the function `plotPriors` for some group; ideally, the group which defines non-differentially expressed data. Figure 4 shows that these data are truly bimodal and thus this is a sensible approach.

```
> plotPriors(SDP.NBML, group = 1)
```




Figure 4: Density of the (log) values estimated for the mean of the data by Negative Binomial methods. The data are bimodally distributed, suggesting that some regions have no true expression.

## References

- [1] Thomas J. Hardcastle and Krystyna A. Kelly. *Empirical Bayesian methods for differential expression in count data*. In submission. 2009.
- [2] Mark Robinson `edgeR`: *Methods for differential expression in digital gene expression datasets*. Bioconductor.