

Advanced analysis using baySeq; generic distribution definitions

Thomas J. Hardcastle

April 16, 2015

1 Generic Prior Distributions

baySeq now offers complete user-specification of underlying distributions. This vignette describes using baySeq under this protocol. Familiarity with standard (negative-binomial) baySeq is assumed; please consult the other vignettes for a description of this approach.

Analysis is carried out through specification of a `densityFunction` class. The primary value in a `densityFunction` object is the `@density` slot, a user-defined function that should take variables `x`, `observables` and `parameters`. `x` corresponds to a row of data in a `countData` object. `observables` is a list object of observed values that may influence the model. By default, the `@libsizes` and `@seglens` values of the `countData` object will be internally appended to this list, unless objects with these names are otherwise specified by the user. `parameters` is a list object of parameters to be empirically estimated from the data with the `getPriors` function and used to estimate likelihoods with the `getLikelihoods` function. The `@dist` function should return a vector of log-likelihood values (or NA for invalid parameter choices) of the same length as the input variable `x`.

Other required slots of the `densityFunction` object are `initiatingValues`, a vector of initiating values to be used in optimisation of the parameters to be used in the `@dist` slot (and thus defining the length of the parameter object) and `equalOverReplicates`, a specification of which parameters are fixed for every replicate group and which may vary for different experimental conditions. If only one parameter is variable over experimental conditions, the Nelder-Mead optimisation used may be unstable, and one-dimensional optimisation with user defined functionally specified lower and upper bounds may (optionally) be provided; otherwise, Nelder-Mead will be attempted.

Optionally a function may be provided in `@stratifyFunction` to stratify the data and improve prior estimation in the tails where the `samplesize` argument in the `getPriors` function is less than the row dimension of the `countData` object. If this is provided, the `@stratifyBreaks` slot should give the number of strata to be used.

Below a model is constructed based on the normal distribution. The standard deviation is assumed to be constant for a given row of data across all experimental conditions, while the means (normalised by library scaling factor) are allowed to vary across experimental conditions.

If parallelisation is available, it is useful to use it.

```
> if(require("parallel")) cl <- makeCluster(8) else cl <- NULL
> library(baySeq)
> normDensityFunction <- function(x, observables, parameters) {
+   if(any(sapply(parameters, function(x) any(x < 0)))) return(rep(NA, length(x)))
+   dnorm(x, mean = parameters[[2]] * observables$libsizes, sd = parameters[[1]], log = TRUE)
+ }
> normDensity <- new("densityFunction", density = normDensityFunction, initiatingValues = c(0.1, 1),
+   equalOverReplicates = c(TRUE, FALSE),
+   lower = function(x) 0, upper = function(x) 1 + max(x) * 2,
+   stratifyFunction = rowMeans, stratifyBreaks = 10)
```

We construct the `countData` object as before.

```
> data(simData)
> CD <- new("countData", data = simData,
+   replicates = c("simA", "simA", "simA", "simA", "simA",
+   "simB", "simB", "simB", "simB", "simB"),
```

```
+      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))
+    )
> libsizes(CD) <- getLibsizes(CD)
> densityFunction(CD) <- normDensity
```

We can then fit priors and calculate posterior likelihoods based on our specified distributional model. The distributional model is specified in the 'getPriors' function and will be automatically used in the 'getLikelihoods' function

```
> normCD <- getPriors(CD, cl = cl)
> normCD <- getLikelihoods(normCD, cl = cl)
```

.

Similarly, we can construct a generic version of the negative-binomial model.

```
> nbinomDensityFunction <- function(x, observables, parameters) {
+   if(any(sapply(parameters, function(x) any(x < 0)))) return(NA)
+   dnbinom(x, mu = parameters[[1]] * observables$libsizes * observables$seglens, size = 1 / parameters
+ }
> densityFunction(CD) <- new("densityFunction", density = nbinomDensityFunction, initiatingValues = c(0,
+   equalOverReplicates = c(FALSE, TRUE),
+   lower = function(x) 0, upper = function(x) 1 + max(x) * 2,
+   stratifyFunction = rowMeans, stratifyBreaks = 10)
> nbCD <- getPriors(CD, cl = cl)
> nbCD <- getLikelihoods(nbCD, cl = cl)
```

.

We can compare this to the standard analysis of these data.

```
> CD <- getPriors.NB(CD, cl = cl)
> CD <- getLikelihoods(CD, cl = cl)
```

.

The generic negative-binomial data performs almost identically to standard baySeq. The methods differ in that the standard baySeq uses quasi-maximum-likelihood to estimate the priors, while generic baySeq uses maximum-likelihood (since no generic method exists for quasi-maximum-likelihood on arbitrary distributions).

2 Paired Data Analysis

We illustrate the possibilities of 'null' data, in which two separate models are applied to data equivalently expressed across all samples. The process for analysing paired data follows approximately the same steps as for analysing unpaired data, however, two different types of differential expression can exist within paired data. Firstly, we can find differential expression between replicate groups, as before. However, we can also find (consistent) differential expression between pairs; this would occur when for a single row of data, the first member of each pair differs from the second member of each pair. baySeq can identify both these types of differential expression simultaneously, and we implement this procedure below.

We begin by loading a simulated dataset containing counts for four paired datasets.

```
> data(pairData)
```

The first four columns in these data are paired with the second four columns. We construct a count data containing paired data in a similar fashion to the countData object. Note that the data are now three dimensional; for each row and each sample there are two observations.

```
> pairCD <- new("countData", data = array(c(pairData[,1:4], pairData[,5:8]), dim = c(nrow(pairData), 4,
+   replicates = c(1,1,2,2),
+   groups = list(NDE = c(1,1,1,1), DE = c(1,1,2,2)),
+   densityFunction = bbDensity)
```

We can find the library sizes for the data with the getLibsizes function.

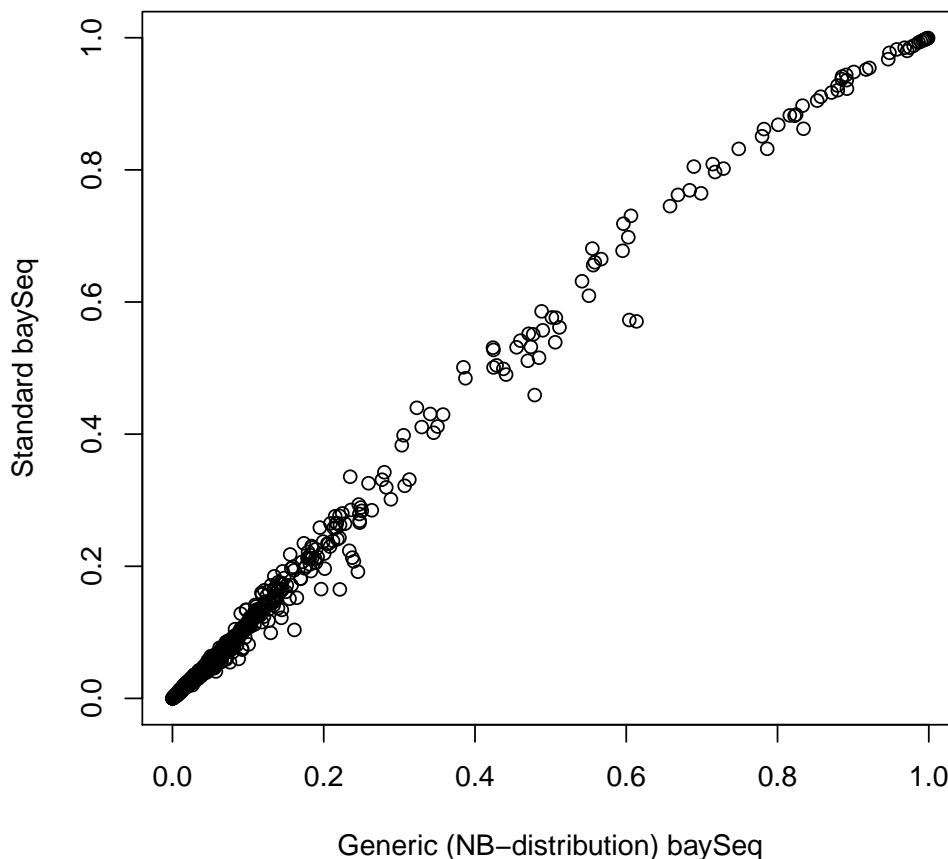


Figure 1: Likelihoods of DE estimated by standard/generic baySeq"

```
> libsizes(pairCD) <- getLibsizes(pairCD)
```

We estimate an empirical distribution on the parameters of a beta-binomial distribution by bootstrapping from the data, taking individual counts and finding the maximum likelihood parameters for a beta-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.

```
> pairCD <- getPriors(pairCD, samplesize = 1000, cl = cl)
```

We then acquire posterior likelihoods as before. The use of 'nullData = TRUE' in this context allows us to identify pairs which show no differential expression between replicate groups, but does show deviation from a one-to-one ratio of data between pairs.

```
> pairCD <- getLikelihoods(pairCD, pET = 'BIC', nullData = TRUE, cl = cl)
```

```
.
```

We can ask for the top candidates for differential expression between replicate groups using the topCounts function as before.

```
> topCounts(pairCD, group = 2)
```

	rowID	X1.1	X1.2	X2.1	X2.2	Likelihood	ordering	FDR.DE	FWER.DE
1	5	159:73	44:24	0:49	0:68	0.9978820	1>2	0.002118030	0.002118030
2	35	53:12	19:7	0:77	0:6	0.9949682	1>2	0.003574908	0.007139159

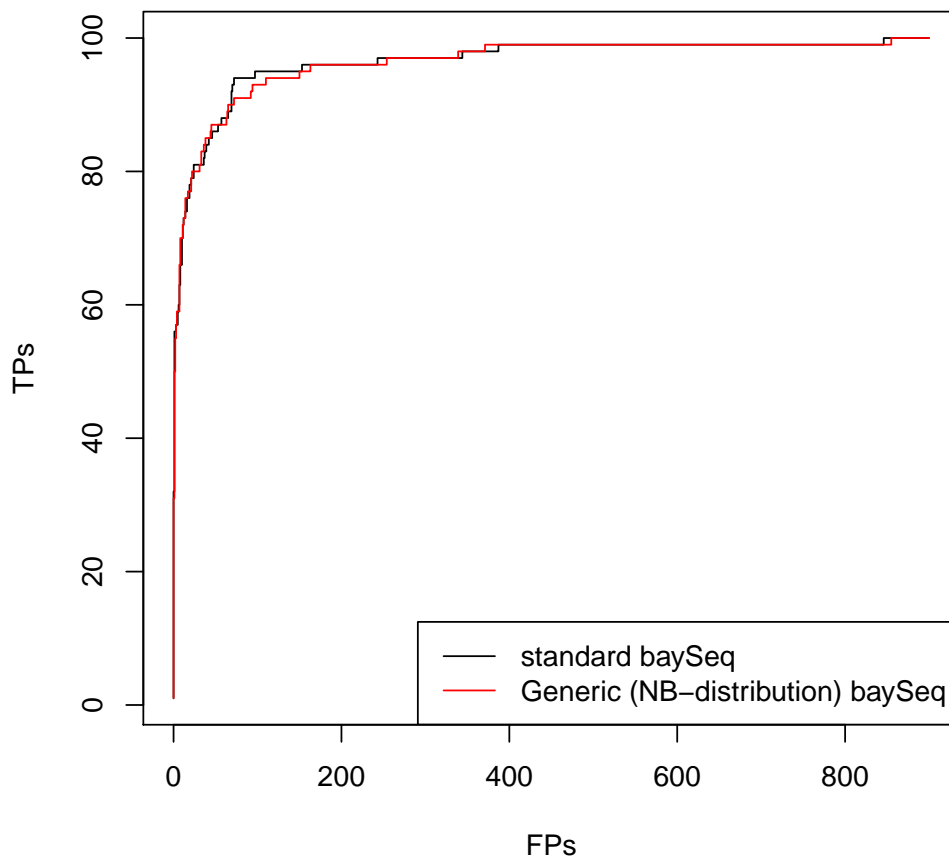


Figure 2: ROC curves estimated by standard/generic baySeq"

3	53	709:0	895:0	373:191	124:60	0.9934156	1>2	0.004578084	0.013676587
4	96	25:0	73:0	8:3	36:13	0.9907188	1>2	0.005753873	0.022830893
5	65	80:0	48:0	36:50	12:3	0.9877494	1>2	0.007053218	0.034801800
6	24	63:0	21:0	47:80	6:13	0.9850613	1>2	0.008367472	0.049220643
7	90	268:0	39:0	74:107	98:36	0.9808500	1>2	0.009907831	0.067428052
8	68	123:63	38:36	1198:179	350:18	0.9660930	2>1	0.012907726	0.099048764
9	71	8:0	15:0	21:16	2:1	0.9655855	1>2	0.015297370	0.130054569
10	50	43:19	44:46	106:6	133:5	0.9649216	2>1	0.017275471	0.160570848

However, we can also look for consistent differential expression between the pairs.

```
> topCounts(pairCD, group = 1)
```

	rowID	X1.1	X1.2	X2.1	X2.2	Likelihood	FDR.NDE	FWER.NDE
1	116	17:70	1:40	9:117	3:45	0.9924393	0.007560749	0.007560749
2	166	1027:27	835:8	1155:29	138:0	0.9900925	0.008734133	0.017393357
3	146	1:38	0:68	0:28	0:26	0.9860240	0.010481413	0.031126241
4	123	1:4	1:11	0:5	1:14	0.9827331	0.012177797	0.047855736
5	180	1:2	1:16	2:41	0:2	0.9787196	0.013998313	0.068117725
6	193	69:1	10:1	119:17	53:5	0.9781786	0.015302159	0.088452689
7	138	0:12	0:4	0:4	0:13	0.9659482	0.017980678	0.119492512
8	101	0:30	0:5	0:60	0:24	0.9657387	0.020015756	0.149659841
9	144	0:4	0:21	0:2	0:12	0.9656936	0.021603601	0.178831921
10	127	0:3	0:12	0:15	0:4	0.9656223	0.022881006	0.207061749

3 Different Model Priors

It is now possible to use different model priors for different subsets of the countData object. If we expect a certain class of genes (for example) to have a different prior likelihood towards differential expression than another such class, we can separate the two sets and estimate (or set) the model priors independently.

Let us suppose that we have reason to believe that the first hundred genes in the 'CD' object are likely to behave differently to the remaining genes. Then

```
> CDv <- getLikelihoods(nbCD, modelPriorSets = list(A = 1:100, B = 101:1000), cl = cl)
```

.

The model priors used are recorded in the @priorModels slot.

```
> CDv@priorModels
```

```
$A
```

```
NDE DE
```

```
0.2 0.8
```

```
$B
```

```
NDE DE
```

```
0.97 0.03
```

We can see the difference in performance by computing the ROC curves as before. Using different model priors can substantially improve performance, although obviously we have cheated here by splitting exactly those data simulated as DE and those as none-DE. It should also be recognised that this approach may bias downstream analyses; e.g. GO enrichment analysis.

Several pre-existing distributions are built into baySeq. Here we use a pre-developed zero-inflated negative binomial distribution to analyse zero-inflated data.

```
> data(zimData)
> CD <- new("countData", data = zimData,
+         replicates = c("simA", "simA", "simA", "simA", "simA",
+         "simB", "simB", "simB", "simB", "simB"),
+         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))
+         )
> libsizes(CD) <- getLibsizes(CD)
> densityFunction(CD) <- nbinomDensity
> CD <- getPriors(CD, cl = cl)
> CD <- getLikelihoods(CD, cl = cl)
.
> CDz <- CD
> densityFunction(CDz) <- ZINBDensity
> CDz <- getPriors(CDz, cl = cl)
> CDz <- getLikelihoods(CDz, cl = cl)
.
```

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

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

Session Info

```
> sessionInfo()
```

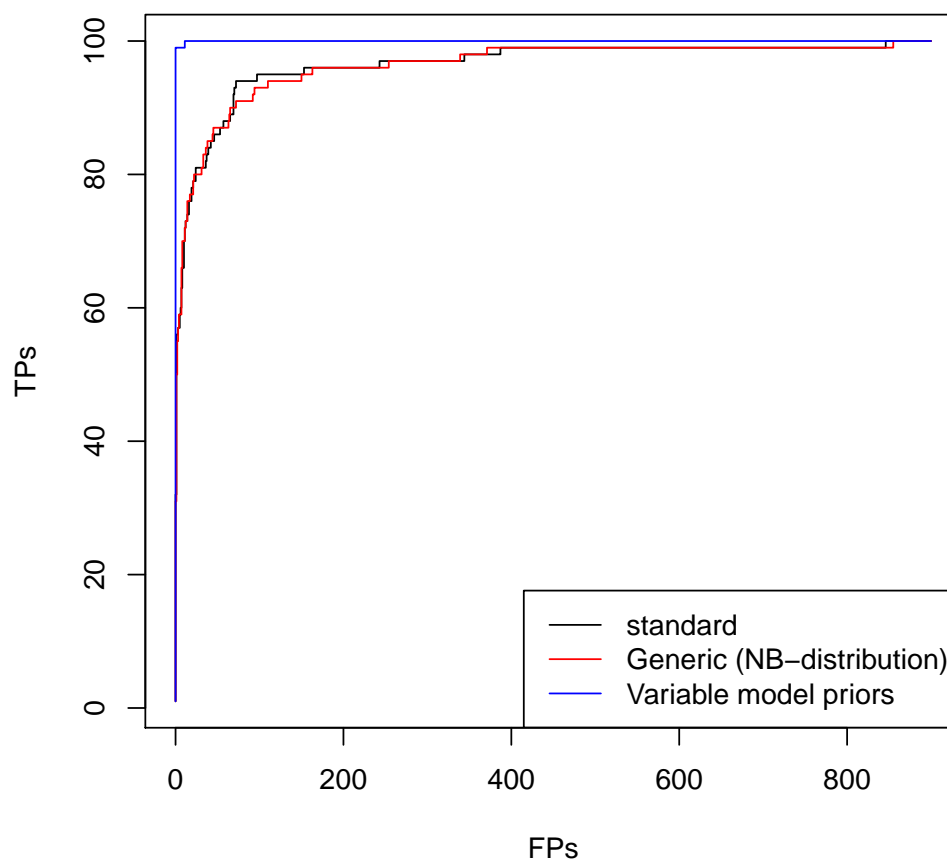


Figure 3: ROC curves estimated by standard/generic/variable model priors baySeq”

```
R version 3.2.0 RC (2015-04-08 r68161)
Platform: x86_64-apple-darwin10.8.0 (64-bit)
Running under: OS X 10.6.8 (Snow Leopard)
```

```
locale:
```

```
[1] en_US.UTF-8/en_US.UTF-8/en_US.UTF-8/C/en_US.UTF-8/en_US.UTF-8
```

```
attached base packages:
```

```
[1] stats4    parallel  stats      graphics  grDevices  utils      datasets  methods
[9] base
```

```
other attached packages:
```

```
[1] baySeq_2.2.0      perm_1.0-0.0      abind_1.4-3       GenomicRanges_1.20.0
[5] GenomeInfoDb_1.4.0 IRanges_2.2.0     S4Vectors_0.6.0   BiocGenerics_0.14.0
```

```
loaded via a namespace (and not attached):
```

```
[1] BiocStyle_1.6.0 XVector_0.8.0     tools_3.2.0
```