

Differential analysis of count data – the DESeq2 package

Michael I. Love¹, Simon Anders², and Wolfgang Huber³

¹Department of Biostatistics, Dana-Farber Cancer Institute and Harvard TH Chan School of Public Health, Boston, US;

²Institute for Molecular Medicine Finland (FIMM), Helsinki, Finland;

³European Molecular Biology Laboratory (EMBL), Heidelberg, Germany

October 17, 2016

Abstract

A basic task in the analysis of count data from RNA-seq is the detection of differentially expressed genes. The count data are presented as a table which reports, for each sample, the number of sequence fragments that have been assigned to each gene. Analogous data also arise for other assay types, including comparative ChIP-Seq, HiC, shRNA screening, mass spectrometry. An important analysis question is the quantification and statistical inference of systematic changes between conditions, as compared to within-condition variability. The package *DESeq2* provides methods to test for differential expression by use of negative binomial generalized linear models; the estimates of dispersion and logarithmic fold changes incorporate data-driven prior distributions¹. This vignette explains the use of the package and demonstrates typical workflows. An RNA-seq workflow² on the Bioconductor website covers similar material to this vignette but at a slower pace, including the generation of count matrices from FASTQ files.

¹Other *Bioconductor* packages with similar aims are [edgeR](#), [limma](#), [DSS](#), [EBSeq](#) and [bay-Seq](#).

²<http://www.bioconductor.org/help/workflows/rnaseqGene/>

Package

DESeq2 1.14.0

If you use *DESeq2* in published research, please cite:

M. I. Love, W. Huber, S. Anders: **Moderated estimation of fold change and dispersion for RNA-seq data with DESeq2.** *Genome Biology* 2014, **15**:550.
<http://dx.doi.org/10.1186/s13059-014-0550-8>

Contents

| | | |
|-------|--|----|
| 1 | Standard workflow | 5 |
| 1.1 | Quick start | 5 |
| 1.2 | How to get help. | 5 |
| 1.3 | Input data | 5 |
| 1.3.1 | Why un-normalized counts? | 5 |
| 1.3.2 | <i>SummarizedExperiment</i> input | 6 |
| 1.3.3 | Count matrix input | 7 |
| 1.3.4 | tximport: transcript abundance summarized to gene-level | 9 |
| 1.3.5 | <i>HTSeq</i> input. | 10 |
| 1.3.6 | Pre-filtering | 11 |
| 1.3.7 | Note on factor levels | 11 |
| 1.3.8 | Collapsing technical replicates | 12 |
| 1.3.9 | About the pasilla dataset | 12 |
| 1.4 | Differential expression analysis | 12 |
| 1.5 | Exploring and exporting results | 15 |
| 1.5.1 | MA-plot | 15 |
| 1.5.2 | Plot counts | 17 |
| 1.5.3 | More information on results columns | 18 |
| 1.5.4 | Rich visualization and reporting of results | 18 |
| 1.5.5 | Exporting results to CSV files | 19 |
| 1.6 | Multi-factor designs | 20 |
| 2 | Data transformations and visualization | 23 |
| 2.1 | Count data transformations | 23 |
| 2.1.1 | Blind dispersion estimation | 23 |
| 2.1.2 | Extracting transformed values | 24 |
| 2.1.3 | Regularized log transformation. | 24 |
| 2.1.4 | Variance stabilizing transformation | 25 |
| 2.1.5 | Effects of transformations on the variance | 26 |
| 2.2 | Data quality assessment by sample clustering and visualization | 26 |
| 2.2.1 | Heatmap of the count matrix. | 27 |
| 2.2.2 | Heatmap of the sample-to-sample distances. | 27 |
| 2.2.3 | Principal component plot of the samples | 28 |
| 3 | Variations to the standard workflow | 30 |
| 3.1 | Wald test individual steps | 30 |
| 3.2 | Contrasts. | 30 |
| 3.3 | Interactions. | 30 |
| 3.4 | Time-series experiments | 32 |
| 3.5 | Likelihood ratio test | 32 |

Differential analysis of count data – the DESeq2 package

| | | |
|--------|---|----|
| 3.6 | Approach to count outliers | 33 |
| 3.7 | Dispersion plot and fitting alternatives | 34 |
| 3.7.1 | Local or mean dispersion fit | 36 |
| 3.7.2 | Supply a custom dispersion fit | 36 |
| 3.8 | Independent filtering of results | 36 |
| 3.9 | Tests of log ₂ fold change above or below a threshold | 38 |
| 3.10 | Access to all calculated values | 39 |
| 3.11 | Sample-/gene-dependent normalization factors | 42 |
| 3.12 | “Model matrix not full rank” | 43 |
| 3.12.1 | Linear combinations | 43 |
| 3.12.2 | Levels without samples | 46 |
| 4 | Theory behind DESeq2 | 49 |
| 4.1 | The DESeq2 model. | 49 |
| 4.2 | Changes compared to the <i>DESeq</i> package | 49 |
| 4.3 | Methods changes since the 2014 DESeq2 paper | 50 |
| 4.4 | Count outlier detection | 51 |
| 4.5 | Contrasts. | 52 |
| 4.6 | Expanded model matrices | 52 |
| 4.7 | Independent filtering and multiple testing | 52 |
| 4.7.1 | Filtering criteria | 52 |
| 4.7.2 | Why does it work? | 54 |
| 5 | Frequently asked questions | 54 |
| 5.1 | How can I get support for DESeq2? | 54 |
| 5.2 | Why are some <i>p</i> values set to NA? | 55 |
| 5.3 | How can I get unfiltered DESeq results? | 55 |
| 5.4 | How do I use the variance stabilized or rlog transformed data for differential testing? | 55 |
| 5.5 | Can I use DESeq2 to analyze paired samples? | 55 |
| 5.6 | If I have multiple groups, should I run all together or split into pairs of groups? | 56 |
| 5.7 | Can I run DESeq2 to contrast the levels of 100 groups? | 57 |
| 5.8 | Can I use DESeq2 to analyze a dataset without replicates? | 57 |
| 5.9 | How can I include a continuous covariate in the design formula? | 57 |
| 5.10 | Will the log fold change shrinkage “overshrink” large differences? | 57 |
| 5.11 | I ran a likelihood ratio test, but <code>results()</code> only gives me one comparison. | 59 |

Differential analysis of count data – the DESeq2 package

| | | |
|------|--|----|
| 5.12 | What are the exact steps performed by <code>DESeq()</code> ? | 59 |
| 5.13 | Is there an official Galaxy tool for DESeq2? | 59 |
| 5.14 | I want to benchmark DESeq2 comparing to other DE tools. 59 | |
| 6 | Acknowledgments | 60 |
| 7 | Session Info | 60 |

1 Standard workflow

1.1 Quick start

Here we show the most basic steps for a differential expression analysis. These steps require you have a *RangedSummarizedExperiment* object `se` which contains the counts and information about samples. The `design` indicates that we want to measure the effect of condition, controlling for batch differences. The two factor variables `batch` and `condition` should be columns of `colData(se)`.

```
dds <- DESeqDataSet(se, design = ~ batch + condition)
dds <- DESeq(dds)
res <- results(dds, contrast=c("condition","trt","con"))
```

If you have a count matrix and sample information table, the first line would use `DESeqDataSetFromMatrix` instead of `DESeqDataSet`, as shown in Section 1.3.3.

1.2 How to get help

All *DESeq2* questions should be posted to the Bioconductor support site: <https://support.bioconductor.org>, which serves as a repository of questions and answers. See the first question in the list of Frequently Asked Questions (Section 5) for more information about how to construct an informative post.

1.3 Input data

1.3.1 Why un-normalized counts?

As input, the *DESeq2* package expects count data as obtained, e. g., from RNA-seq or another high-throughput sequencing experiment, in the form of a matrix of integer values. The value in the i -th row and the j -th column of the matrix tells how many reads can be assigned to gene i in sample j . Analogously, for other types of assays, the rows of the matrix might correspond e. g. to binding regions (with ChIP-Seq) or peptide sequences (with quantitative mass spectrometry). We will list method for obtaining count matrices in sections below.

The values in the matrix should be un-normalized counts of sequencing reads (for single-end RNA-seq) or fragments (for paired-end RNA-seq). The [RNA-seq workflow](#) describes multiple techniques for preparing such count matrices. It is important to provide count matrices as input for *DESeq2*'s statistical model [1] to hold, as only the count values allow assessing the measurement precision correctly. The *DESeq2* model internally corrects for library size, so transformed or normalized values such as counts scaled by library size should not be used as input.

Differential analysis of count data – the DESeq2 package

1.3.2 SummarizedExperiment input

The class used by the *DESeq2* package to store the read counts is *DESeqDataSet* which extends the *RangedSummarizedExperiment* class of the *SummarizedExperiment* package. This facilitates preparation steps and also downstream exploration of results. For counting aligned reads in genes, the `summarizeOverlaps` function of *GenomicAlignments* with `mode="Union"` is encouraged, resulting in a *RangedSummarizedExperiment* object. Other methods for obtaining count matrices are described in the next section.

An example of the steps to produce a *RangedSummarizedExperiment* can be found in an RNA-seq workflow on the Bioconductor website: <http://www.bioconductor.org/help/workflows/rnaseqGene/> and in the vignette for the data package *airway*. Here we load the *RangedSummarizedExperiment* from that package in order to build a *DESeqDataSet*.

```
library("airway")
data("airway")
se <- airway
```

A *DESeqDataSet* object must have an associated design formula. The design formula expresses the variables which will be used in modeling. The formula should be a tilde (~) followed by the variables with plus signs between them (it will be coerced into an *formula* if it is not already). An intercept is included, representing the base mean of counts. The design can be changed later, however then all differential analysis steps should be repeated, as the design formula is used to estimate the dispersions and to estimate the log₂ fold changes of the model. The constructor function below shows the generation of a *DESeqDataSet* from a *RangedSummarizedExperiment* *se*.

Note: In order to benefit from the default settings of the package, you should put the variable of interest at the end of the formula and make sure the control level is the first level.

```
library("DESeq2")
ddsSE <- DESeqDataSet(se, design = ~ cell + dex)
ddsSE

## class: DESeqDataSet
## dim: 64102 8
## metadata(2): '' version
## assays(1): counts
## rownames(64102): ENSG000000000003 ENSG000000000005 ... LRG_98
##   LRG_99
## rowData names(0):
## colnames(8): SRR1039508 SRR1039509 ... SRR1039520
##   SRR1039521
## colData names(9): SampleName cell ... Sample BioSample
```

Differential analysis of count data – the DESeq2 package

1.3.3 Count matrix input

Alternatively, the function `DESeqDataSetFromMatrix` can be used if you already have a matrix of read counts prepared from another source. Another method for quickly producing count matrices from alignment files is the `featureCounts` function in the `Rsubread` package. To use `DESeqDataSetFromMatrix`, the user should provide the counts matrix, the information about the samples (the columns of the count matrix) as a `DataFrame` or `data.frame`, and the design formula.

To demonstrate the use of `DESeqDataSetFromMatrix`, we will read in count data from the `pasilla` package. We read in a count matrix, which we will name `countData`, and the sample information table, which we will name `colData`. Further below we describe how to extract these objects from, e.g. `featureCounts` output.

```
library("pasilla")
pasCts <- system.file("extdata", "pasilla_gene_counts.tsv",
                     package="pasilla", mustWork=TRUE)
pasAnno <- system.file("extdata", "pasilla_sample_annotation.csv",
                      package="pasilla", mustWork=TRUE)
countData <- as.matrix(read.csv(pasCts, sep="\t", row.names="gene_id"))
colData <- read.csv(pasAnno, row.names=1)
colData <- colData[,c("condition", "type")]
```

We examine the count matrix and column data to see if they are consistent:

```
head(countData)

##           untreated1 untreated2 untreated3 untreated4 treated1
## FBgn0000003           0           0           0           0           0
## FBgn0000008           92          161           76           70          140
## FBgn0000014            5            1            0            0            4
## FBgn0000015            0            2            1            2            1
## FBgn0000017          4664          8714          3564          3150          6205
## FBgn0000018           583           761           245           310           722
##
##           treated2 treated3
## FBgn0000003           0            1
## FBgn0000008           88            70
## FBgn0000014            0            0
## FBgn0000015            0            0
## FBgn0000017          3072          3334
## FBgn0000018           299            308

head(colData)

##           condition           type
## treated1fb      treated single-read
## treated2fb      treated paired-end
## treated3fb      treated paired-end
## untreated1fb   untreated single-read
## untreated2fb   untreated single-read
```

Differential analysis of count data – the DESeq2 package

```
## untreated3fb untreated paired-end
```

Note that these are not in the same order with respect to samples! It is critical that the columns of the count matrix and the rows of the column data (information about samples) are in the same order.. We should re-arrange one or the other so that they are consistent in terms of sample order (if we do not, later functions would produce an error). We additionally need to chop off the "fb" of the row names of `colData`, so the naming is consistent.

```
rownames(colData) <- sub("fb", "", rownames(colData))
all(rownames(colData) %in% colnames(countData))

## [1] TRUE

countData <- countData[, rownames(colData)]
all(rownames(countData) == colnames(countData))

## [1] TRUE
```

If you have used the `featureCounts` function in the *Rsubread* package, the matrix of read counts can be directly provided from the "counts" element in the list output. The count matrix and column data can typically be read into R from flat files using base R functions such as `read.csv` or `read.delim`. For *HTSeq* count files, see the dedicated input function below.

With the count matrix, `countData`, and the sample information, `colData`, we can construct a *DESeqDataSet*:

```
dds <- DESeqDataSetFromMatrix(countData = countData,
                              colData = colData,
                              design = ~ condition)

dds

## class: DESeqDataSet
## dim: 14599 7
## metadata(1): version
## assays(1): counts
## rownames(14599): FBgn0000003 FBgn0000008 ... FBgn0261574
##   FBgn0261575
## rowData names(0):
## colnames(7): treated1 treated2 ... untreated3 untreated4
## colData names(2): condition type
```

If you have additional feature data, it can be added to the *DESeqDataSet* by adding to the metadata columns of a newly constructed object. (Here we add redundant data just for demonstration, as the gene names are already the rownames of the `dds`.)

Differential analysis of count data – the DESeq2 package

```
featureData <- data.frame(gene=rownames(countData))
(mcols(dds) <- DataFrame(mcols(dds), featureData))

## DataFrame with 14599 rows and 1 column
##           gene
##      <factor>
## 1    FBgn0000003
## 2    FBgn0000008
## 3    FBgn0000014
## 4    FBgn0000015
## 5    FBgn0000017
## ...          ...
## 14595 FBgn0261571
## 14596 FBgn0261572
## 14597 FBgn0261573
## 14598 FBgn0261574
## 14599 FBgn0261575
```

1.3.4 tximport: transcript abundance summarized to gene-level

Users can create gene-level count matrices for use with *DESeq2* by importing information using the *tximport* package. This workflow allows users to import transcript abundance estimates from a variety of external software, including the following methods:

- [Sailfish](#) [2]
- [Salmon](#) [3]
- [kallisto](#) [4]
- [RSEM](#) [5]

Some advantages of using the above methods for transcript abundance estimation are: (i) this approach corrects for potential changes in gene length across samples (e.g. from differential isoform usage) [6], (ii) some of these methods (*Sailfish*, *Salmon*, *kallisto*) are substantially faster and require less memory and disk usage compared to alignment-based methods that require creation and storage of BAM files, and (iii) it is possible to avoid discarding those fragments that can align to multiple genes with homologous sequence, thus increasing sensitivity [7].

Full details on the motivation and methods for importing transcript level abundance and count estimates, summarizing to gene-level count matrices and producing an offset which corrects for potential changes in average transcript length across samples are described in [8]. The *tximport*→*DESeq2* approach uses rounded estimated gene counts (but not normalized) instead of the raw count of fragments which can be unambiguously assigned to a gene.

Here, we demonstrate how to import transcript abundances and construct of a gene-level *DESeqDataSet* object from *Sailfish* `quant.sf` files, which are stored in the *tximportData* package. Note that, instead of locating `dir` using `system.file`,

Differential analysis of count data – the DESeq2 package

a user would typically just provide a path, e.g. `/path/to/quant/files`. For further details on use of `tximport`, including the construction of the `tx2gene` table for linking transcripts to genes, please refer to the `tximport` package vignette.

```
library("tximport")
library("readr")
library("tximportData")
dir <- system.file("extdata", package="tximportData")
samples <- read.table(file.path(dir,"samples.txt"), header=TRUE)
files <- file.path(dir,"salmon", samples$run, "quant.sf")
names(files) <- paste0("sample",1:6)
tx2gene <- read.csv(file.path(dir, "tx2gene.csv"))
txi <- tximport(files, type="salmon", tx2gene=tx2gene, reader=read_tsv)
```

Next we create an condition vector to demonstrate building an `DESeqDataSet`. For a typical use, this information would already be present as a column of the `samples` table. The best practice is to read `colData` from a CSV or TSV file, and to construct `files` from a column of `colData`, as shown in the code chunk above.

```
coldata <- data.frame(condition=factor(rep(c("A","B"),each=3)))
rownames(coldata) <- colnames(txi$counts)
ddsTxi <- DESeqDataSetFromTximport(txi, colData=coldata,
                                  design=~ condition)
```

The `ddsTxi` object can then be used as `dds` in the following analysis steps.

1.3.5 HTSeq input

You can use the function `DESeqDataSetFromHTSeqCount` if you have `htseq-count` from the `HTSeq` python package³. For an example of using the python scripts, see the `pasilla` data package. First you will want to specify a variable which points to the directory in which the `HTSeq` output files are located.

³available from <http://www-huber.embl.de/users/anders/HTSeq>, described in [9]

```
directory <- "/path/to/your/files/"
```

However, for demonstration purposes only, the following line of code points to the directory for the demo `HTSeq` output files packages for the `pasilla` package.

```
directory <- system.file("extdata", package="pasilla", mustWork=TRUE)
```

We specify which files to read in using `list.files`, and select those files which contain the string `"treated"` using `grep`. The `sub` function is used to chop up the sample filename to obtain the condition status, or you might alternatively read in a phenotypic table using `read.table`.

```
sampleFiles <- grep("treated",list.files(directory),value=TRUE)
sampleCondition <- sub(".*treated).*","\\1",sampleFiles)
```

Differential analysis of count data – the DESeq2 package

```
sampleTable <- data.frame(sampleName = sampleFiles,
                           fileName = sampleFiles,
                           condition = sampleCondition)
ddsHTSeq <- DESeqDataSetFromHTSeqCount(sampleTable = sampleTable,
                                       directory = directory,
                                       design= ~ condition)

ddsHTSeq

## class: DESeqDataSet
## dim: 70463 7
## metadata(1): version
## assays(1): counts
## rownames(70463): FBgn0000003:001 FBgn0000008:001 ...
##   FBgn0261575:001 FBgn0261575:002
## rowData names(0):
## colnames(7): treated1fb.txt treated2fb.txt ...
##   untreated3fb.txt untreated4fb.txt
## colData names(1): condition
```

1.3.6 Pre-filtering

While it is not necessary to pre-filter low count genes before running the *DESeq2* functions, there are two reasons which make pre-filtering useful: by removing rows in which there are no reads or nearly no reads, we reduce the memory size of the `dds` data object and we increase the speed of the transformation and testing functions within *DESeq2*. Here we perform a minimal pre-filtering to remove rows that have only 0 or 1 read. Note that more strict filtering to increase power is *automatically* applied via independent filtering on the mean of normalized counts within the `results` function, which will be discussed in Section 3.8.

```
dds <- dds[ rowSums(counts(dds)) > 1, ]
```

1.3.7 Note on factor levels

By default, R will choose a *reference level* for factors based on alphabetical order. Then, if you never tell the *DESeq2* functions which level you want to compare against (e.g. which level represents the control group), the comparisons will be based on the alphabetical order of the levels. There are two solutions: you can either explicitly tell `results` which comparison to make using the `contrast` argument (this will be shown later), or you can explicitly set the factors levels. Setting the factor levels can be done in two ways, either using `factor`:

```
dds$condition <- factor(dds$condition, levels=c("untreated", "treated"))
```

...or using `relevel`, just specifying the reference level:

Differential analysis of count data – the DESeq2 package

```
dds$condition <- relevel(dds$condition, ref="untreated")
```

If you need to subset the columns of a *DESeqDataSet*, i.e., when removing certain samples from the analysis, it is possible that all the samples for one or more levels of a variable in the design formula would be removed. In this case, the `droplevels` function can be used to remove those levels which do not have samples in the current *DESeqDataSet*:

```
dds$condition <- droplevels(dds$condition)
```

1.3.8 Collapsing technical replicates

DESeq2 provides a function `collapseReplicates` which can assist in combining the counts from technical replicates into single columns of the count matrix. The term “technical replicate” implies multiple sequencing runs of the same library. You should not collapse biological replicates using this function. See the manual page for an example of the use of `collapseReplicates`.

1.3.9 About the pasilla dataset

We continue with the *pasilla* data constructed from the count matrix method above. This data set is from an experiment on *Drosophila melanogaster* cell cultures and investigated the effect of RNAi knock-down of the splicing factor *pasilla* [10]. The detailed transcript of the production of the *pasilla* data is provided in the vignette of the data package *pasilla*.

1.4 Differential expression analysis

The standard differential expression analysis steps are wrapped into a single function, `DESeq`. The estimation steps performed by this function are described in Section 4.1, in the manual page for `?DESeq` and in the Methods section of the *DESeq2* publication [1]. The individual sub-functions which are called by `DESeq` are still available, described in Section 3.1.

Results tables are generated using the function `results`, which extracts a results table with \log_2 fold changes, p values and adjusted p values. With no arguments to `results`, the results will be for the last variable in the design formula, and if this is a factor, the comparison will be the last level of this variable over the first level. Details about the comparison are printed to the console. The text, `condition treated vs untreated`, tells you that the estimates are of the logarithmic fold change $\log_2(\text{treated}/\text{untreated})$.

```
dds <- DESeq(dds)
res <- results(dds)
```

Differential analysis of count data – the DESeq2 package

```
res

## log2 fold change (MAP): condition treated vs untreated
## Wald test p-value: condition treated vs untreated
## DataFrame with 11638 rows and 6 columns
##           baseMean log2FoldChange   lfcSE      stat    pvalue
##           <numeric>      <numeric> <numeric> <numeric> <numeric>
## FBgn0000008      95.144      0.00175   0.1825   0.0096   0.9923
## FBgn0000014       1.057     -0.02335   0.0983  -0.2375   0.8123
## FBgn0000015       0.847     -0.09134   0.1022  -0.8940   0.3713
## FBgn0000017    4352.593     -0.22390   0.1176  -1.9043   0.0569
## FBgn0000018     418.615     -0.09532   0.1348  -0.7069   0.4796
## ...           ...           ...           ...           ...
## FBgn0261570     3208.4      0.2753    0.118   2.326    0.020
## FBgn0261572       6.2     -0.2556    0.211  -1.213    0.225
## FBgn0261573     2241.0      0.0119    0.107   0.112    0.911
## FBgn0261574     4857.7      0.0131    0.165   0.079    0.937
## FBgn0261575      10.7      0.0338    0.188   0.180    0.857
##           padj
##           <numeric>
## FBgn0000008      0.997
## FBgn0000014      NA
## FBgn0000015      NA
## FBgn0000017      0.286
## FBgn0000018      0.828
## ...           ...
## FBgn0261570      0.142
## FBgn0261572      0.622
## FBgn0261573      0.983
## FBgn0261574      0.989
## FBgn0261575      0.968
```

These steps should take less than 30 seconds for most analyses. For experiments with many samples (e.g. 100 samples), one can take advantage of parallelized computation. Both of the above functions have an argument `parallel` which if set to `TRUE` can be used to distribute computation across cores specified by the `register` function of *BiocParallel*. For example, the following chunk (not evaluated here), would register 4 cores, and then the two functions above, with `parallel=TRUE`, would split computation over these cores.

```
library("BiocParallel")
register(MulticoreParam(4))
```

We can order our results table by the smallest adjusted *p* value:

```
resOrdered <- res[order(res$padj),]
```

We can summarize some basic tallies using the `summary` function.

Differential analysis of count data – the DESeq2 package

```
summary(res)

##
## out of 11638 with nonzero total read count
## adjusted p-value < 0.1
## LFC > 0 (up)      : 516, 4.4%
## LFC < 0 (down)   : 538, 4.6%
## outliers [1]     : 14, 0.12%
## low counts [2]   : 3152, 27%
## (mean count < 6)
## [1] see 'cooksCutoff' argument of ?results
## [2] see 'independentFiltering' argument of ?results
```

How many adjusted p -values were less than 0.1?

```
sum(res$padj < 0.1, na.rm=TRUE)

## [1] 1054
```

The `results` function contains a number of arguments to customize the results table which is generated. Note that the `results` function automatically performs independent filtering based on the mean of normalized counts for each gene, optimizing the number of genes which will have an adjusted p value below a given FDR cutoff, `alpha`. Independent filtering is further discussed in Section 3.8. By default the argument `alpha` is set to 0.1. If the adjusted p value cutoff will be a value other than 0.1, `alpha` should be set to that value:

```
res05 <- results(dds, alpha=0.05)
summary(res05)

##
## out of 11638 with nonzero total read count
## adjusted p-value < 0.05
## LFC > 0 (up)      : 408, 3.5%
## LFC < 0 (down)   : 431, 3.7%
## outliers [1]     : 14, 0.12%
## low counts [2]   : 3152, 27%
## (mean count < 6)
## [1] see 'cooksCutoff' argument of ?results
## [2] see 'independentFiltering' argument of ?results

sum(res05$padj < 0.05, na.rm=TRUE)

## [1] 839
```

A generalization of the idea of p value filtering is to *weight* hypotheses to optimize power. A new Bioconductor package, *IHW*, is now available that implements the method of *Independent Hypothesis Weighting* [11]. Here we show the use of *IHW* for p value adjustment of *DESeq2* results. For more details, please see the vignette of the *IHW* package. Note that the *IHW* result object is stored in the metadata.

Differential analysis of count data – the DESeq2 package

```
library("IHW")
resIHW <- results(dds, filterFun=ihw)
summary(resIHW)

##
## out of 11638 with nonzero total read count
## adjusted p-value < 0.1
## LFC > 0 (up)      : 510, 4.4%
## LFC < 0 (down)    : 548, 4.7%
## outliers [1]      : 14, 0.12%
## [1] see 'cooksCutoff' argument of ?results
## [2] see metadata(res)$ihwResult on hypothesis weighting

sum(resIHW$padj < 0.1, na.rm=TRUE)

## [1] 1058

metadata(resIHW)$ihwResult

## ihwResult object with 11638 hypothesis tests
## Nominal FDR control level: 0.1
## Split into 7 bins, based on an ordinal covariate
```

If a multi-factor design is used, or if the variable in the design formula has more than two levels, the `contrast` argument of `results` can be used to extract different comparisons from the `DESeqDataSet` returned by `DESeq`. Multi-factor designs are discussed further in Section 1.6, and the use of the `contrast` argument is discussed in Section 3.2.

For advanced users, note that all the values calculated by the `DESeq2` package are stored in the `DESeqDataSet` object, and access to these values is discussed in Section 3.10.

1.5 Exploring and exporting results

1.5.1 MA-plot

In `DESeq2`, the function `plotMA` shows the log₂ fold changes attributable to a given variable over the mean of normalized counts. Points will be colored red if the adjusted *p* value is less than 0.1. Points which fall out of the window are plotted as open triangles pointing either up or down.

```
plotMA(res, main="DESeq2", ylim=c(-2,2))
```

After calling `plotMA`, one can use the function `identify` to interactively detect the row number of individual genes by clicking on the plot. One can then recover the gene identifiers by saving the resulting indices:

Differential analysis of count data – the DESeq2 package




Figure 1: MA-plot. These plots show the log₂ fold changes from the treatment over the mean of normalized counts, i.e. the average of counts normalized by size factors. The left plot shows the “unshrunk” log₂ fold changes, while the right plot, produced by the code above, shows the shrinkage of log₂ fold changes resulting from the incorporation of zero-centered normal prior. The shrinkage is greater for the log₂ fold change estimates from genes with low counts and high dispersion, as can be seen by the narrowing of spread of leftmost points in the right plot.

```
idx <- identify(res$baseMean, res$log2FoldChange)
rownames(res)[idx]
```

The MA-plot of log₂ fold changes returned by *DESeq2* allows us to see how the shrinkage of fold changes works for genes with low counts. You can still obtain results tables which include the “unshrunk” log₂ fold changes (for a simple comparison, the ratio of the mean normalized counts in the two groups). A column `lfcMLE` with the unshrunk maximum likelihood estimate (MLE) for the log₂ fold change will be added with an additional argument to `results`:

```
resMLE <- results(dds, addMLE=TRUE)
head(resMLE, 4)

## log2 fold change (MAP): condition treated vs untreated
## Wald test p-value: condition treated vs untreated
## DataFrame with 4 rows and 7 columns
##           baseMean log2FoldChange  lfcMLE  lfcSE  stat
##           <numeric>  <numeric> <numeric> <numeric> <numeric>
## FBgn0000008    95.144     0.00175  0.00215  0.1825  0.0096
## FBgn0000014     1.057    -0.02335 -0.49669  0.0983 -0.2375
## FBgn0000015     0.847    -0.09134 -1.88276  0.1022 -0.8940
## FBgn0000017  4352.593    -0.22390 -0.24003  0.1176 -1.9043
##           pvalue  padj
##           <numeric> <numeric>
## FBgn0000008  0.9923  0.997
## FBgn0000014  0.8123  NA
## FBgn0000015  0.3713  NA
## FBgn0000017  0.0569  0.286
```


Differential analysis of count data – the DESeq2 package

One can make an MA-plot of the unshrunk estimates like so:

```
plotMA(resMLE, MLE=TRUE, main="unshrunk LFC", ylim=c(-2,2))
```

1.5.2 Plot counts

It can also be useful to examine the counts of reads for a single gene across the groups. A simple function for making this plot is `plotCounts`, which normalizes counts by sequencing depth and adds a pseudocount of $\frac{1}{2}$ to allow for log scale plotting. The counts are grouped by the variables in `intgroup`, where more than one variable can be specified. Here we specify the gene which had the smallest p value from the results table created above. You can select the gene to plot by rowname or by numeric index.

```
plotCounts(dds, gene=which.min(res$padj), intgroup="condition")
```

For customized plotting, an argument `returnData` specifies that the function should only return a `data.frame` for plotting with `ggplot`.

```
d <- plotCounts(dds, gene=which.min(res$padj), intgroup="condition",  
               returnData=TRUE)  
library("ggplot2")  
ggplot(d, aes(x=condition, y=count)) +  
  geom_point(position=position_jitter(w=0.1,h=0)) +  
  scale_y_log10(breaks=c(25,100,400))
```




Figure 2: Plot of counts for one gene. The plot of normalized counts (plus a pseudocount of $\frac{1}{2}$) either made using the `plotCounts` function (left) or using another plotting library (right, using `ggplot2`).

Differential analysis of count data – the DESeq2 package

1.5.3 More information on results columns

Information about which variables and tests were used can be found by calling the function `mcols` on the results object.

```
mcols(res)$description
## [1] "mean of normalized counts for all samples"
## [2] "log2 fold change (MAP): condition treated vs untreated"
## [3] "standard error: condition treated vs untreated"
## [4] "Wald statistic: condition treated vs untreated"
## [5] "Wald test p-value: condition treated vs untreated"
## [6] "BH adjusted p-values"
```

For a particular gene, a log₂ fold change of -1 for `condition treated vs untreated` means that the treatment induces a multiplicative change in observed gene expression level of $2^{-1} = 0.5$ compared to the untreated condition. If the variable of interest is continuous-valued, then the reported log₂ fold change is per unit of change of that variable.

Note on p-values set to NA: some values in the results table can be set to NA for one of the following reasons:

1. If within a row, all samples have zero counts, the `baseMean` column will be zero, and the log₂ fold change estimates, *p* value and adjusted *p* value will all be set to NA.
2. If a row contains a sample with an extreme count outlier then the *p* value and adjusted *p* value will be set to NA. These outlier counts are detected by Cook's distance. Customization of this outlier filtering and description of functionality for replacement of outlier counts and refitting is described in Section 3.6,
3. If a row is filtered by automatic independent filtering, for having a low mean normalized count, then only the adjusted *p* value will be set to NA. Description and customization of independent filtering is described in Section 3.8.

1.5.4 Rich visualization and reporting of results

ReportingTools. An HTML report of the results with plots and sortable/filterable columns can be generated using the `ReportingTools` package on a `DESeqDataSet` that has been processed by the `DESeq` function. For a code example, see the “RNA-seq differential expression” vignette at the `ReportingTools` page, or the manual page for the `publish` method for the `DESeqDataSet` class.

regionReport. An HTML and PDF summary of the results with plots can also be generated using the `regionReport` package. The `DESeq2Report` function should be run on a `DESeqDataSet` that has been processed by the `DESeq` function. For more details see the manual page for `DESeq2Report` and an example vignette in the `regionReport` package.

Differential analysis of count data – the DESeq2 package

Glimma. Interactive visualization of *DESeq2* output, including MA-plots (also called MD-plot) can be generated using the *Glimma* package. See the manual page for `glMDPlot.DESeqResults`.

pcaExplorer. Interactive visualization of *DESeq2* output, including PCA plots, boxplots of counts and other useful summaries can be generated using the *pca-Explorer* package. See the “Launching the application” section of the package vignette.

1.5.5 Exporting results to CSV files

A plain-text file of the results can be exported using the base *R* functions `write.csv` or `write.delim`. We suggest using a descriptive file name indicating the variable and levels which were tested.

```
write.csv(as.data.frame(resOrdered),
          file="condition_treated_results.csv")
```

Exporting only the results which pass an adjusted *p* value threshold can be accomplished with the `subset` function, followed by the `write.csv` function.

```
resSig <- subset(resOrdered, padj < 0.1)
resSig

## log2 fold change (MAP): condition treated vs untreated
## Wald test p-value: condition treated vs untreated
## DataFrame with 1054 rows and 6 columns
##           baseMean log2FoldChange      lfcSE      stat      pvalue
##           <numeric>      <numeric> <numeric> <numeric> <numeric>
## FBgn0039155      731          -4.10      0.143     -28.7 2.90e-181
## FBgn0025111     1501           2.70      0.118      22.9 5.16e-116
## FBgn0029167     3706          -2.11      0.093     -22.7 1.04e-113
## FBgn0003360     4343          -2.90      0.131     -22.2 6.88e-109
## FBgn0039827      262          -3.37      0.176     -19.1 1.20e-81
## ...           ...           ...           ...           ...           ...
## FBgn0032900      29.5         -0.561     0.224     -2.51 0.0122
## FBgn0030026     212.2          0.460     0.184      2.50 0.0123
## FBgn0031183     428.5         -0.319     0.128     -2.50 0.0124
## FBgn0038874     103.8         -0.509     0.204     -2.50 0.0124
## FBgn0053329     602.6         -0.423     0.169     -2.50 0.0124
##           padj
##           <numeric>
## FBgn0039155 2.46e-177
## FBgn0025111 2.19e-112
## FBgn0029167 2.95e-110
## FBgn0003360 1.46e-105
## FBgn0039827 2.04e-78
## ...           ...
## FBgn0032900 0.0988
```

Differential analysis of count data – the DESeq2 package

```
## FBgn0030026    0.0991
## FBgn0031183    0.0996
## FBgn0038874    0.0996
## FBgn0053329    0.0996
```

1.6 Multi-factor designs

Experiments with more than one factor influencing the counts can be analyzed using design formula that include the additional variables. By adding these to the design, one can control for additional variation in the counts. For example, if the condition samples are balanced across experimental batches, by including the `batch` factor to the design, one can increase the sensitivity for finding differences due to `condition`. There are multiple ways to analyze experiments when the additional variables are of interest and not just controlling factors (see Section 3.3 on interactions).

The data in the *pasilla* package have a condition of interest (the column `condition`), as well as information on the type of sequencing which was performed (the column `type`), as we can see below:

```
colData(dds)
## DataFrame with 7 rows and 3 columns
##           condition      type sizeFactor
##           <factor>    <factor> <numeric>
## treated1    treated single-read    1.636
## treated2    treated paired-end    0.761
## treated3    treated paired-end    0.833
## untreated1  untreated single-read    1.138
## untreated2  untreated single-read    1.793
## untreated3  untreated paired-end    0.650
## untreated4  untreated paired-end    0.752
```

We create a copy of the *DESeqDataSet*, so that we can rerun the analysis using a multi-factor design.

```
ddsMF <- dds
```

We can account for the different types of sequencing, and get a clearer picture of the differences attributable to the treatment. As `condition` is the variable of interest, we put it at the end of the formula. Thus the `results` function will by default pull the `condition` results unless `contrast` or `name` arguments are specified. Then we can re-run `DESeq`:

```
design(ddsMF) <- formula(~ type + condition)
ddsMF <- DESeq(ddsMF)
```

Again, we access the results using the `results` function.

Differential analysis of count data – the DESeq2 package

```
resMF <- results(ddsMF)
head(resMF)

## log2 fold change (MAP): condition treated vs untreated
## Wald test p-value: condition treated vs untreated
## DataFrame with 6 rows and 6 columns
##           baseMean log2FoldChange      lfcSE      stat      pvalue
##           <numeric>      <numeric> <numeric> <numeric> <numeric>
## FBgn0000008      95.144      -0.0238    0.1790   -0.133    0.8940
## FBgn0000014       1.057      -0.0239    0.0979   -0.244    0.8071
## FBgn0000015       0.847      -0.0791    0.0928   -0.853    0.3937
## FBgn0000017    4352.593      -0.2405    0.1053   -2.283    0.0224
## FBgn0000018    418.615      -0.0639    0.1213   -0.527    0.5984
## FBgn0000024       6.406       0.0634    0.2026    0.313    0.7542
##           padj
##           <numeric>
## FBgn0000008      0.963
## FBgn0000014      NA
## FBgn0000015      NA
## FBgn0000017      0.128
## FBgn0000018      0.849
## FBgn0000024      0.914
```

It is also possible to retrieve the log₂ fold changes, *p* values and adjusted *p* values of the `type` variable. The `contrast` argument of the function `results` takes a character vector of length three: the name of the variable, the name of the factor level for the numerator of the log₂ ratio, and the name of the factor level for the denominator. The `contrast` argument can also take other forms, as described in the help page for `results` and in Section 3.2.

```
resMFtype <- results(ddsMF,
                    contrast=c("type", "single-read", "paired-end"))
head(resMFtype)

## log2 fold change (MAP): type single-read vs paired-end
## Wald test p-value: type single-read vs paired-end
## DataFrame with 6 rows and 6 columns
##           baseMean log2FoldChange      lfcSE      stat      pvalue
##           <numeric>      <numeric> <numeric> <numeric> <numeric>
## FBgn0000008      95.144      -0.1926    0.1621   -1.188    0.2348
## FBgn0000014       1.057       0.1016    0.0654    1.552    0.1207
## FBgn0000015       0.847      -0.0219    0.0619   -0.353    0.7241
## FBgn0000017    4352.593      -0.0893    0.1022   -0.873    0.3824
## FBgn0000018    418.615       0.2046    0.1159    1.766    0.0775
## FBgn0000024       6.406       0.0465    0.1463    0.318    0.7506
##           padj
##           <numeric>
## FBgn0000008      0.531
```

Differential analysis of count data – the DESeq2 package

```
## FBgn0000014      NA
## FBgn0000015      NA
## FBgn0000017      0.679
## FBgn0000018      0.281
## FBgn0000024      NA
```

If the variable is continuous or an interaction term (see Section 3.3) then the results can be extracted using the `name` argument to `results`, where the name is one of elements returned by `resultsNames(dds)`.

2 Data transformations and visualization

2.1 Count data transformations

In order to test for differential expression, we operate on raw counts and use discrete distributions as described in the previous Section 1.4. However for other downstream analyses – e.g. for visualization or clustering – it might be useful to work with transformed versions of the count data.

Maybe the most obvious choice of transformation is the logarithm. Since count values for a gene can be zero in some conditions (and non-zero in others), some advocate the use of *pseudocounts*, i. e. transformations of the form

$$\mathbf{1} \quad y = \log_2(n + 1) \quad \text{or more generally,} \quad y = \log_2(n + n_0),$$

where n represents the count values and n_0 is a positive constant.

In this section, we discuss two alternative approaches that offer more theoretical justification and a rational way of choosing the parameter equivalent to n_0 above. The *regularized logarithm* or *rlog* incorporates a prior on the sample differences [1], and the other uses the concept of variance stabilizing transformations (VST) [12, 13, 14]. Both transformations produce transformed data on the \log_2 scale which has been normalized with respect to library size.

The point of these two transformations, the *rlog* and the VST, is to remove the dependence of the variance on the mean, particularly the high variance of the logarithm of count data when the mean is low. Both *rlog* and VST use the experiment-wide trend of variance over mean, in order to transform the data to remove the experiment-wide trend. Note that we do not require or desire that all the genes have *exactly* the same variance after transformation. Indeed, in Figure 4 below, you will see that after the transformations the genes with the same mean do not have exactly the same standard deviations, but that the experiment-wide trend has flattened. It is those genes with row variance above the trend which will allow us to cluster samples into interesting groups.

Note on running time: if you have many samples (e.g. 100s), the `rlog` function might take too long, and the `varianceStabilizingTransformation` is a faster choice. The *rlog* and VST have similar properties, but the *rlog* requires fitting a shrinkage term for each sample and each gene which takes time. See the *DESeq2* paper for more discussion on the differences [1]. In addition, a new function `vst` provides an even faster version of the `varianceStabilizingTransformation` but calculating the global dispersion trend on a subset of the genes (default 1000). `vst` may be attractive for interactive EDA.

2.1.1 Blind dispersion estimation

The two functions, `rlog` and `varianceStabilizingTransformation`, have an argument `blind`, for whether the transformation should be blind to the sample information specified by the design formula. When `blind` equals `TRUE` (the default), the functions will re-estimate the dispersions using only an intercept (design formula

Differential analysis of count data – the DESeq2 package

~ 1). This setting should be used in order to compare samples in a manner wholly unbiased by the information about experimental groups, for example to perform sample QA (quality assurance) as demonstrated below.

However, blind dispersion estimation is not the appropriate choice if one expects that many or the majority of genes (rows) will have large differences in counts which are explainable by the experimental design, and one wishes to transform the data for downstream analysis. In this case, using blind dispersion estimation will lead to large estimates of dispersion, as it attributes differences due to experimental design as unwanted “noise”, and will result in overly shrinking the transformed values towards each other. By setting `blind` to `FALSE`, the dispersions already estimated will be used to perform transformations, or if not present, they will be estimated using the current design formula. Note that only the fitted dispersion estimates from mean-dispersion trend line are used in the transformation (the global dependence of dispersion on mean for the entire experiment). So setting `blind` to `FALSE` is still for the most part not using the information about which samples were in which experimental group in applying the transformation.

2.1.2 Extracting transformed values

These functions return an object of class *DESeqTransform* which is a subclass of *RangedSummarizedExperiment*. For ~ 20 samples, running on a newly created *DESeqDataSet*, `rlog` may take 30 seconds, `varianceStabilizingTransformation` may take 5 seconds, and `vst` less than 1 second (by subsetting to 1000 genes for calculating the global dispersion trend). However, the running times are shorter and more similar with `blind=FALSE` and if the function `DESeq` has already been run, because then it is not necessary to re-estimate the dispersion values. The `assay` function is used to extract the matrix of normalized values.

```
rld <- rlog(dds, blind=FALSE)
vsd <- varianceStabilizingTransformation(dds, blind=FALSE)
vsd.fast <- vst(dds, blind=FALSE)
head(assay(rld), 3)

##           treated1 treated2 treated3 untreated1 untreated2
## FBgn0000008      6.505   6.670   6.500       6.478   6.531
## FBgn0000014      0.178   0.149   0.149       0.200   0.154
## FBgn0000015     -0.287  -0.294  -0.294      -0.295  -0.280
##           untreated3 untreated4
## FBgn0000008      6.674   6.552
## FBgn0000014      0.150   0.149
## FBgn0000015     -0.277  -0.264
```

2.1.3 Regularized log transformation

The function `rlog`, stands for *regularized log*, transforming the original count data to the \log_2 scale by fitting a model with a term for each sample and a prior distribution on the coefficients which is estimated from the data. This is the same

Differential analysis of count data – the DESeq2 package




Figure 3: VST and log2. Graphs of the variance stabilizing transformation for sample 1, in blue, and of the transformation $f(n) = \log_2(n/s_1)$, in black. n are the counts and s_1 is the size factor for the first sample.

kind of shrinkage (sometimes referred to as regularization, or moderation) of log fold changes used by the `DESeq` and `nbinomWaldTest`, as seen in Figure 1. The resulting data contains elements defined as:

$$\log_2(q_{ij}) = \beta_{i0} + \beta_{ij}$$

where q_{ij} is a parameter proportional to the expected true concentration of fragments for gene i and sample j (see Section 4.1), β_{i0} is an intercept which does not undergo shrinkage, and β_{ij} is the sample-specific effect which is shrunk toward zero based on the dispersion-mean trend over the entire dataset. The trend typically captures high dispersions for low counts, and therefore these genes exhibit higher shrinkage from the `rlog`.

Note that, as q_{ij} represents the part of the mean value μ_{ij} after the size factor s_j has been divided out, it is clear that the `rlog` transformation inherently accounts for differences in sequencing depth. Without priors, this design matrix would lead to a non-unique solution, however the addition of a prior on non-intercept betas allows for a unique solution to be found. The regularized log transformation is preferable to the variance stabilizing transformation if the size factors vary widely.

2.1.4 Variance stabilizing transformation

Above, we used a parametric fit for the dispersion. In this case, the closed-form expression for the variance stabilizing transformation is used by `varianceStabilizingTransformation`, which is derived in the file `vst.pdf`, that is distributed in the package alongside this vignette. If a local fit is used (option `fitType="locfit"` to `estimateDispersions`) a numerical integration is used instead.

The resulting variance stabilizing transformation is shown in Figure 3. The code that produces the figure is hidden from this vignette for the sake of brevity, but can be seen in the `.Rnw` or `.R` source file. Note that the vertical axis in such plots is the

Differential analysis of count data – the DESeq2 package




Figure 4: Per-gene standard deviation (taken across samples), against the rank of the mean. (a) for the shifted logarithm $\log_2(n + 1)$, the regularized log transformation (b) and the variance stabilizing transformation (c).

square root of the variance over all samples, so including the variance due to the experimental conditions. While a flat curve of the square root of variance over the mean may seem like the goal of such transformations, this may be unreasonable in the case of datasets with many true differences due to the experimental conditions.

2.1.5 Effects of transformations on the variance

Figure 4 plots the standard deviation of the transformed data, across samples, against the mean, using the shifted logarithm transformation [1](#), the regularized log transformation and the variance stabilizing transformation. The shifted logarithm has elevated standard deviation in the lower count range, and the regularized log to a lesser extent, while for the variance stabilized data the standard deviation is roughly constant along the whole dynamic range.

```
library("vsn")
notAllZero <- (rowSums(counts(dds))>0)
meanSdPlot(log2(counts(dds,normalized=TRUE)[notAllZero,] + 1))
meanSdPlot(assay(rld[notAllZero,]))
meanSdPlot(assay(vsd[notAllZero,]))
```

2.2 Data quality assessment by sample clustering and visualization

Data quality assessment and quality control (i. e. the removal of insufficiently good data) are essential steps of any data analysis. These steps should typically be performed very early in the analysis of a new data set, preceding or in parallel to the differential expression testing.

We define the term *quality* as *fitness for purpose*⁴. Our purpose is the detection of

⁴http://en.wikipedia.org/wiki/Quality_%28business%29

Differential analysis of count data – the DESeq2 package

differentially expressed genes, and we are looking in particular for samples whose experimental treatment suffered from an anomaly that renders the data points obtained from these particular samples detrimental to our purpose.

2.2.1 Heatmap of the count matrix

To explore a count matrix, it is often instructive to look at it as a heatmap. Below we show how to produce such a heatmap for various transformations of the data.

```
library("pheatmap")
select <- order(rowMeans(counts(dds,normalized=TRUE)),
               decreasing=TRUE)[1:20]

nt <- normTransform(dds) # defaults to log2(x+1)
log2.norm.counts <- assay(nt)[select,]
df <- as.data.frame(colData(dds)[,c("condition","type")])
pheatmap(log2.norm.counts, cluster_rows=FALSE, show_rownames=FALSE,
         cluster_cols=FALSE, annotation_col=df)
pheatmap(assay(rld)[select,], cluster_rows=FALSE, show_rownames=FALSE,
         cluster_cols=FALSE, annotation_col=df)
pheatmap(assay(vsd)[select,], cluster_rows=FALSE, show_rownames=FALSE,
         cluster_cols=FALSE, annotation_col=df)
```




Figure 5: Heatmaps showing the expression data of the 20 most highly expressed genes. The data is of log₂ normalized counts (left), from regularized log transformation (center) and from variance stabilizing transformation (right).

2.2.2 Heatmap of the sample-to-sample distances

Another use of the transformed data is sample clustering. Here, we apply the `dist` function to the transpose of the transformed count matrix to get sample-to-sample distances. We could alternatively use the variance stabilized transformation here.

Differential analysis of count data – the DESeq2 package

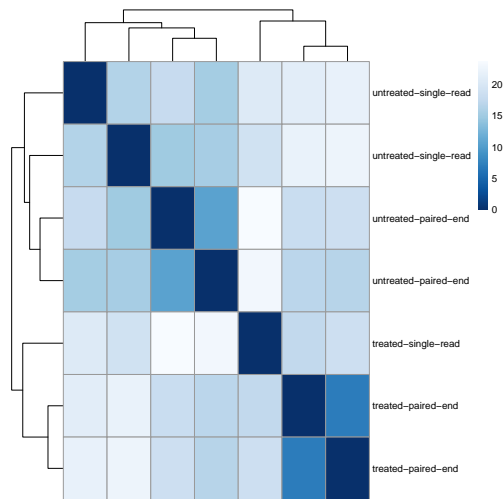


Figure 6: Sample-to-sample distances. Heatmap showing the Euclidean distances between the samples as calculated from the regularized log transformation.

```
sampleDists <- dist(t(assay(rld)))
```

A heatmap of this distance matrix gives us an overview over similarities and dissimilarities between samples (Figure 6): We have to provide a hierarchical clustering `hc` to the heatmap function based on the sample distances, or else the heatmap function would calculate a clustering based on the distances between the rows/columns of the distance matrix.

```
library("RColorBrewer")
sampleDistMatrix <- as.matrix(sampleDists)
rownames(sampleDistMatrix) <- paste(rld$condition, rld$type, sep="-")
colnames(sampleDistMatrix) <- NULL
colors <- colorRampPalette( rev(brewer.pal(9, "Blues")) )(255)
pheatmap(sampleDistMatrix,
          clustering_distance_rows=sampleDists,
          clustering_distance_cols=sampleDists,
          col=colors)
```

2.2.3 Principal component plot of the samples

Related to the distance matrix of Section 2.2.2 is the PCA plot of the samples, which we obtain as follows (Figure 7).

```
plotPCA(rld, intgroup=c("condition", "type"))
```

It is also possible to customize the PCA plot using the `ggplot` function.

Differential analysis of count data – the DESeq2 package




Figure 7: PCA plot. PCA plot. The 7 samples shown in the 2D plane spanned by their first two principal components. This type of plot is useful for visualizing the overall effect of experimental covariates and batch effects.

```
data <- plotPCA(rld, intgroup=c("condition", "type"), returnData=TRUE)
percentVar <- round(100 * attr(data, "percentVar"))
ggplot(data, aes(PC1, PC2, color=condition, shape=type)) +
  geom_point(size=3) +
  xlab(paste0("PC1: ", percentVar[1], "% variance")) +
  ylab(paste0("PC2: ", percentVar[2], "% variance")) +
  coord_fixed()
```




Figure 8: PCA plot. PCA plot customized using the *ggplot2* library.

3 Variations to the standard workflow

3.1 Wald test individual steps

The function `DESeq` runs the following functions in order:

```
dds <- estimateSizeFactors(dds)
dds <- estimateDispersions(dds)
dds <- nbinomWaldTest(dds)
```

3.2 Contrasts

A contrast is a linear combination of estimated log₂ fold changes, which can be used to test if differences between groups are equal to zero. The simplest use case for contrasts is an experimental design containing a factor with three levels, say A, B and C. Contrasts enable the user to generate results for all 3 possible differences: log₂ fold change of B vs A, of C vs A, and of C vs B. The `contrast` argument of `results` function is used to extract test results of log₂ fold changes of interest, for example:

```
results(dds, contrast=c("condition", "C", "B"))
```

Log₂ fold changes can also be added and subtracted by providing a `list` to the `contrast` argument which has two elements: the names of the log₂ fold changes to add, and the names of the log₂ fold changes to subtract. The names used in the list should come from `resultsNames(dds)`.

Alternatively, a numeric vector of the length of `resultsNames(dds)` can be provided, for manually specifying the linear combination of terms. Demonstrations of the use of contrasts for various designs can be found in the examples section of the help page for the `results` function. The mathematical formula that is used to generate the contrasts can be found in Section 4.5.

3.3 Interactions

Interaction terms can be added to the design formula, in order to test, for example, if the log₂ fold change attributable to a given condition is *different* based on another factor, for example if the condition effect differs across genotype.

Many users begin to add interaction terms to the design formula, when in fact a much simpler approach would give all the results tables that are desired. We will explain this approach first, because it is much simpler to perform. If the comparisons of interest are, for example, the effect of a condition for different sets of samples, a simpler approach than adding interaction terms explicitly to the design formula is to perform the following steps:

Differential analysis of count data – the DESeq2 package

1. combine the factors of interest into a single factor with all combinations of the original factors
2. change the design to include just this factor, e.g. `~ group`

Using this design is similar to adding an interaction term, in that it models multiple condition effects which can be easily extracted with `results`. Suppose we have two factors `genotype` (with values I, II, and III) and `condition` (with values A and B), and we want to extract the condition effect specifically for each genotype. We could use the following approach to obtain, e.g. the condition effect for genotype I:

```
dds$group <- factor(paste0(dds$genotype, dds$condition))
design(dds) <- ~ group
dds <- DESeq(dds)
resultsNames(dds)
results(dds, contrast=c("group", "IB", "IA"))
```

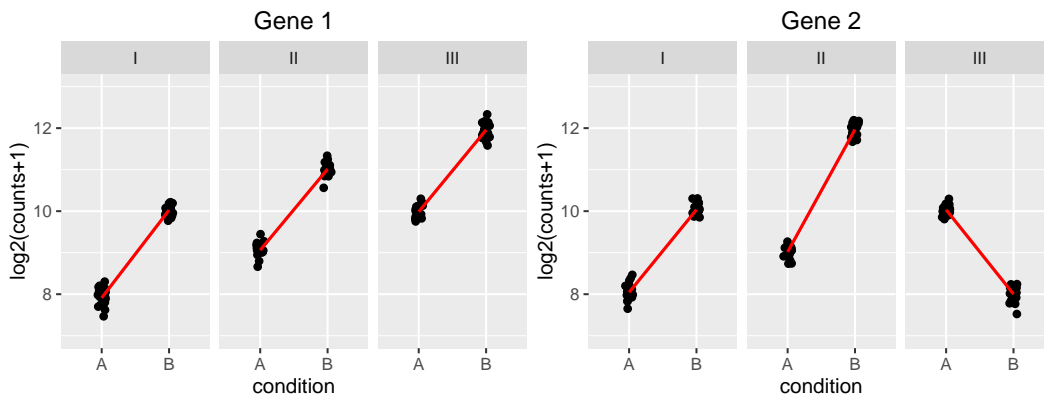


Figure 9: Genotype-specific condition effects. Here, the y-axis represents $\log_2(\text{counts} + 1)$, and each group has 20 samples (black dots). A red line connects the mean of the groups within each genotype. On the left side (Gene 1), note that the condition effect is consistent across genotypes. Although condition A has a different baseline for I, II, and III, the condition effect is a \log_2 fold change of about 2 for each genotype. Using a model with an interaction term `genotype:condition`, the interaction terms for genotype II and genotype III will be nearly 0. On the right side (Gene 2), we can see that the condition effect is not consistent across genotype. Here the main condition effect (the effect for the reference genotype I) is again 2. However, this time the interaction terms will be around 1 for genotype II and -4 for genotype III. This is because the condition effect is higher by 1 for genotype II compared to genotype I, and lower by 4 for genotype III compared to genotype I. The condition effect for genotype II (or III) is obtained by adding the main condition effect and the interaction term for that genotype. Such a plot can be made using the `plotCounts` function (Section 1.5.2).

Now we will continue to explain the use of interactions in order to test for *differences* in condition effects. We continue with the example of condition effects across three genotypes (I, II, and III). For a diagram of how interactions might look across genotypes please refer to Figure 9.

The key point to remember about designs with interaction terms is that, unlike for a design `~ genotype + condition`, where the condition effect represents the *overall* effect controlling for differences due to genotype, by adding `geno`

Differential analysis of count data – the DESeq2 package

`type:condition`, the main condition effect only represents the effect of condition for the *reference level* of genotype (I, or whichever level was defined by the user as the reference level). The interaction terms `genotypeII.conditionB` and `genotypeIII.conditionB` give the *difference* between the condition effect for a given genotype and the condition effect for the reference genotype.

This genotype-condition interaction example is examined in further detail in Example 3 in the help page for `results`, which can be found by typing `?results`. In particular, we show how to test for differences in the condition effect across genotype, and we show how to obtain the condition effect for non-reference genotypes. Note that in *DESeq2* version 1.10, the `DESeq` function will turn off log fold change shrinkage (setting `betaPrior=FALSE`), for designs which contain an interaction term. Turning off the log fold change shrinkage allows the software to use standard model matrices (as would be produced by `model.matrix`), where the interaction coefficients are easier to interpret.

3.4 Time-series experiments

There are a number of ways to analyze time-series experiments, depending on the biological question of interest. In order to test for any differences over multiple time points, one can use a design including the time factor, and then test using the likelihood ratio test as described in Section 3.5, where the time factor is removed in the reduced formula. For a control and treatment time series, one can use a design formula containing the condition factor, the time factor, and the interaction of the two. In this case, using the likelihood ratio test with a reduced model which does not contain the interaction terms will test whether the condition induces a change in gene expression at any time point after the reference level time point (time 0). An example of the later analysis is provided in an RNA-seq workflow on the Bioconductor website: <http://www.bioconductor.org/help/workflows/rnaseqGene/>.

3.5 Likelihood ratio test

DESeq2 offers two kinds of hypothesis tests: the Wald test, where we use the estimated standard error of a log₂ fold change to test if it is equal to zero, and the likelihood ratio test (LRT). The LRT examines two models for the counts, a *full* model with a certain number of terms and a *reduced* model, in which some of the terms of the *full* model are removed. The test determines if the increased likelihood of the data using the extra terms in the *full* model is more than expected if those extra terms are truly zero.

The LRT is therefore useful for testing multiple terms at once, for example testing 3 or more levels of a factor at once, or all interactions between two variables. The LRT for count data is conceptually similar to an analysis of variance (ANOVA) calculation in linear regression, except that in the case of the Negative Binomial GLM, we use an analysis of deviance (ANODEV), where the *deviance* captures the difference in likelihood between a full and a reduced model.

Differential analysis of count data – the DESeq2 package

The likelihood ratio test can be performed by specifying `test="LRT"` when using the `DESeq` function, and providing a reduced design formula, e.g. one in which a number of terms from `design(dds)` are removed. The degrees of freedom for the test is obtained from the difference between the number of parameters in the two models. A simple likelihood ratio test, if the full design was `condition` would look like:

```
dds <- DESeq(dds, test="LRT", reduced=~1)
res <- results(dds)
```

If the full design contained other variables, such as a batch variable, then the likelihood ratio test would look like:

```
dds <- DESeq(dds, test="LRT", reduced=~batch)
res <- results(dds)
```

3.6 Approach to count outliers

RNA-seq data sometimes contain isolated instances of very large counts that are apparently unrelated to the experimental or study design, and which may be considered outliers. There are many reasons why outliers can arise, including rare technical or experimental artifacts, read mapping problems in the case of genetically differing samples, and genuine, but rare biological events. In many cases, users appear primarily interested in genes that show a consistent behavior, and this is the reason why by default, genes that are affected by such outliers are set aside by `DESeq2`, or if there are sufficient samples, outlier counts are replaced for model fitting. These two behaviors are described below.

The `DESeq` function calculates, for every gene and for every sample, a diagnostic test for outliers called *Cook's distance*. Cook's distance is a measure of how much a single sample is influencing the fitted coefficients for a gene, and a large value of Cook's distance is intended to indicate an outlier count. The Cook's distances are stored as a matrix available in `assays(dds)[["cooks"]]`.

The `results` function automatically flags genes which contain a Cook's distance above a cutoff for samples which have 3 or more replicates. The p values and adjusted p values for these genes are set to `NA`. At least 3 replicates are required for flagging, as it is difficult to judge which sample might be an outlier with only 2 replicates. This filtering can be turned off with `results(dds, cooksCutoff=FALSE)`.

With many degrees of freedom – i. e., many more samples than number of parameters to be estimated – it is undesirable to remove entire genes from the analysis just because their data include a single count outlier. When there are 7 or more replicates for a given sample, the `DESeq` function will automatically replace counts with large Cook's distance with the trimmed mean over all samples, scaled up by the size factor or normalization factor for that sample. This approach is conservative, it will not lead to false positives, as it replaces the outlier value with the

Differential analysis of count data – the DESeq2 package

value predicted by the null hypothesis. This outlier replacement only occurs when there are 7 or more replicates, and can be turned off with `DESeq(dds, minReplicatesForReplace=Inf)`.

The default Cook's distance cutoff for the two behaviors described above depends on the sample size and number of parameters to be estimated. The default is to use the 99% quantile of the $F(p, m - p)$ distribution (with p the number of parameters including the intercept and m number of samples). The default for gene flagging can be modified using the `cooksCutoff` argument to the `results` function. For outlier replacement, DESeq preserves the original counts in `counts(dds)` saving the replacement counts as a matrix named `replaceCounts` in `assays(dds)`. Note that with continuous variables in the design, outlier detection and replacement is not automatically performed, as our current methods involve a robust estimation of within-group variance which does not extend easily to continuous covariates. However, users can examine the Cook's distances in `assays(dds)[["cooks"]]`, in order to perform manual visualization and filtering if necessary.

Note on many outliers: if there are very many outliers (e.g. many hundreds or thousands) reported by `summary(res)`, one might consider further exploration to see if a single sample or a few samples should be removed due to low quality. The automatic outlier filtering/replacement is most useful in situations which the number of outliers is limited. When there are thousands of reported outliers, it might make more sense to turn off the outlier filtering/replacement (DESeq with `minReplicatesForReplace=Inf` and `results` with `cooksCutoff=FALSE`) and perform manual inspection: First it would be advantageous to make a PCA plot using the code example in Section 2.2.3 to spot individual sample outliers; Second, one can make a boxplot of the Cook's distances to see if one sample is consistently higher than others:

```
par(mar=c(8,5,2,2))
boxplot(log10(assays(dds)[["cooks"]]), range=0, las=2)
```

3.7 Dispersion plot and fitting alternatives

Plotting the dispersion estimates is a useful diagnostic. The dispersion plot in Figure 11 is typical, with the final estimates shrunk from the gene-wise estimates towards the fitted estimates. Some gene-wise estimates are flagged as outliers and not shrunk towards the fitted value, (this outlier detection is described in the manual page for `estimateDispersionsMAP`). The amount of shrinkage can be more or less than seen here, depending on the sample size, the number of coefficients, the row mean and the variability of the gene-wise estimates.

Differential analysis of count data – the DESeq2 package




Figure 10: Boxplot of Cook's distances. Here we can look to see if one sample has much higher Cook's distances than the other samples. In this case, the samples all have comparable range of Cook's distances.




Figure 11: Dispersion plot. The dispersion estimate plot shows the gene-wise estimates (black), the fitted values (red), and the final maximum *a posteriori* estimates used in testing (blue).

```
plotDispEsts(dds)
```

Differential analysis of count data – the DESeq2 package

3.7.1 Local or mean dispersion fit

A local smoothed dispersion fit is automatically substituted in the case that the parametric curve doesn't fit the observed dispersion mean relationship. This can be prespecified by providing the argument `fitType="local"` to either `DESeq` or `estimateDispersions`. Additionally, using the mean of gene-wise dispersion estimates as the fitted value can be specified by providing the argument `fitType="mean"`.

3.7.2 Supply a custom dispersion fit

Any fitted values can be provided during dispersion estimation, using the lower-level functions described in the manual page for `estimateDispersionsGeneEst`. In the code chunk below, we store the gene-wise estimates which were already calculated and saved in the metadata column `dispGeneEst`. Then we calculate the median value of the dispersion estimates above a threshold, and save these values as the fitted dispersions, using the replacement function for `dispersionFunction`. In the last line, the function `estimateDispersionsMAP`, uses the fitted dispersions to generate maximum *a posteriori* (MAP) estimates of dispersion.

```
ddsCustom <- dds
useForMedian <- mcols(ddsCustom)$dispGeneEst > 1e-7
medianDisp <- median(mcols(ddsCustom)$dispGeneEst[useForMedian],
                    na.rm=TRUE)
dispersionFunction(ddsCustom) <- function(mu) medianDisp
ddsCustom <- estimateDispersionsMAP(ddsCustom)
```

3.8 Independent filtering of results

The `results` function of the `DESeq2` package performs independent filtering by default using the mean of normalized counts as a filter statistic. A threshold on the filter statistic is found which optimizes the number of adjusted p values lower than a significance level `alpha` (we use the standard variable name for significance level, though it is unrelated to the dispersion parameter α). The theory behind independent filtering is discussed in greater detail in Section 4.7. The adjusted p values for the genes which do not pass the filter threshold are set to `NA`.

The independent filtering is performed using the `filtered_p` function of the `gene-filter` package, and all of the arguments of `filtered_p` can be passed to the `results` function. The filter threshold value and the number of rejections at each quantile of the filter statistic are available as metadata of the object returned by `results`. For example, we can visualize the optimization by plotting the `filterNumRej` attribute of the results object, as seen in Figure 12.

```
metadata(res)$alpha
## [1] 0.1
```

Differential analysis of count data – the DESeq2 package

```
metadata(res)$filterThreshold
## 27.1%
## 5.56

plot(metadata(res)$filterNumRej,
      type="b", ylab="number of rejections",
      xlab="quantiles of filter")
lines(metadata(res)$lo.fit, col="red")
abline(v=metadata(res)$filterTheta)
```




Figure 12: Independent filtering. The `results` function maximizes the number of rejections (adjusted p value less than a significance level), over the quantiles of a filter statistic (the mean of normalized counts). The threshold chosen (vertical line) is the lowest quantile of the filter for which the number of rejections is within 1 residual standard deviation to the peak of a curve fit to the number of rejections over the filter quantiles.

Independent filtering can be turned off by setting `independentFiltering` to `FALSE`.

```
resNoFilt <- results(dds, independentFiltering=FALSE)
addmargins(table(filtering=(res$padj < .1),
                 noFiltering=(resNoFilt$padj < .1)))

##           noFiltering
## filtering FALSE TRUE Sum
##    FALSE  7418    0 7418
##    TRUE    111  943 1054
##    Sum    7529  943 8472
```

3.9 Tests of log2 fold change above or below a threshold

It is also possible to provide thresholds for constructing Wald tests of significance. Two arguments to the `results` function allow for threshold-based Wald tests: `lfcThreshold`, which takes a numeric of a non-negative threshold value, and `altHypothesis`, which specifies the kind of test. Note that the *alternative hypothesis* is specified by the user, i.e. those genes which the user is interested in finding, and the test provides p values for the null hypothesis, the complement of the set defined by the alternative. The `altHypothesis` argument can take one of the following four values, where β is the log2 fold change specified by the `name` argument:

- `greaterAbs` - $|\beta| > \text{lfcThreshold}$ - tests are two-tailed
- `lessAbs` - $|\beta| < \text{lfcThreshold}$ - p values are the maximum of the upper and lower tests
- `greater` - $\beta > \text{lfcThreshold}$
- `less` - $\beta < -\text{lfcThreshold}$

The test `altHypothesis="lessAbs"` requires that the user have run `DESeq` with the argument `betaPrior=FALSE`. To understand the reason for this requirement, consider that during hypothesis testing, the null hypothesis is favored unless the data provide strong evidence to reject the null. For this test, including a zero-centered prior on log fold change would favor the alternative hypothesis, shrinking log fold changes toward zero. Removing the prior on log fold changes for tests of small log fold change allows for detection of only those genes where the data alone provides evidence against the null.

The four possible values of `altHypothesis` are demonstrated in the following code and visually by MA-plots in Figure 13. First we run `DESeq` and specify `betaPrior=FALSE` in order to demonstrate `altHypothesis="lessAbs"`.

```
ddsNoPrior <- DESeq(dds, betaPrior=FALSE)
```

In order to produce results tables for the following tests, the same arguments (except `ylim`) would be provided to the `results` function.

```
par(mfrow=c(2,2),mar=c(2,2,1,1))
yl <- c(-2.5,2.5)

resGA <- results(dds, lfcThreshold=.5, altHypothesis="greaterAbs")
resLA <- results(ddsNoPrior, lfcThreshold=.5, altHypothesis="lessAbs")
resG <- results(dds, lfcThreshold=.5, altHypothesis="greater")
resL <- results(dds, lfcThreshold=.5, altHypothesis="less")

plotMA(resGA, ylim=yl)
abline(h=c(-.5, .5),col="dodgerblue",lwd=2)
plotMA(resLA, ylim=yl)
abline(h=c(-.5, .5),col="dodgerblue",lwd=2)
plotMA(resG, ylim=yl)
```

Differential analysis of count data – the DESeq2 package

```
abline(h=.5,col="dodgerblue",lwd=2)
plotMA(resL, ylim=yl)
abline(h=-.5,col="dodgerblue",lwd=2)
```




Figure 13: MA-plots of tests of log₂ fold change with respect to a threshold value. Going left to right across rows, the tests are for `altHypothesis = "greaterAbs"`, `"lessAbs"`, `"greater"`, and `"less"`.

3.10 Access to all calculated values

All row-wise calculated values (intermediate dispersion calculations, coefficients, standard errors, etc.) are stored in the `DESeqDataSet` object, e.g. `dds` in this vignette. These values are accessible by calling `mcols` on `dds`. Descriptions of the columns are accessible by two calls to `mcols`.

```
mcols(dds,use.names=TRUE)[1:4,1:4]
## DataFrame with 4 rows and 4 columns
##           gene baseMean baseVar allZero
##           <factor> <numeric> <numeric> <logical>
```

Differential analysis of count data – the DESeq2 package

```
## FBgn0000008 FBgn0000008 95.144 2.25e+02 FALSE
## FBgn0000014 FBgn0000014 1.057 2.96e+00 FALSE
## FBgn0000015 FBgn0000015 0.847 1.01e+00 FALSE
## FBgn0000017 FBgn0000017 4352.593 3.62e+05 FALSE

# here using substr() only for display purposes
substr(names(mcols(dds)),1,10)

## [1] "gene" "baseMean" "baseVar" "allZero"
## [5] "dispGeneEs" "dispFit" "dispersion" "dispIter"
## [9] "dispOutlie" "dispMAP" "Intercept" "conditionu"
## [13] "conditiont" "SE_Interce" "SE_conditi" "SE_conditi"
## [17] "MLE_Interc" "MLE_condit" "WaldStatis" "WaldStatis"
## [21] "WaldStatis" "WaldPvalue" "WaldPvalue" "WaldPvalue"
## [25] "betaConv" "betaIter" "deviance" "maxCooks"

mcols(mcols(dds), use.names=TRUE)[1:4,]

## DataFrame with 4 rows and 2 columns
## type
## <character>
## gene input
## baseMean intermediate
## baseVar intermediate
## allZero intermediate
## description
## <character>
## gene
## baseMean mean of normalized counts for all samples
## baseVar variance of normalized counts for all samples
## allZero all counts for a gene are zero
```

The mean values $\mu_{ij} = s_j q_{ij}$ and the Cook's distances for each gene and sample are stored as matrices in the assays slot:

```
head(assays(dds)[["mu"]])

## treated1 treated2 treated3 untreated1 untreated2
## FBgn0000008 154.37 71.853 78.590 107.305 169.03
## FBgn0000014 1.83 0.850 0.930 1.292 2.04
## FBgn0000015 1.28 0.595 0.651 0.948 1.49
## FBgn0000017 6491.79 3021.568 3304.895 5276.388 8311.41
## FBgn0000018 660.84 307.586 336.428 491.323 773.93
## FBgn0000024 10.83 5.040 5.513 7.195 11.33
## untreated3 untreated4
## FBgn0000008 61.233 70.862
## FBgn0000014 0.737 0.853
## FBgn0000015 0.541 0.626
## FBgn0000017 3010.958 3484.435
## FBgn0000018 280.372 324.461
```


Differential analysis of count data – the DESeq2 package

```
## FBgn0000024      4.106      4.751

head(assays(dds)[["cooks"]])

##           treated1 treated2 treated3 untreated1 untreated2
## FBgn0000008  0.08803  0.30431  0.077569    0.0984    0.01370
## FBgn0000014  1.15575  0.26448  0.304818    2.3881    0.13034
## FBgn0000015  0.00989  0.12687  0.133641    0.1076    0.01755
## FBgn0000017  0.01853  0.00256  0.000717    0.0749    0.01318
## FBgn0000018  0.08725  0.00627  0.059497    0.1914    0.00168
## FBgn0000024  0.02330  0.19266  0.012765    0.2081    0.00216
##           untreated3 untreated4
## FBgn0000008    0.1888    0.000541
## FBgn0000014    0.1355    0.166619
## FBgn0000015    0.0436    0.344032
## FBgn0000017    0.1835    0.050460
## FBgn0000018    0.0736    0.009682
## FBgn0000024    0.0418    0.099273
```

The dispersions α_i can be accessed with the `dispersions` function.

```
head(dispersions(dds))

## [1] 0.0304 2.8630 2.2096 0.0128 0.0156 0.2386

# which is the same as
head(mcols(dds)$dispersion)

## [1] 0.0304 2.8630 2.2096 0.0128 0.0156 0.2386
```

The size factors s_j are accessible via `sizeFactors`:

```
sizeFactors(dds)

##   treated1   treated2   treated3 untreated1 untreated2 untreated3
##   1.636     0.761     0.833     1.138     1.793     0.650
## untreated4
##   0.752
```

For advanced users, we also include a convenience function `coef` for extracting the matrix of coefficients $[\beta_{i,r}]$ for all genes i and parameters r , as in the formula in Section 4.1. This function can also return a matrix of standard errors, see `?coef`. The columns of this matrix correspond to the effects returned by `resultsNames`. Note that the `results` function is best for building results tables with p values and adjusted p values.

```
head(coef(dds))

##           Intercept conditionuntreated conditiontreated
## FBgn0000008    6.560           -0.000875           0.000876
## FBgn0000014    0.171           0.011678           -0.011677
```

Differential analysis of count data – the DESeq2 package

```
## FBgn0000015    -0.309          0.045671    -0.045671
## FBgn0000017    12.067          0.111950    -0.111949
## FBgn0000018     8.706          0.047663    -0.047662
## FBgn0000024     2.694          -0.033432     0.033432
```

The beta prior variance σ_r^2 is stored as an attribute of the *DESeqDataSet*:

```
attr(dds, "betaPriorVar")
##          Intercept conditionuntreated  conditiontreated
##          1.0e+06          1.1e-01          1.1e-01
```

The dispersion prior variance σ_d^2 is stored as an attribute of the dispersion function:

```
dispersionFunction(dds)
## function (q)
## coefs[1] + coefs[2]/q
## <environment: 0xe210658>
## attr("coefficients")
## asymptDisp  extraPois
##          0.014          2.721
## attr("fitType")
## [1] "parametric"
## attr("varLogDispEsts")
## [1] 0.989
## attr("dispPriorVar")
## [1] 0.499
attr(dispersionFunction(dds), "dispPriorVar")
## [1] 0.499
```

The version of *DESeq2* which was used to construct the *DESeqDataSet* object, or the version used when *DESeq* was run, is stored here:

```
metadata(dds)[["version"]]
## [1] '1.14.0'
```

3.11 Sample-/gene-dependent normalization factors

In some experiments, there might be gene-dependent dependencies which vary across samples. For instance, GC-content bias or length bias might vary across samples coming from different labs or processed at different times. We use the terms “normalization factors” for a gene \times sample matrix, and “size factors” for a single number per sample. Incorporating normalization factors, the mean parameter μ_{ij} from Section 4.1 becomes:

Differential analysis of count data – the DESeq2 package

$$\mu_{ij} = NF_{ij}q_{ij}$$

with normalization factor matrix NF having the same dimensions as the counts matrix K . This matrix can be incorporated as shown below. We recommend providing a matrix with row-wise geometric means of 1, so that the mean of normalized counts for a gene is close to the mean of the unnormalized counts. This can be accomplished by dividing out the current row geometric means.

```
normFactors <- normFactors / exp(rowMeans(log(normFactors)))  
normalizationFactors(dds) <- normFactors
```

These steps then replace `estimateSizeFactors` in the steps described in Section 3.1. Normalization factors, if present, will always be used in the place of size factors.

The methods provided by the `cqn` or `EDASeq` packages can help correct for GC or length biases. They both describe in their vignettes how to create matrices which can be used by `DESeq2`. From the formula above, we see that normalization factors should be on the scale of the counts, like size factors, and unlike offsets which are typically on the scale of the predictors (i.e. the logarithmic scale for the negative binomial GLM). At the time of writing, the transformation from the matrices provided by these packages should be:

```
cqnOffset <- cqnObject$glm.offset  
cqnNormFactors <- exp(cqnOffset)  
EDASeqNormFactors <- exp(-1 * EDASeqOffset)
```

3.12 “Model matrix not full rank”

While most experimental designs run easily using design formula, some design formulas can cause problems and result in the `DESeq` function returning an error with the text: “the model matrix is not full rank, so the model cannot be fit as specified.” There are two main reasons for this problem: either one or more columns in the model matrix are linear combinations of other columns, or there are levels of factors or combinations of levels of multiple factors which are missing samples. We address these two problems below and discuss possible solutions:

3.12.1 Linear combinations

The simplest case is the linear combination, or linear dependency problem, when two variables contain exactly the same information, such as in the following sample table. The software cannot fit an effect for `batch` and `condition`, because they produce identical columns in the model matrix. This is also referred to as “perfect confounding”. A unique solution of coefficients (the β_i in the formula in Section 4.1) is not possible.

Differential analysis of count data – the DESeq2 package

```
##  batch condition
## 1    1          A
## 2    1          A
## 3    2          B
## 4    2          B
```

Another situation which will cause problems is when the variables are not identical, but one variable can be formed by the combination of other factor levels. In the following example, the effect of batch 2 vs 1 cannot be fit because it is identical to a column in the model matrix which represents the condition C vs A effect.

```
##  batch condition
## 1    1          A
## 2    1          A
## 3    1          B
## 4    1          B
## 5    2          C
## 6    2          C
```

In both of these cases above, the batch effect cannot be fit and must be removed from the model formula. There is just no way to tell apart the condition effects and the batch effects. The options are either to assume there is no batch effect (which we know is highly unlikely given the literature on batch effects in sequencing datasets) or to repeat the experiment and properly balance the conditions across batches. A balanced design would look like:

```
##  batch condition
## 1    1          A
## 2    1          B
## 3    1          C
## 4    2          A
## 5    2          B
## 6    2          C
```

Finally, there is a case where we can in fact perform inference. Consider an experiment with grouped individuals, where we seek to test the group-specific effect of a treatment, while controlling for individual effects. A simple example of such a design is:

```
(coldata <- data.frame(grp=factor(rep(c("X","Y"),each=4)),
                      ind=factor(rep(1:4,each=2)),
                      cnd=factor(rep(c("A","B"),4))))

##  grp ind cnd
## 1  X  1  A
## 2  X  1  B
## 3  X  2  A
## 4  X  2  B
```

Differential analysis of count data – the DESeq2 package

```
## 5 Y 3 A
## 6 Y 3 B
## 7 Y 4 A
## 8 Y 4 B
```

This design can be analyzed by *DESeq2* but requires a bit of refactoring in order to fit the model terms. Here we will use a trick described in the *edgeR* user guide, from the section “Comparisons Both Between and Within Subjects”. If we try to analyze with a formula such as, `~ ind + grp*cnd`, we will obtain an error, because the effect for group is a linear combination of the individuals.

However, the following steps allow for an analysis of group-specific condition effects, while controlling for differences in individual. For object construction, use a dummy design, such as `~ 1`. Then add a column `ind.n` which distinguishes the individuals “nested” within a group. Here, we add this column to `coldata`, but in practice you would add this column to `dds`.

```
coldata$ind.n <- factor(rep(rep(1:2,each=2),2))
coldata
##  grp ind cnd ind.n
## 1  X  1  A    1
## 2  X  1  B    1
## 3  X  2  A    2
## 4  X  2  B    2
## 5  Y  3  A    1
## 6  Y  3  B    1
## 7  Y  4  A    2
## 8  Y  4  B    2
```

Now we can reassign our *DESeqDataSet* a design of `~ grp + grp:ind.n + grp:cnd`, before we call *DESeq*. This new design will result in the following model matrix:

```
model.matrix(~ grp + grp:ind.n + grp:cnd, coldata)
##  (Intercept) grpY grpX:ind.n2 grpY:ind.n2 grpX:cndB grpY:cndB
## 1          1  0      0      0      0      0
## 2          1  0      0      0      1      0
## 3          1  0      1      0      0      0
## 4          1  0      1      0      1      0
## 5          1  1      0      0      0      0
## 6          1  1      0      0      0      1
## 7          1  1      0      1      0      0
## 8          1  1      0      1      0      1
## attr(,"assign")
## [1] 0 1 2 2 3 3
## attr(,"contrasts")
## attr(,"contrasts")$grp
## [1] "contr.treatment"
```

Differential analysis of count data – the DESeq2 package

```
##  
## attr(,"contrasts")$ind.n  
## [1] "contr.treatment"  
##  
## attr(,"contrasts")$cnd  
## [1] "contr.treatment"
```

where the terms `grpX.cndB` and `grpY.cndB` give the group-specific condition effects. These can be extracted using `results` with the `name` argument. Furthermore, `grpX.cndB` and `grpY.cndB` can be contrasted using the `contrast` argument, in order to test if the condition effect is different across group:

```
results(dds, contrast=list("grpY.cndB", "grpX.cndB"))
```

3.12.2 Levels without samples

The base R function for creating model matrices will produce a column of zeros if a level is missing from a factor or a combination of levels is missing from an interaction of factors. The solution to the first case is to call `droplevels` on the column, which will remove levels without samples. This was shown in the beginning of this vignette.

The second case is also solvable, by manually editing the model matrix, and then providing this to `DESeq`. Here we construct an example dataset to illustrate:

```
group <- factor(rep(1:3,each=6))  
condition <- factor(rep(rep(c("A","B","C"),each=2),3))  
(d <- data.frame(group, condition)[-c(17,18),])  
  
##   group condition  
## 1     1         A  
## 2     1         A  
## 3     1         B  
## 4     1         B  
## 5     1         C  
## 6     1         C  
## 7     2         A  
## 8     2         A  
## 9     2         B  
## 10    2         B  
## 11    2         C  
## 12    2         C  
## 13    3         A  
## 14    3         A  
## 15    3         B  
## 16    3         B
```

Differential analysis of count data – the DESeq2 package

Note that if we try to estimate all interaction terms, we introduce a column with all zeros, as there are no condition C samples for group 3. (Here, `unnname` is used to display the matrix concisely.)

```
m1 <- model.matrix(~ condition*group, d)
colnames(m1)

## [1] "(Intercept)"      "conditionB"      "conditionC"
## [4] "group2"            "group3"          "conditionB:group2"
## [7] "conditionC:group2" "conditionB:group3" "conditionC:group3"

unnname(m1)

##      [,1] [,2] [,3] [,4] [,5] [,6] [,7] [,8] [,9]
## [1,]  1  0  0  0  0  0  0  0  0
## [2,]  1  0  0  0  0  0  0  0  0
## [3,]  1  1  0  0  0  0  0  0  0
## [4,]  1  1  0  0  0  0  0  0  0
## [5,]  1  0  1  0  0  0  0  0  0
## [6,]  1  0  1  0  0  0  0  0  0
## [7,]  1  0  0  1  0  0  0  0  0
## [8,]  1  0  0  1  0  0  0  0  0
## [9,]  1  1  0  1  0  1  0  0  0
## [10,] 1  1  0  1  0  1  0  0  0
## [11,] 1  0  1  1  0  0  1  0  0
## [12,] 1  0  1  1  0  0  1  0  0
## [13,] 1  0  0  0  1  0  0  0  0
## [14,] 1  0  0  0  1  0  0  0  0
## [15,] 1  1  0  0  1  0  0  1  0
## [16,] 1  1  0  0  1  0  0  1  0
## attr(,"assign")
## [1] 0 1 1 2 2 3 3 3 3
## attr(,"contrasts")
## attr(,"contrasts")$condition
## [1] "contr.treatment"
##
## attr(,"contrasts")$group
## [1] "contr.treatment"
```

We can remove this column like so:

```
m1 <- m1[, -9]
unnname(m1)

##      [,1] [,2] [,3] [,4] [,5] [,6] [,7] [,8]
## [1,]  1  0  0  0  0  0  0  0
## [2,]  1  0  0  0  0  0  0  0
## [3,]  1  1  0  0  0  0  0  0
## [4,]  1  1  0  0  0  0  0  0
## [5,]  1  0  1  0  0  0  0  0
## [6,]  1  0  1  0  0  0  0  0
```

Differential analysis of count data – the DESeq2 package

```
## [7,] 1 0 0 1 0 0 0 0
## [8,] 1 0 0 1 0 0 0 0
## [9,] 1 1 0 1 0 1 0 0
## [10,] 1 1 0 1 0 1 0 0
## [11,] 1 0 1 1 0 0 1 0
## [12,] 1 0 1 1 0 0 1 0
## [13,] 1 0 0 0 1 0 0 0
## [14,] 1 0 0 0 1 0 0 0
## [15,] 1 1 0 0 1 0 0 1
## [16,] 1 1 0 0 1 0 0 1
```

Now this matrix `m1` can be provided to the `full` argument of `DESeq`. For a likelihood ratio test of interactions, a model matrix using a reduced design such as `~ condition + group` can be given to the `reduced` argument. Wald tests can also be generated instead of the likelihood ratio test, but for user-supplied model matrices, the argument `betaPrior` must be set to `FALSE`.

4 Theory behind DESeq2

4.1 The DESeq2 model

The *DESeq2* model and all the steps taken in the software are described in detail in our publication [1], and we include the formula and descriptions in this section as well. The differential expression analysis in *DESeq2* uses a generalized linear model of the form:

$$\begin{aligned} K_{ij} &\sim \text{NB}(\mu_{ij}, \alpha_i) \\ \mu_{ij} &= s_j q_{ij} \\ \log_2(q_{ij}) &= x_j \beta_i \end{aligned}$$

where counts K_{ij} for gene i , sample j are modeled using a negative binomial distribution with fitted mean μ_{ij} and a gene-specific dispersion parameter α_i . The fitted mean is composed of a sample-specific size factor s_j ⁵ and a parameter q_{ij} proportional to the expected true concentration of fragments for sample j . The coefficients β_i give the log₂ fold changes for gene i for each column of the model matrix X .

The dispersion parameter α_i defines the relationship between the variance of the observed count and its mean value. In other words, how far do we expect the observed count will be from the mean value, which depends both on the size factor s_j and the covariate-dependent part q_{ij} as defined above.

$$\text{Var}(K_{ij}) = E[(K_{ij} - \mu_{ij})^2] = \mu_{ij} + \alpha_i \mu_{ij}^2$$

The log₂ fold changes in β_i are the maximum *a posteriori* estimates after incorporating a zero-centered Normal prior – in the software referred to as a β -prior – hence *DESeq2* provides “moderated” log₂ fold change estimates. Dispersions are estimated using expected mean values from the maximum likelihood estimate of log₂ fold changes, and optimizing the Cox-Reid adjusted profile likelihood, as first implemented for RNA-seq data in *edgeR* [15, 16]. The steps performed by the *DESeq* function are documented in its manual page; briefly, they are:

1. estimation of size factors s_j by `estimateSizeFactors`
2. estimation of dispersion α_i by `estimateDispersions`
3. negative binomial GLM fitting for β_i and Wald statistics by `nbinomWaldTest`

For access to all the values calculated during these steps, see Section 3.10

4.2 Changes compared to the *DESeq* package

The main changes in the package *DESeq2*, compared to the (older) version *DESeq*, are as follows:

⁵The model can be generalized to use sample- and gene-dependent normalization factors, see Appendix 3.11.

Differential analysis of count data – the DESeq2 package

- *RangedSummarizedExperiment* is used as the superclass for storage of input data, intermediate calculations and results.
- Maximum *a posteriori* estimation of GLM coefficients incorporating a zero-centered Normal prior with variance estimated from data (equivalent to Tikhonov/ridge regularization). This adjustment has little effect on genes with high counts, yet it helps to moderate the otherwise large variance in log₂ fold change estimates for genes with low counts or highly variable counts.
- Maximum *a posteriori* estimation of dispersion replaces the `sharingMode` options `fit-only` or `maximum` of the previous version of the package. This is similar to the dispersion estimation methods of DSS [17].
- All estimation and inference is based on the generalized linear model, which includes the two condition case (previously the *exact test* was used).
- The Wald test for significance of GLM coefficients is provided as the default inference method, with the likelihood ratio test of the previous version still available.
- It is possible to provide a matrix of sample-/gene-dependent normalization factors (Section 3.11).
- Automatic independent filtering on the mean of normalized counts (Section 4.7).
- Automatic outlier detection and handling (Section 4.4).

4.3 Methods changes since the 2014 DESeq2 paper

- For the calculation of the beta prior variance, instead of matching the empirical quantile to the quantile of a Normal distribution, *DESeq2*() now uses the weighted quantile function of the *Hmisc* package. The weighting is described in the manual page for `nbinomWaldTest`. The weights are the inverse of the expected variance of log counts (as used in the diagonals of the matrix *W* in the GLM). The effect of the change is that the estimated prior variance is robust against noisy estimates of log fold change from genes with very small counts. This change was introduced in version 1.6 (October 2014).
- For designs with interaction terms, the solution described in the paper is no longer used (log fold change shrinkage only applied to interaction terms). Instead, *DESeq2* now turns off log fold change shrinkage for all terms if an interaction term is present (`betaPrior=FALSE`). While the inference on interaction terms was correct with `betaPrior=TRUE`, the interpretation of the individual terms and the extraction of contrasts was too confusing. This change was introduced in version 1.10 (October 2015).
- A small change to the independent filtering routine: instead of taking the quantile of the filter (the mean of normalized counts) which directly *maximizes* the number of rejections, the threshold chosen is the lowest quantile of the filter for which the number of rejections is close to the peak of a curve fit

Differential analysis of count data – the DESeq2 package

to the number of rejections over the filter quantiles. “Close to” is defined as within 1 residual standard deviation. This change was introduced in version 1.10 (October 2015).

For a list of all changes since version 1.0.0, see the NEWS file included in the package.

4.4 Count outlier detection

DESeq2 relies on the negative binomial distribution to make estimates and perform statistical inference on differences. While the negative binomial is versatile in having a mean and dispersion parameter, extreme counts in individual samples might not fit well to the negative binomial. For this reason, we perform automatic detection of count outliers. We use Cook’s distance, which is a measure of how much the fitted coefficients would change if an individual sample were removed [18]. For more on the implementation of Cook’s distance see Section 3.6 and the manual page for the `results` function. Below we plot the maximum value of Cook’s distance for each row over the rank of the test statistic to justify its use as a filtering criterion.

```
W <- res$stat
maxCooks <- apply(assays(dds)[["cooks"]], 1, max)
idx <- !is.na(W)
plot(rank(W[idx]), maxCooks[idx], xlab="rank of Wald statistic",
     ylab="maximum Cook's distance per gene",
     ylim=c(0,5), cex=.4, col=rgb(0,0,0,.3))
m <- ncol(dds)
p <- 3
abline(h=qf(.99, p, m - p))
```




Figure 14: Cook’s distance. Plot of the maximum Cook’s distance per gene over the rank of the Wald statistics for the condition. The two regions with small Cook’s distances are genes with a single count in one sample. The horizontal line is the default cutoff used for 7 samples and 3 estimated parameters.

4.5 Contrasts

Contrasts can be calculated for a *DESeqDataSet* object for which the GLM coefficients have already been fit using the Wald test steps (`DESeq` with `test="Wald"` or using `nbinomWaldTest`). The vector of coefficients β is left multiplied by the contrast vector c to form the numerator of the test statistic. The denominator is formed by multiplying the covariance matrix Σ for the coefficients on either side by the contrast vector c . The square root of this product is an estimate of the standard error for the contrast. The contrast statistic is then compared to a normal distribution as are the Wald statistics for the *DESeq2* package.

$$W = \frac{c^t \beta}{\sqrt{c^t \Sigma c}}$$

4.6 Expanded model matrices

DESeq2 uses “expanded model matrices” with the log₂ fold change prior, in order to produce shrunken log₂ fold change estimates and test results which are independent of the choice of reference level. Another way of saying this is that the shrinkage is *symmetric* with respect to all the levels of the factors in the design. The expanded model matrices differ from the standard model matrices, in that they have an indicator column (and therefore a coefficient) for each level of factors in the design formula in addition to an intercept. Note that in version 1.10 and onward, standard model matrices are used for designs with interaction terms, as the shrinkage of log₂ fold changes is not recommended for these designs.

The expanded model matrices are not full rank, but a coefficient vector β_i can still be found due to the zero-centered prior on non-intercept coefficients. The prior variance for the log₂ fold changes is calculated by first generating maximum likelihood estimates for a standard model matrix. The prior variance for each level of a factor is then set as the average of the mean squared maximum likelihood estimates for each level and every possible contrast, such that that this prior value will be reference-level-independent. The `contrast` argument of the `results` function is used in order to generate comparisons of interest.

4.7 Independent filtering and multiple testing

4.7.1 Filtering criteria

The goal of independent filtering is to filter out those tests from the procedure that have no, or little chance of showing significant evidence, without even looking at their test statistic. Typically, this results in increased detection power at the same experiment-wide type I error. Here, we measure experiment-wide type I error in terms of the false discovery rate.

A good choice for a filtering criterion is one that

1. is statistically independent from the test statistic under the null hypothesis,

Differential analysis of count data – the DESeq2 package

2. is correlated with the test statistic under the alternative, and
3. does not notably change the dependence structure –if there is any– between the tests that pass the filter, compared to the dependence structure between the tests before filtering.

The benefit from filtering relies on property 2, and we will explore it further in Section 4.7.2. Its statistical validity relies on property 1 – which is simple to formally prove for many combinations of filter criteria with test statistics– and 3, which is less easy to theoretically imply from first principles, but rarely a problem in practice. We refer to [19] for further discussion of this topic.

A simple filtering criterion readily available in the results object is the mean of normalized counts irrespective of biological condition (Figure 15), and so this is the criterion which is used automatically by the `results` function to perform independent filtering. Genes with very low counts are not likely to see significant differences typically due to high dispersion. For example, we can plot the $-\log_{10} p$ values from all genes over the normalized mean counts.

```
plot(res$baseMean+1, -log10(res$pvalue),  
     log="x", xlab="mean of normalized counts",  
     ylab=expression(-log[10](pvalue)),  
     ylim=c(0,30),  
     cex=.4, col=rgb(0,0,0,.3))
```




Figure 15: Mean counts as a filter statistic. The mean of normalized counts provides an independent statistic for filtering the tests. It is independent because the information about the variables in the design formula is not used. By filtering out genes which fall on the left side of the plot, the majority of the low p values are kept.

Differential analysis of count data – the DESeq2 package

4.7.2 Why does it work?

Consider the p value histogram in Figure 16. It shows how the filtering ameliorates the multiple testing problem – and thus the severity of a multiple testing adjustment – by removing a background set of hypotheses whose p values are distributed more or less uniformly in $[0,1]$.

```
use <- res$baseMean > metadata(res)$filterThreshold
h1 <- hist(res$pvalue[!use], breaks=0:50/50, plot=FALSE)
h2 <- hist(res$pvalue[use], breaks=0:50/50, plot=FALSE)
colori <- c(`do not pass`="khaki", `pass`="powderblue")
```

```
barplot(height = rbind(h1$counts, h2$counts), beside = FALSE,
        col = colori, space = 0, main = "", ylab="frequency")
text(x = c(0, length(h1$counts)), y = 0, label = paste(c(0,1)),
     adj = c(0.5,1.7), xpd=NA)
legend("topright", fill=rev(colori), legend=rev(names(colori)))
```




Figure 16: Histogram of p values for all tests. The area shaded in blue indicates the subset of those that pass the filtering, the area in khaki those that do not pass.

5 Frequently asked questions

5.1 How can I get support for DESeq2?

We welcome questions about our software, and want to ensure that we eliminate issues if and when they appear. We have a few requests to optimize the process:

- all questions should take place on the Bioconductor support site: <https://support.bioconductor.org>, which serves as a repository of questions and answers. This helps to save the developers' time in responding to similar questions. Make sure to tag your post with "deseq2". It is often very helpful in addition to describe the aim of your experiment.

Differential analysis of count data – the DESeq2 package

- before posting, first search the Bioconductor support site mentioned above for past threads which might have answered your question.
- if you have a question about the behavior of a function, read the sections of the manual page for this function by typing a question mark and the function name, e.g. `?results`. We spend a lot of time documenting individual functions and the exact steps that the software is performing.
- include all of your R code, especially the creation of the *DESeqDataSet* and the design formula. Include complete warning or error messages, and conclude your message with the full output of `sessionInfo()`.
- if possible, include the output of `as.data.frame(colData(dds))`, so that we can have a sense of the experimental setup. If this contains confidential information, you can replace the levels of those factors using `levels()`.

5.2 Why are some p values set to NA?

See the details in Section 1.5.3.

5.3 How can I get unfiltered DESeq results?

Users can obtain unfiltered GLM results, i.e. without outlier removal or independent filtering with the following call:

```
dds <- DESeq(dds, minReplicatesForReplace=Inf)
res <- results(dds, cooksCutoff=FALSE, independentFiltering=FALSE)
```

In this case, the only p values set to NA are those from genes with all counts equal to zero.

5.4 How do I use the variance stabilized or rlog transformed data for differential testing?

The variance stabilizing and rlog transformations are provided for applications other than differential testing, for example clustering of samples or other machine learning applications. For differential testing we recommend the `DESeq` function applied to raw counts as outlined in Section 1.4.

5.5 Can I use DESeq2 to analyze paired samples?

Yes, you should use a multi-factor design which includes the sample information as a term in the design formula. This will account for differences between the samples while estimating the effect due to the condition. The condition of interest should go at the end of the design formula. See Section 1.6.

5.6 If I have multiple groups, should I run all together or split into pairs of groups?

Typically, we recommend users to run samples from all groups together, and then use the `contrast` argument of the `results` function to extract comparisons of interest after fitting the model using `DESeq`.

The model fit by `DESeq` estimates a single dispersion parameter for each gene, which defines how far we expect the observed count for a sample will be from the mean value from the model given its size factor and its condition group. See Section 4.1 and the *DESeq2* paper for full details. Having a single dispersion parameter for each gene is usually sufficient for analyzing multi-group data, as the final dispersion value will incorporate the within-group variability across all groups.

However, for some datasets, exploratory data analysis (EDA) plots as outlined in Section 2.2.3 could reveal that one or more groups has much higher within-group variability than the others. A simulated example of such a set of samples is shown in Figure 17. This is case where, by comparing groups A and B separately – subsetting a *DESeqDataSet* to only samples from those two groups and then running `DESeq` on this subset – will be more sensitive than a model including all samples together. It should be noted that such an extreme range of within-group variability is not common, although it could arise if certain treatments produce an extreme reaction (e.g. cell death). Again, this can be easily detected from the EDA plots such as PCA described in this vignette.




Figure 17: Extreme range of within-group variability. Typically, it is recommended to run `DESeq` across samples from all groups, for datasets with multiple groups. However, this simulated dataset shows a case where it would be preferable to compare groups A and B by creating a smaller dataset without the C samples. Group C has much higher within-group variability, which would inflate the per-gene dispersion estimate for groups A and B as well.

5.7 Can I run DESeq2 to contrast the levels of 100 groups?

DESeq2 will work with any kind of design specified using the R formula. We encourage users to consider exploratory data analysis such as principal components analysis as described in Section 2.2.3, rather than performing statistical testing of all combinations of dozens of groups.

As a speed concern with fitting very large models, note that each additional level of a factor in the design formula adds another parameter to the GLM which is fit by *DESeq2*. Users might consider first removing genes with very few reads, e.g. genes with row sum of 1, as this will speed up the fitting procedure.

5.8 Can I use DESeq2 to analyze a dataset without replicates?

If a *DESeqDataSet* is provided with an experimental design without replicates, a warning is printed, that the samples are treated as replicates for estimation of dispersion. This kind of analysis is only useful for exploring the data, but will not provide the kind of proper statistical inference on differences between groups. Without biological replicates, it is not possible to estimate the biological variability of each gene. More details can be found in the manual page for `?DESeq`.

5.9 How can I include a continuous covariate in the design formula?

Continuous covariates can be included in the design formula in the same manner as factorial covariates. Continuous covariates might make sense in certain experiments, where a constant fold change might be expected for each unit of the covariate. However, in many cases, more meaningful results can be obtained by cutting continuous covariates into a factor defined over a small number of bins (e.g. 3-5). In this way, the average effect of each group is controlled for, regardless of the trend over the continuous covariates. In R, *numeric* vectors can be converted into *factors* using the function `cut`.

5.10 Will the log fold change shrinkage “overshrink” large differences?

For most datasets, the application of a prior to the log fold changes is a good choice, providing log fold change estimates that are more stable across the entire range of mean counts than the maximum likelihood estimates (see Figure 1 and the *DESeq2* paper). One situation in which the prior on log fold changes might “overshrink” the estimates is if nearly all genes show no difference across condition, a very small set of genes have extremely large differences, and no genes in between. A simulated example of such a dataset is Figure 18. This is not likely to be the

Differential analysis of count data – the DESeq2 package

case for most experiments, where typically there is a range of differences by size: some genes with medium-to-large differences across treatment, and some with small differences.




Figure 18: Example of a dataset with where the log fold change prior should be turned off. Here we show a simulated MA-plot, where nearly all of the log fold changes are falling near the x-axis, with three genes that have very large log fold changes (note the y-axis is from -10 to 10 on the log₂ scale). This would indicate a dataset where the log fold change prior would “overshrink” the large fold changes, and so `betaPrior=FALSE` should be used.

There could be experiments in which only a few genes have very large log fold changes, and the rest of the genes are nearly constant across treatment. Or, there could be artificially constructed libraries fitting this description, e.g. technical replicates where the only difference across libraries is the concentration of a few spiked-in genes. “Overshrinking” of a few large log fold changes can be assessed by running `results` with `addMLE=TRUE`, which will print a results table with columns for the shrunken and unshrunken (MLE) log fold changes. The two estimates can be visually compared by running `plotMA` with `MLE=TRUE` and `MLE=FALSE`. If “overshrinking” very large log fold changes is a concern, it is better to turn off the log fold change prior by running `DESeq` with `betaPrior=FALSE`.

Even more detail: how do we avoid overshrinking on typical datasets? The answer is that we estimate the width of the log fold change prior in a robust way to accommodate the very largest log fold changes, and so to avoid overshrinking. The details of the prior estimation are described in the manual page for `nbinomWaldTest`. Briefly, a weighted upper quantile is used to match the width of the log fold change prior to the upper 5% of the MLE log fold changes, weighting by the expected sampling variability of the estimated log fold changes given the mean count for each gene. Note that this does not equal an assumption that 5% of genes are differentially expressed, but that a reasonable width of a log fold change distribution can be obtained from the upper 5% of MLE log fold changes.

5.11 I ran a likelihood ratio test, but `results()` only gives me one comparison.

“... How do I get the p values for all of the variables/levels that were removed in the reduced design?”

This is explained in the help page for `?results` in the section about likelihood ratio test p -values, but we will restate the answer here. When one performs a likelihood ratio test, the p values and the test statistic (the `stat` column) are values for the test that removes all of the variables which are present in the full design and not in the reduced design. This tests the null hypothesis that all the coefficients from these variables and levels of these factors are equal to zero.

The likelihood ratio test p values therefore represent a test of *all the variables and all the levels of factors* which are among these variables. However, the results table only has space for one column of log fold change, so a single variable and a single comparison is shown (among the potentially multiple log fold changes which were tested in the likelihood ratio test). This is indicated at the top of the results table with the text, e.g.: “log2 fold change (MLE): condition C vs A” followed by “LRT p -value: '~ batch + condition' vs '~ batch' ”. This indicates that the p value is for the likelihood ratio test of *all the variables and all the levels*, while the log fold change is a single comparison from among those variables and levels. See the help page for `results` for more details.

5.12 What are the exact steps performed by `DESeq()`?

See the manual page for `DESeq`, which links to the subfunctions which are called in order, where complete details are listed.

5.13 Is there an official Galaxy tool for DESeq2?

Yes. The repository for the `DESeq2` tool is <https://github.com/galaxyproject/tools-iuc/tree/master/tools/deseq2> and a link to its location in the Tool Shed is <https://toolshed.g2.bx.psu.edu/view/iuc/deseq2/d983d19fbbab>.

5.14 I want to benchmark DESeq2 comparing to other DE tools.

One aspect which can cause problems for comparison is that, by default, `DESeq2` outputs NA values for adjusted p values based on independent filtering of genes which have low counts. This is a way for the `DESeq2` to give extra information on why the adjusted p value for this gene is not small. Additionally, p values can be set to NA based on extreme count outlier detection (see Section 1.5.3 for full

Differential analysis of count data – the DESeq2 package

details). These `NA` values should be considered negatives for purposes of estimating sensitivity and specificity. The easiest way to work with the adjusted p values in a benchmarking context is probably to convert these `NA` values to 1:

```
res$padj <- ifelse(is.na(res$padj), 1, res$padj)
```

6 Acknowledgments

We have benefited in the development of *DESeq2* from the help and feedback of many individuals, including but not limited to: The Bioconductor Core Team, Alejandro Reyes, Andrzej Oleś, Aleksandra Pekowska, Felix Klein, Nikolaos Ignatiadis, Vince Carey, Owen Solberg, Ruping Sun, Devon Ryan, Steve Lianoglou, Jessica Larson, Christina Chaivorapol, Pan Du, Richard Bourgon, Willem Talloen, Elin Videvall, Hanneke van Deutekom, Todd Burwell, Jesse Rowley, Igor Dolgalev, Stephen Turner, Ryan C Thompson, Tyr Wiesner-Hanks, Konrad Rudolph, David Robinson, Mingxiang Teng, Mathias Lesche, Sonali Arora, Jordan Ramilowski, Ian Dworkin, Björn Grüning, Ryan McMinds, Paul Gordon, Leonardo Collado Torres, Enrico Ferrero.

7 Session Info

- R version 3.3.1 (2016-06-21), x86_64-pc-linux-gnu
- Locale: LC_CTYPE=en_US.UTF-8, LC_NUMERIC=C, LC_TIME=en_US.UTF-8, LC_COLLATE=C, LC_MONETARY=en_US.UTF-8, LC_MESSAGES=en_US.UTF-8, LC_PAPER=en_US.UTF-8, LC_NAME=C, LC_ADDRESS=C, LC_TELEPHONE=C, LC_MEASUREMENT=en_US.UTF-8, LC_IDENTIFICATION=C
- Base packages: base, datasets, grDevices, graphics, methods, parallel, stats, stats4, utils
- Other packages: Biobase 2.34.0, BiocGenerics 0.20.0, DESeq2 1.14.0, GenomInfoDb 1.10.0, GenomicRanges 1.26.0, IHW 1.2.0, IRanges 2.8.0, RColorBrewer 1.1-2, S4Vectors 0.12.0, SummarizedExperiment 1.4.0, airway 0.107.2, ggplot2 2.1.0, hexbin 1.27.1, knitr 1.14, pasilla 1.0.2, pheatmap 1.0.8, readr 1.0.0, tximport 1.2.0, tximportData 1.1.2, vsn 3.42.0
- Loaded via a namespace (and not attached): AnnotationDbi 1.36.0, BiocInstaller 1.24.0, BiocParallel 1.8.0, BiocStyle 2.2.0, DBI 0.5-1, Formula 1.2-1, Hmisc 3.17-4, Matrix 1.2-7.1, RCurl 1.95-4.8, RSQLite 1.0.0, Rcpp 0.12.7, XML 3.98-1.4, XVector 0.14.0, acepack 1.3-3.3, affy 1.52.0, affyio 1.44.0, annotate 1.52.0, assertthat 0.1, bitops 1.0-6, chron 2.3-47, cluster 2.0.5, codetools 0.2-15, colorspace 1.2-7, data.table 1.9.6, digest 0.6.10, evaluate 0.10, fdrtool 1.2.15, foreign 0.8-67, formatR 1.4, genefilter 1.56.0, geneplotter 1.52.0, grid 3.3.1, gridExtra 2.2.1, gtable 0.2.0, highr 0.6, labeling 0.3, lattice 0.20-34, latticeExtra 0.6-28, limma 3.30.0, locfit 1.5-9.1, lpsymphony 1.2.0,

Differential analysis of count data – the DESeq2 package

magrittr 1.5, munsell 0.4.3, nnet 7.3-12, plyr 1.8.4, preprocessCore 1.36.0, rpart 4.1-10, scales 0.4.0, slam 0.1-38, splines 3.3.1, stringi 1.1.2, stringr 1.1.0, survival 2.39-5, tibble 1.2, tools 3.3.1, xtable 1.8-2, zlibbioc 1.20.0

References

- [1] Michael I. Love, Wolfgang Huber, and Simon Anders. Moderated estimation of fold change and dispersion for RNA-seq data with DESeq2. *Genome Biology*, 15:550, 2014. URL: <http://dx.doi.org/10.1186/s13059-014-0550-8>.
- [2] Rob Patro, Stephen M. Mount, and Carl Kingsford. Sailfish enables alignment-free isoform quantification from RNA-seq reads using lightweight algorithms. *Nature Biotechnology*, 32:462–464, 2014. URL: <http://dx.doi.org/10.1038/nbt.2862>.
- [3] Rob Patro, Geet Duggal, and Carl Kingsford. Salmon: Accurate, versatile and ultrafast quantification from rna-seq data using lightweight-alignment. *bioRxiv*, 2015. URL: <http://biorxiv.org/content/early/2015/06/27/021592>.
- [4] Nicolas Bray, Harold Pimentel, Pall Melsted, and Lior Pachter. Near-optimal rna-seq quantification. *arXiv*, 2015. URL: <http://arxiv.org/abs/1505.02710>.
- [5] Bo Li and Colin N. Dewey. RSEM: accurate transcript quantification from RNA-Seq data with or without a reference genome. *BMC Bioinformatics*, 12:323+, 2011. URL: <http://dx.doi.org/10.1186/1471-2105-12-323>, doi:10.1186/1471-2105-12-3231.
- [6] Cole Trapnell, David G Hendrickson, Martin Sauvageau, Loyal Goff, John L Rinn, and Lior Pachter. Differential analysis of gene regulation at transcript resolution with RNA-seq. *Nature Biotechnology*, 2013. URL: <http://dx.doi.org/10.1038/nbt.2450>, doi:10.1038/nbt.2450.
- [7] Christelle Robert and Mick Watson. Errors in RNA-Seq quantification affect genes of relevance to human disease. *Genome Biology*, 2015. URL: <http://dx.doi.org/10.1186/s13059-015-0734-x>, doi:10.1186/s13059-015-0734-x.
- [8] Charlotte Sonesson, Michael I. Love, and Mark Robinson. Differential analyses for RNA-seq: transcript-level estimates improve gene-level inferences. *F1000Research*, 4, 2015. URL: <http://dx.doi.org/10.12688/f1000research.7563.1>.
- [9] Simon Anders, Paul Theodor Pyl, and Wolfgang Huber. HTSeq – A Python framework to work with high-throughput sequencing data. *Bioinformatics*, 2014. URL: <http://dx.doi.org/10.1093/bioinformatics/btu638>.
- [10] A. N. Brooks, L. Yang, M. O. Duff, K. D. Hansen, J. W. Park, S. Dudoit, S. E. Brenner, and B. R. Graveley. Conservation of an RNA regulatory map between Drosophila and mammals. *Genome Research*, pages 193–202, 2011. URL: <http://genome.cshlp.org/cgi/doi/10.1101/gr.108662.110>, doi:10.1101/gr.108662.110.

Differential analysis of count data – the DESeq2 package

- [11] Nikolaos Ignatiadis, Bernd Klaus, Judith Zaugg, and Wolfgang Huber. Data-driven hypothesis weighting increases detection power in big data analytics. *bioRxiv*, 2015. URL: <http://dx.doi.org/10.1101/034330>.
- [12] Robert Tibshirani. Estimating transformations for regression via additivity and variance stabilization. *Journal of the American Statistical Association*, 83:394–405, 1988.
- [13] Wolfgang Huber, Anja von Heydebreck, Holger Sültmann, Annemarie Poustka, and Martin Vingron. Parameter estimation for the calibration and variance stabilization of microarray data. *Statistical Applications in Genetics and Molecular Biology*, 2(1):Article 3, 2003.
- [14] Simon Anders and Wolfgang Huber. Differential expression analysis for sequence count data. *Genome Biology*, 11:R106, 2010. URL: <http://genomebiology.com/2010/11/10/R106>.
- [15] D. R. Cox and N. Reid. Parameter orthogonality and approximate conditional inference. *Journal of the Royal Statistical Society, Series B*, 49(1):1–39, 1987. URL: <http://www.jstor.org/stable/2345476>.
- [16] Davis J McCarthy, Yunshun Chen, and Gordon K Smyth. Differential expression analysis of multifactor RNA-Seq experiments with respect to biological variation. *Nucleic Acids Research*, 40:4288–4297, January 2012. URL: <http://www.ncbi.nlm.nih.gov/pubmed/22287627>, doi: 10.1093/nar/gks042.
- [17] Hao Wu, Chi Wang, and Zhijin Wu. A new shrinkage estimator for dispersion improves differential expression detection in RNA-seq data. *Biostatistics*, September 2012. URL: <http://dx.doi.org/10.1093/biostatistics/kxs033>, doi: 10.1093/biostatistics/kxs033.
- [18] R. Dennis Cook. Detection of Influential Observation in Linear Regression. *Technometrics*, February 1977.
- [19] Richard Bourgon, Robert Gentleman, and Wolfgang Huber. Independent filtering increases detection power for high-throughput experiments. *PNAS*, 107(21):9546–9551, 2010. URL: <http://www.pnas.org/content/107/21/9546.long>.