

subSeq Package Vignette (Version 1.38.0)

David G. Robinson and John D. Storey

April 15, 2025

Contents

1	Introduction	1
2	Quick Start Guide	2
2.1	The <code>subsamples</code> object	3
2.2	Summarizing Output by Depth	5
2.2.1	Estimating False Discovery Proportion	6
3	Plotting	7
4	Writing Your Own Analysis Method	10
4.1	Advanced Uses of Handlers	11
5	Reproducing or Adding to a Previous Run	11
5.1	Adding Methods or Depths	12
5.2	Examining a Matrix More Closely	13
6	Note on Subsampling	14

1 Introduction

This is a vignette for the `subSeq` package, which performs subsampling of sequencing data to test the effect of depth on your conclusions. When you use a RNA-Seq differential expression method, such as `edgeR` or `DESeq2`, you can answer a couple of biological questions:

1. What genes are differentially expressed?
2. What classes of genes show differential expression?

However, what if we're interested in questions of experimental design:

1. Do I have enough reads to detect most of the biologically relevant differences?

2. If I run a similar experiment, should I run additional lanes, or can I multiplex even more samples and work with fewer reads?

One way to help answer these questions is to pretend you have *fewer* reads than you do, and to see how your results (the number of significant genes, your estimates of their effects, and so on) change. If you can achieve the same results with just 10% of your reads, it indicates that (when using your particular analysis method to answer your particular question) the remaining 90% of the reads added very little. In turn, if your conclusions changed considerably between 80% and 100% of your reads, it is likely they would change more if you added additional reads.

`subSeq` is also designed to work with *any analysis method you want*, as long as it takes as input a matrix of count data per gene. `edgeR`, `DESeq2`, `voom` and `DEXSeq` are provided as default examples, but if you prefer to use a different package, to use these packages with a different set of options, or even to use your own method, it is very easy to do so and still take advantage of this package's subsampling methods. See [Writing Your Own Analysis Method](#) for more.

2 Quick Start Guide

We demonstrate the use of this method on a subset of the Hammer et al 2010 dataset [4]. This dataset (as provided by the ReCount database [3]), comes built in to the `subSeq` package:

```
library(subSeq)
data(hammer)
```

Alternatively we could have downloaded the RNA-Seq counts directly from the ReCount database:

```
load(url("http://bowtie-bio.sourceforge.net/recount/ExpressionSets/hammer_eset.RData"))
hammer = hammer.eset
```

We then filter such that no genes had fewer than 5 reads:

```
hammer.counts = Biobase::exprs(hammer)[, 1:4]
hammer.design = Biobase::pData(hammer)[1:4, ]
hammer.counts = hammer.counts[rowSums(hammer.counts) >=
5, ]
```

First we should decide which proportions to use. In general, it's a good idea to pick subsamplings on a logarithmic scale, since there tends to be a rapid increase in power for low read depths that then slows down for higher read depths.

```
proportions = 10^seq(-2, 0, 0.5)
proportions

## [1] 0.01000000 0.03162278 0.10000000 0.31622777
## [5] 1.00000000
```

The more proportions you use for subsampling, the greater your resolution and the better your understanding of the trend. However, it will also take longer to run. We give these proportions, along with the matrix and the design, to the `subsample` function, which performs the subsampling:

```
subsamples = subsample(hammer.counts, proportions,
  method = c("edgeR", "voomLimma"), treatment = hammer.design$protocol)
```

In this command, we are giving it a count matrix (`hammer.counts`), a vector of proportions to sample, and a vector describing the treatment design (`hammer.design$protocol`), and telling it to try two methods (edgeR and voom/limma) and report the results of each. (If there are a large number of proportions, this step may take a few minutes. If it is being run interactively, it will show a progress bar with an estimated time remaining).

2.1 The subsamples object

The `subsample` function returns a `subsamples` object. This also inherits the `data.table` class, as it is a table with one row *for each gene, in each subsample, with each method*:

```
options(width = 40)
subsamples

##           method proportion
##           <fctr>         <num>
##    1:      edgeR         0.01
##    2:      edgeR         0.01
##    3:      edgeR         0.01
##    4:      edgeR         0.01
##    5:      edgeR         0.01
##    ---
## 150206: voomLimma         1.00
## 150207: voomLimma         1.00
## 150208: voomLimma         1.00
## 150209: voomLimma         1.00
## 150210: voomLimma         1.00
##           replication coefficient
##           <int>         <num>
##    1:           1 0.0000000
```

```

##      2:      1  0.0000000
##      3:      1  0.0000000
##      4:      1  2.2920623
##      5:      1  0.0000000
##      ---
## 150206:      1 -0.4907982
## 150207:      1  0.9883726
## 150208:      1  2.7152540
## 150209:      1 -1.0073941
## 150210:      1 -0.5339148
##                pvalue                ID
##                <num>                <char>
##      1: 1.000000000 ENSRNOG000000000001
##      2: 1.000000000 ENSRNOG000000000007
##      3: 1.000000000 ENSRNOG000000000008
##      4: 1.000000000 ENSRNOG000000000010
##      5: 1.000000000 ENSRNOG000000000012
##      ---
## 150206: 0.010546926 ENSRNOG000000043503
## 150207: 0.126636226 ENSRNOG000000043504
## 150208: 0.002538513 ENSRNOG000000043507
## 150209: 0.103295778 ENSRNOG000000043512
## 150210: 0.012900077 ENSRNOG000000043513
##      count      depth      qvalue
##      <int>      <int>      <num>
##      1:      0    198283  1.000000000
##      2:      0    198283  1.000000000
##      3:      0    198283  1.000000000
##      4:      1    198283  1.000000000
##      5:      0    198283  1.000000000
##      ---
## 150206:  1404 19862441 0.009464488
## 150207:   17 19862441 0.065719919
## 150208:    6 19862441 0.003240602
## 150209:   12 19862441 0.056127854
## 150210:  430 19862441 0.011053328

```

The fields it contains are:

- **coefficient:** The estimated effect size at each gene. For differential expression between two conditions, this is the log2 fold-change.
- **pvalue:** The p-value estimated by the method
- **ID:** The ID of the gene, normally provided by the rownames of the count matrix

- **count**: The number of reads for *this particular gene* at this depth
- **depth**: The total depth (across all reads) at this level of sampling
- **replication**: Which subsampling replication (in this subsampling run, only one replicate was performed; use the **replications** argument to **subsample** to perform multiple replications)
- **method**: The name of the differential expression method used
- **qvalue**: A q-value estimated from the p-value distribution using the **qvalue** package. If you consider all genes with a q-value below q as candidates, you expect to achieve a false discovery rate of q .

2.2 Summarizing Output by Depth

This object gives the differential analysis results per gene, but we are probably more interested in the results per sampling depth: to see how the overall conclusions change at lower read depths. Performing **summary** on the subsamplings object shows this:

```
subsamples.summary = summary(subsamples)
subsamples.summary
```

##		depth	proportion	method
##		<int>	<num>	<char>
##	1:	198283	0.01000000	edgeR
##	2:	198283	0.01000000	voomLimma
##	3:	628394	0.03162278	edgeR
##	4:	628394	0.03162278	voomLimma
##	5:	1988324	0.10000000	edgeR
##	6:	1988324	0.10000000	voomLimma
##	7:	6277163	0.31622777	edgeR
##	8:	6277163	0.31622777	voomLimma
##	9:	19862441	1.00000000	edgeR
##	10:	19862441	1.00000000	voomLimma
##		replication	significant	pearson
##		<int>	<int>	<num>
##	1:	1	180	0.4453952
##	2:	1	0	0.4871767
##	3:	1	882	0.5561873
##	4:	1	0	0.6067309
##	5:	1	2436	0.6953136
##	6:	1	1489	0.7460535
##	7:	1	5083	0.8396533
##	8:	1	5029	0.8806891
##	9:	1	8258	1.0000000

```
## 10:          1          8614 1.0000000
##      spearman concordance      MSE
##      <num>      <num>      <num>
## 1: 0.4562906    0.3509033 2.4654393
## 2: 0.4794714    0.4769176 0.9324971
## 3: 0.5894677    0.4964425 1.6138540
## 4: 0.6098967    0.6024104 0.7321934
## 5: 0.7422582    0.6644645 0.9382205
## 6: 0.7543837    0.7454932 0.4690254
## 7: 0.8762327    0.8318243 0.4252754
## 8: 0.8798326    0.8806848 0.2224153
## 9: 1.0000000    1.0000000 0.0000000
## 10: 1.0000000    1.0000000 0.0000000
##      estFDP      rFDP      percent
##      <num>      <num>      <num>
## 1: 3.534231e-05 0.000000000 0.02542732
## 2: 0.000000e+00 0.000000000 0.00000000
## 3: 6.737492e-03 0.007936508 0.11183538
## 4: 0.000000e+00 0.000000000 0.00000000
## 5: 1.292033e-02 0.013136289 0.29584051
## 6: 1.801057e-02 0.021490934 0.17432400
## 7: 2.263099e-02 0.020263624 0.60407569
## 8: 3.722373e-02 0.037780871 0.56365754
## 9: 5.091755e-02 0.000000000 1.00000000
## 10: 5.043905e-02 0.000000000 1.00000000
```

As one example: the first row of this summary table indicates that at a read depth of 198283, edgeR found 180 genes significant (at an FDR of .05), and that its fold change estimates had a 0.4562906 Spearman correlation and a 0.4453952 Pearson correlation with the full experiment.

Note that different methods are compared *on the same subsampled matrix* for each depth. That is, at depth 198283, results from **edgeR** are being compared to results from **voomLimma** on the same matrix.

2.2.1 Estimating False Discovery Proportion

One question that **subSeq** can answer is whether decreasing your read depth ever *introduces* false positives. This would be a worrisome sign, as it suggests that *increasing* your read depth even more might prove that you are finding many false positives. **subSeq** has two ways of estimating the false discovery proportion at each depth, each of which requires choosing an FDR threshold to control for (default 5%):

- **eFDP**: The estimated false discovery proportion at each depth when controlling for FDR. This is found by calculating the local false discovery rate of each gene in the oracle, (which we will call the "oracle lFDR"), then

finding the mean oracle lFDR of the genes found significant at each depth. This is effectively using the best information available (the full, oracle experiment) to estimate how successful your FDR control is at each depth. This estimate will converge to the desired FDR threshold (e.g. 5%) at the full depth.

- **rFDP**: the "relative" false discovery proportion at each depth, where a relative false discovery is defined as a gene found significant at a subsampled depth that was not found significant at the full depth. This estimate will converge to 0% at the full depth.

Generally, we recommend the **eFDP** over the **rFDP** (and the **eFDP** is the default metric plotted), as a) it is a less noisy estimate, b) It converges to 5% rather than dropping back down to 0, and c) it takes into account *how* unlikely each hypothesis appeared to be in the oracle, rather than simply the binary question of whether it fell above a threshold. We report the **rFDP** only because it a simpler metric that does not require the estimation of local FDR.

In either case, it is important to note that the full experiment is not perfect and has false positives and negatives of its own, so these metrics should not be viewed as an absolute false discovery rate (i.e., "proof" that an experiment is correct). Instead, they examine whether decreasing read depth tends to introduce false results.

3 Plotting

The best way to understand your subsampling results intuitively is to plot them. You can do this by plotting the summary object (not the original subsamples). `plot(subsamples.summary)` creates Figure 1, which plots some useful metrics of quality by read depth.

Creating your own plot by read depth from the summary object is easy using `ggplot2`. For instance, if you would like to plot the percentage of oracle (full experiment) genes that were found significant at each depth, you could do:

```
library(ggplot2)
ggplot(subsamples.summary, aes(x = depth,
  y = percent, col = method)) + geom_line()
```






Figure 1: The default output of `plot(subsamples.summary)`. This shows four plots: (i) The number of significant genes found significant at each depth; (ii) the estimated false discovery at this depth, calculated as the average oracle local FDR of the genes found significant at this depth; (iii) The Spearman correlation of estimates at each depth with the estimates at the full experiment; (iv) the mean-squared error of the estimates at each depth with the estimates at the full experiment.



If you'd like to focus on the Pearson correlations of the estimates, you could do:

```
ggplot(subsamples.summary, aes(x = depth,
  y = pearson, col = method)) + geom_line()
```



4 Writing Your Own Analysis Method

There are five methods of RNA-Seq differential expression analysis built in to `subSeq`:

- `edgeR` Uses edgeR's exact negative binomial test to compare two conditions [6]
- `edgeR.glm` Uses edgeR's generalized linear model, which allows fitting a continuous variable and including covariates.
- `voomLimma` Uses `voom` from the `edgeR` package to normalize the data, then apply's limma's linear model with empirical Bayesian shrinkage [5].
- `DESeq2` Applies the negative binomial Wald test from the `DESeq2` package [1]. DESeq2 is notable for including shrinkage on the parameters.
- `DEXSeq` Uses a negative binomial generalized linear model to look for differential exon usage in mapped exon data (see ??) [2].

These handlers are provided so you can perform `subSeq` analyses "out-of-the-box". However, it is very likely that your RNA-Seq analysis does not fit into one of these methods! It may use different published software, include options you need to set manually, or it could be a method you wrote yourself. No problem: all you need to do is write a function, called a *handler*, that takes in a count matrix and any additional options, then performs your analysis and returns the result as a `data.frame`.

Start by writing up the code you would use to analyze a count matrix to produce (i) coefficients (such as the log fold-change) and (ii) p-values. But put your code inside a function, and have the function take as its first argument the (possibly subsampled) count matrix, and return the results as a two-column `data.frame` with one row per gene (in the same order as the matrix):

```
myMethod = function(count.matrix, treatment) {  
  # calculate your coefficients based on the input count matrix  
  coefficients = <some code>  
  # calculate your p-values  
  pvalues = <some more code>  
  
  # return them as a data.frame  
  return(data.frame(coefficient=coefficients, pvalue=pvalues))  
}
```

Your function can be as long or complicated as you want, as long as its first argument is a `count.matrix` and it returns a data frame with `coefficient` and `pvalue` columns.

Now that you've defined this function, you can pass its name to `subsample` just like you would one of the built-in ones:

```
subsamples = subsample(hammer.counts, proportions,
  method = c("edgeR", "DESeq2", "myMethod"),
  treatment = hammer.design$protocol)
```

4.1 Advanced Uses of Handlers

Your handler can do more than take these two arguments and calculate a p-value:

- Your handler can return more columns than a coefficient and p-value (for example, you can return a column of dispersion estimates, confidence interval boundaries, or anything else on which you would like to examine the effect of depth). These columns will be included in the output.

If you provide multiple methods, including one that returns an extra columns and one that doesn't, the column will be filled with NA for the method that doesn't provide the column.

- Your handler can take more arguments than just a count matrix and treatment vector. Any arguments passed at the end of the 'subsample' function will be passed on to it. *However*, you cannot use multiple methods where some take extra arguments and some do not. Instead, perform them separately (using the same seed; see [Reproducing or Adding to a Previous Run](#)) and then use `combineSubsamples`.
- Your handler does not necessarily have to return one row per gene. For example, you could have a handler that performs gene set enrichment analysis (GSEA) on your data, which then returns one gene set per row. If so, it must return an additional column: ID (those were otherwise acquired from the rownames of the count matrix). Note that `count` column (noting the read depth per gene) will become NA unless it is provided by the handler function.

5 Reproducing or Adding to a Previous Run

If you have a `subsamples` or `subsample` summary object that either you or someone else ran, you might be interested in working with the original data further. For example:

- You may want to perform additional methods, and be able to compare them directly to the original
- You may want to sample additional depths, and combine them together
- You may have noticed a strange discrepancy in one of the results, and want to further examine the count matrix at that specific depth

All of these are possible, as `subSeq` stores the random seed used for the run in the results.

5.1 Adding Methods or Depths


If you already have some methods performed (as we do in `subsamples`), we may want to add additional methods. For example you might want to analyze the same depths with the `voom` method. To do this, use the `getSeed` function to retrieve the random seed from the `subsamples` results, and then provide that seed to the `subsample` function:

```
seed = getSeed(subsamples)

subsamples.more = subsample(hammer.counts,
  proportions, method = c("voomLimma"),
  treatment = hammer.design$protocol, seed = seed)
```

After that, you can combine the two objects using `combineSubsamples`:

```
subsamples.combined = combineSubsamples(subsamples,
  subsamples.more)
plot(summary(subsamples.combined))
```



5.2 Examining a Matrix More Closely

Say that after your analysis, you are surprised that a particular method performed how it did at a particular depth, and you wish to examine that depth further.

You can also use the seed to retrieve the precise matrix used. For instance, if we want to find the matrix used for the .1 proportion subsampling:

```
submatrix = generateSubsampledMatrix(hammer.counts,
  0.1, seed = seed)
dim(submatrix)

## [1] 15021      4

sum(submatrix)
```

6 Note on Subsampling

subSeq performs read subsampling using a random draw from a binomial distribution. Specifically, for a subsampling proportion p_k , each value m, n in the subsampled matrix $Y^{(k)}$ is generated as:

$$Y_{m,n}^{(k)} \sim \text{Binom}(X_{m,n}, p_k)$$

This is equivalent to allowing each original mapped read to have probability p_k of being included in the new counts, as done, for example, by the Picard DownsampleSam function [7].

A computationally intensive alternative is to sample the reads *before* they are mapped, then to perform mapping on the sampled data. When the mapping is done independently and deterministically for each read, as it is e.g. in Bowtie, this is mathematically identical to subsampling the aligned reads or counts, since the inclusion of one read does not affect the mapping of any other. This applies to methods such as Bar-Seq and Tn-Seq, or for RNA-Seq in organisms with no or very few introns (e.g. bacteria or yeast). Note that some spliced read mappers, such as TopHat, do not perform all their mappings entirely independently, since unspliced mappings are first used to determine exonic regions before searching for spliced reads. However, those methods are (not coincidentally) among the most computationally intensive mappers in terms of running time and memory usage, giving subSeq an especially large advantage.

Several other papers that perform subsampling have made the assumption of independence implicitly, but if you are unsure, it may be useful to subsample the original fastq file, then perform the read mapping on each sample individually, and compare the subsampled counts to those from **subSeq** to check that this approximation is reasonable for your data.

References

- [1] Simon Anders and Wolfgang Huber. Differential expression analysis for sequence count data. *Genome Biology*, 11:R106, 2010.
- [2] Simon Anders, Alejandro Reyes, and Wolfgang Huber. Detecting differential usage of exons from RNA-seq data. *Genome research*, 22(10):2008–2017, October 2012.
- [3] Alyssa C Frazee, Ben Langmead, and Jeffrey T Leek. ReCount: a multi-experiment resource of analysis-ready RNA-seq gene count datasets. *BMC Bioinformatics*, 12:449, 2011.

- [4] Paul Hammer, Michaela S Banck, Ronny Amberg, Cheng Wang, Gabriele Petznick, Shujun Luo, Irina Khrebtukova, Gary P Schroth, Peter Beyerlein, and Andreas S Beutler. mRNA-seq with agnostic splice site discovery for nervous system transcriptomics tested in chronic pain. *Genome research*, 20(6):847–860, June 2010.
- [5] CW Law, Y Chen, W Shi, and GK Smyth. Voom: precision weights unlock linear model analysis tools for RNA-seq read counts. *Genome Biology*, January 2014.
- [6] Mark D Robinson, Davis J McCarthy, and Gordon K Smyth. edgeR: a Bioconductor package for differential expression analysis of digital gene expression data. *Bioinformatics*, 26(1):139–140, January 2010.
- [7] A. Wysoker, K. Tibbetts, and T. Fennell. Picard. <http://picard.sourceforge.net>, 2012.