

Clustering high-throughput sequencing data based on patterns of co-expression

Thomas J. Hardcastle

October 30, 2018

1 Introduction

This vignette outlines two possibilities for clustering gene expression based on patterns of co-expression between samples, optionally accounting for the replicate structure of the data.

The package can be installed as

```
> if (!requireNamespace("BiocManager", quietly=TRUE))
+   install.packages("BiocManager")
> BiocManager::install("clusterSeq")
```

```
> library(clusterSeq)
```

We demonstrate these analyses on a set of time series data in female rat thymus tissues taken from the Rat Bodymap project [1]. By identifying clusters of genes that demonstrate similar patterns of expression over time we can identify patterns of time dependence within the data.

We first load in the processed data of observed read counts at each gene for each sample.

```
> data(ratThymus, package = "clusterSeq")
> head(ratThymus)
```

We define the replicate structure of the data in a vector whose members correspond to the columns of the data matrix.

```
> replicates <- c("2week", "2week", "2week", "2week",
+               "6week", "6week", "6week", "6week",
+               "21week", "21week", "21week", "21week",
+               "104week", "104week", "104week", "104week")
```

Library scaling factors are acquired here using the `baySeq::getLibsizes` function but might be acquired through any other means.

```
> library(baySeq)
> libsizes <- getLibsizes(data = ratThymus)
```

2 K-means based clustering

For k-means based clustering, we assume an approximately log-normal distribution. We adjust the data to remove zeros and rescale by the library scaling factors.

```
> ratThymus[ratThymus == 0] <- 1
> normRT <- log2(ratThymus %*% diag(1/libsizes) * mean(libsizes))
```

For speed purposes, we will consider only the first 1000 genes in the data.

```
> normRT <- normRT[1:1000,]
```

K-means based clustering compares two genes by separately clustering the expression values of each gene using all possible values of k (or, if the sample size is exceptionally large, all values of k from 1 to some large maximum permitted value.). The genes are then compared to find the maximum k for which the average expression of the identified clusters is monotonic between the two genes. For this value of k , the the maximum difference between expression levels observed within a cluster of either gene is reported as a measure of the dissimilarity between the two genes.

However, a problem arises for genes that are non-differentially expressed across all samples. For these genes, the appropriate number of clusters is 1; however, choices of $k > 1$ often lead to false assignment of genes to clusters exhibiting differential expression. To resolve this, the clustering incorporates a bootstrapping stage based on Tibshirani's gap statistic. Bootstrapping uniformly distributed data on the same range as the observed data, we calculate the the dissimilarity score as above, and find those cases for which the gap between the bootstrapped mean dissimilarity and the observed dissimilarity for $k = 1$ exceeds that for $k = 2$ by more than some multiple of the standard error of the bootstrapped dissimilarities of $k = 2$. These cases are forced to be treated as non-differentially expressed by discarding all dissimilarity data for $k > 1$.

We generate a matrix of dissimilarity scores between each pair of genes using the `kCluster` function. The full matrix can be reported to (gzipped) file; the function returns a data.frame which for each gene defines its nearest neighbour of higher row index, and the dissimilarity with that neighbour. This is sufficient for a clustering based on singleton agglomeration (see `hclust`).

```
> kClust <- kCluster(normRT, matrixFile = "kclust_matrix.txt.gz", B = 1000)
> head(kClust)
```

Using the output from 'kCluster' we construct a clustering of genes by singleton agglomeration of those genes with dissimilarity scores lower than a specified threshold.

```
> mkClust <- makeClusters(kClust, normRT, threshold = 1)
```

We can repeat this analysis forcing members of the same replicate groups to cluster together at the k-means stage.

```
> kClustR <- kCluster(normRT, replicates = replicates, matrixFile = "kclust_matrix_newForceReps.txt.gz", B = 1000)
> mkClustR <- makeClusters(kClustR, normRT, threshold = 1)
```

As an alternative to singleton agglomeration, we can use the full matrix and any agglomerative method defined in the `hclust` heirarchical analysis function.

```
> mkClustRC <- makeClustersFF("kclust_matrix_newForceReps.txt.gz", method = "complete", cut.height = 5)
```

3 Posterior likelihood analyses

As an alternative to the K-means based clustering, we can use empirical Bayesian likelihoods on all possible models for the structure of the data at each gene, and from these estimate the likelihood that two genes share the same structure. We estimate these likelihoods using the baySeq package.

We first create a countData object to contain the data and find all possible models that preserve the replicate structure.

```
> library(baySeq)
> cD.ratThymus <- new("countData", data = ratThymus, replicates = c(1,1,1,1,2,2,2,2,3,3,3,3,4,4,4,4))
> densityFunction(cD.ratThymus) <- nbinomDensity
> libsizes(cD.ratThymus) <- getLibsizes(cD.ratThymus)
> cD.ratThymus <- allModels(cD.ratThymus)
```

The commented out code is for the analysis of the data with baySeq. Since this takes some time, this has already been performed and can be loaded in with the analysis already performed.

```
> #cl <- makeCluster(4)
> #cD.ratThymus <- getPriors(cD.ratThymus, consensus = TRUE, cl = cl)
> #cD.ratThymus <- getLikelihoods(cD.ratThymus, cl = cl)
>
> data(cD.ratThymus, package = "clusterSeq")
```

Given the estimated posteriors for each gene to have each model, we estimate the likelihood that two genes have identical models, and that the ordering of average expression in the groups defined by those models is monotonic between the two genes. As before, the function associatePosteriors returns a data.frame which for each gene defines its nearest neighbour of higher row index, and the dissimilarity with that neighbour. This is sufficient for a clustering based on singleton agglomeration using the makeClusters function.

```
> cD.ratThymus <- cD.ratThymus[1:1000,]
> aM <- associatePosteriors(cD.ratThymus)
> sX <- makeClusters(aM, cD.ratThymus, threshold = 0.5)
```

Likelihoods of a pair being clustered together in one clustering given that they are clustered in that way in the other.

```
> wallace(sX, mkClustRC)

P.v1_given_v2 P.v2_given_v1
1.0000000      0.2787628
```

```
> par(mfrow = c(2,3))
> plotCluster(sX[1:6], cD.ratThymus)
```

```
«label=plotClusterBS-fig, fig=TRUE,echo=FALSE»= «plotClusterBS» @
```

Figure 1: Profiles of gene expression in top six clusters acquired from baySeq analysis

Session Info

```
> sessionInfo()

R version 3.5.1 Patched (2018-07-12 r74967)
Platform: x86_64-pc-linux-gnu (64-bit)
Running under: Ubuntu 16.04.5 LTS

Matrix products: default
BLAS: /home/biocbuild/bbs-3.8-bioc/R/lib/libRblas.so
LAPACK: /home/biocbuild/bbs-3.8-bioc/R/lib/libRlapack.so

locale:
 [1] LC_CTYPE=en_US.UTF-8      LC_NUMERIC=C
 [3] LC_TIME=en_US.UTF-8      LC_COLLATE=C
 [5] LC_MONETARY=en_US.UTF-8  LC_MESSAGES=en_US.UTF-8
 [7] LC_PAPER=en_US.UTF-8    LC_NAME=C
 [9] LC_ADDRESS=C             LC_TELEPHONE=C
[11] LC_MEASUREMENT=en_US.UTF-8 LC_IDENTIFICATION=C

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

other attached packages:
[1] clusterSeq_1.6.0 baySeq_2.16.0 abind_1.4-5
[4] GenomicRanges_1.34.0 GenomeInfoDb_1.18.0 IRanges_2.16.0
[7] S4Vectors_0.20.0 BiocGenerics_0.28.0 BiocParallel_1.16.0

loaded via a namespace (and not attached):
 [1] Rcpp_0.12.19 edgeR_3.24.0 knitr_1.20
 [4] XVector_0.22.0 zlibbioc_1.28.0 lattice_0.20-35
 [7] tools_3.5.1 grid_3.5.1 htmltools_0.3.6
[10] yaml_2.2.0 rprojroot_1.3-2 digest_0.6.18
[13] GenomeInfoDbData_1.2.0 BiocManager_1.30.3 bitops_1.0-6
[16] RCurl_1.95-4.11 evaluate_0.12 rmarkdown_1.10
[19] limma_3.38.0 compiler_3.5.1 backports_1.1.2
[22] locfit_1.5-9.1 BiocStyle_2.10.0
```

References

- [1] Yu, Ying and Fuscoe, James C and Zhao, Chen and Guo, Chao and Jia, Meiwen and Qing, Tao and Bannon, Desmond I and Lancashire, Lee and Bao, Wenjun and Du, Tingting and Luo, Heng and Su, Zhenqiang and Jones, Wendell D and Moland, Carrie

L and Branham, William S and Qian, Feng and Ning, Baitang and Li, Yan and Hong, Huixiao and Guo, Lei and Mei, Nan and Shi, Tielu and Wang, Kevin Y and Wolfinger, Russell D and Nikolsky, Yuri and Walker, Stephen J and Duerksen-Hughes, Penelope and Mason, Christopher E and Tong, Weida and Thierry-Mieg, Jean and Thierry-Mieg, Danielle and Shi, Leming and Wang, Charles *A rat RNA-Seq transcriptomic BodyMap across 11 organs and 4 developmental stages*. Nature Communications (2014).