

Advanced analysis using baySeq; generic distribution definitions

Thomas J. Hardcastle

May 29, 2026

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(4) 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[[2]]
+ }
> densityFunction(CD) <- new("densityFunction", density = nbinomDensityFunction, initiatingValues = c(0.1,
+   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)
.

```



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

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).



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

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, 2)),
+      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.

```
> 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)
```

	X1.1	X1.2	X2.1	X2.2	likes	DE	FDR.DE	FWER.DE
NA	159:73	44:24	0:49	0:68	0.9978810	1>2	0.002118967	0.002118967
NA.1	53:12	19:7	0:77	0:6	0.9949700	1>2	0.003574459	0.007138260
NA.2	709:0	895:0	373:191	124:60	0.9934377	1>2	0.004570396	0.013653686
NA.3	25:0	73:0	8:3	36:13	0.9907158	1>2	0.005748848	0.022811128
NA.4	80:0	48:0	36:50	12:3	0.9877493	1>2	0.007049215	0.034782360
NA.5	63:0	21:0	47:80	6:13	0.9850684	1>2	0.008362950	0.049194626
NA.6	268:0	39:0	74:107	98:36	0.9808504	1>2	0.009903899	0.067402163
NA.7	123:63	38:36	1198:179	350:18	0.9661297	2>1	0.012899703	0.098989559
NA.8	8:0	15:0	21:16	2:1	0.9655804	1>2	0.015290807	0.130002012
NA.9	43:19	44:46	106:6	133:5	0.9649350	2>1	0.017268230	0.160508525

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

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

	X1.1	X1.2	X2.1	X2.2	likes	FDR.NDE	FWER.NDE
NA	17:70	1:40	9:117	3:45	0.9924370	0.007563041	0.007563041
NA.1	1027:27	835:8	1155:29	138:0	0.9900925	0.008735279	0.017395626
NA.2	1:38	0:68	0:28	0:26	0.9860246	0.010481969	0.031127867
NA.3	1:4	1:11	0:5	1:14	0.9827342	0.012177921	0.047856196
NA.4	1:2	1:16	2:41	0:2	0.9787194	0.013998455	0.068118378
NA.5	69:1	10:1	119:17	53:5	0.9781863	0.015300994	0.088446157

NA.6	0:12	0:4	0:4	0:13	0.9659482	0.017979683	0.119486223
NA.7	0:30	0:5	0:60	0:24	0.9657384	0.020014926	0.149654051
NA.8	0:4	0:21	0:2	0:12	0.9656937	0.021602858	0.178826286
NA.9	0:3	0:12	0:15	0:4	0.9656224	0.022880330	0.207056253

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

```
> # FAILS Bioc 3.17
> 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
```

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.

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

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) <- getLsizes(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()
```

R version 4.6.0 (2026-04-24)

Platform: x86_64-pc-linux-gnu

Running under: Ubuntu 24.04.4 LTS

Matrix products: default

BLAS: /usr/lib/x86_64-linux-gnu/openblas-pthread/libblas.so.3

LAPACK: /usr/lib/x86_64-linux-gnu/openblas-pthread/libopenblas-p-r0.3.26.so; LAPACK version 3.12.0

locale:

[1] LC_CTYPE=en_US.UTF-8	LC_NUMERIC=C	LC_TIME=en_US.UTF-8
[4] LC_COLLATE=en_US.UTF-8	LC_MONETARY=en_US.UTF-8	LC_MESSAGES=en_US.UTF-8
[7] LC_PAPER=en_US.UTF-8	LC_NAME=C	LC_ADDRESS=C
[10] LC_TELEPHONE=C	LC_MEASUREMENT=en_US.UTF-8	LC_IDENTIFICATION=C

time zone: Etc/UTC

tzcode source: system (glibc)

attached base packages:

[1] parallel stats graphics grDevices utils datasets methods base

other attached packages:

[1] baySeq_2.47.0

loaded via a namespace (and not attached):

[1] cli_3.6.6	knitr_1.51	rlang_1.2.0	xfun_0.57
[5] generics_0.1.4	S4Vectors_0.51.3	statmod_1.5.2	buildtools_1.0.0
[9] BiocStyle_2.41.0	htmltools_0.5.9	maketools_1.3.2	sys_3.4.3
[13] stats4_4.6.0	locfit_1.5-9.12	rmarkdown_2.31	grid_4.6.0
[17] Seqinfo_1.3.0	evaluate_1.0.5	abind_1.4-8	fastmap_1.2.0
[21] yaml_2.3.12	IRanges_2.47.2	BiocManager_1.30.27	compiler_4.6.0
[25] limma_3.69.1	edgeR_4.11.1	lattice_0.22-9	digest_0.6.39
[29] GenomicRanges_1.65.0	tools_4.6.0	BiocGenerics_0.59.6	