Differential analysis of RNA-Seq data at the gene level using the DESeq2 package

Michael Love, Simon Anders, Wolfgang Huber

Minor Modifications by Sean Davis

February 5, 2014

Contents

1 Introduction

In this lab, you will learn how to analyse a count table, such as arising from a summarised RNA-Seq experiment, for differentially expressed genes.

2 Input data

2.1 Experiment data

We read in a prepared *SummarizedExperiment*, which was generated from publicly available data from the article by Felix Haglund et al., "Evidence of a Functional Estrogen Receptor in Parathyroid Adenomas", J Clin Endocrin Metab, Sep 2012, http://www.ncbi.nlm.nih.gov/pubmed/23024189. Details on the generation of this object can be found in the vignette for the *parathyroidSE* package, http://bioconductor.org/packages/release/data/experiment/html/parathyroidSE.html.

The purpose of the experiment was to investigate the role of the estrogen receptor in parathyroid tumors. The investigators derived primary cultures of parathyroid adenoma cells from 4 patients. These primary cultures were treated with diarylpropionitrile (DPN), an estrogen receptor β agonist, or with 4-hydroxytamoxifen (OHT). RNA was extracted at 24 hours and 48 hours from cultures under treatment and control. The blocked design of the experiment allows for statistical analysis of the treatment effects while controlling for patient-to-patient variation.

We first load the *DESeq2* package and the data package *parathyroidSE*, which contains the example data set.

```
library("DESeq2")
library("parathyroidSE")
```

The `data` command loads a data object.

```
data("parathyroidGenesSE")
```

The information in a *SummarizedExperiment* object can be accessed with accessor functions. For example, to see the actual data, i.e., here, the read counts, we use the `assay` function. (The `head` function restricts the output to the first few lines.)

```
head(assay(parathyroidGenesSE))
```

```
##                 [,1] [,2] [,3] [,4] [,5] [,6] [,7] [,8] [,9] [,10] [,11]
## ENSG00000000003  792 1064  444  953  519  855  413  365  278  1173   463
## ENSG00000000005    4    1    2    3    3    1    0    1    0     0     0
## ENSG00000000419  294  282  164  263  179  217  277  204  189   601   257
```

```
## ENSG00000000457  156  184   93  145   75  122  228  171  116   422   182
## ENSG00000000460  396  207  210  212  221  173  611  199  426  1391   286
## ENSG00000000938    3    8    2    5    0    4   13   22    3    38    13
##                 [,12] [,13] [,14] [,15] [,16] [,17] [,18] [,19] [,20]
## ENSG00000000003   316   987   424   305   391   586   714   957   346
## ENSG00000000005     0     0     0     0     0     0     0     1     0
## ENSG00000000419   183   588   275   263   281   406   568   764   288
## ENSG00000000457   122   441   211   131   115   196   266   347   133
## ENSG00000000460   417  1452   238   188   102   389   294   778   162
## ENSG00000000938    10    26    13     7     3    10    18    15     7
##                 [,21] [,22] [,23] [,24] [,25] [,26] [,27]
## ENSG00000000003   433   402   277   511   366   271   492
## ENSG00000000005     0     0     0     0     0     0     0
## ENSG00000000419   259   250   147   271   227   197   363
## ENSG00000000457   168   148    83   184   136   118   195
## ENSG00000000460    85   339    75   154   314   117   233
## ENSG00000000938     8     7     5    13     8     7     8
```

In this count table, each row represents an Ensembl gene, each column a sequenced RNA library, and the values give the raw numbers of sequencing reads that were mapped to the respective gene in each library.

Question 1: For how many genes are there counts in this table?

Answer 1:

```
nrow(parathyroidGenesSE)
```

```
## [1] 63193
```

We also have metadata on each of the samples (the “columns” of the count table):

```
colData(parathyroidGenesSE)
```

```
## DataFrame with 27 rows and 8 columns
##             run experiment  patient treatment     time submission
##     <character>   <factor> <factor>  <factor> <factor>   <factor>
## 1     SRR479052  SRX140503        1   Control      24h  SRA051611
## 2     SRR479053  SRX140504        1   Control      48h  SRA051611
## 3     SRR479054  SRX140505        1       DPN      24h  SRA051611
## 4     SRR479055  SRX140506        1       DPN      48h  SRA051611
## 5     SRR479056  SRX140507        1       OHT      24h  SRA051611
## ...         ...        ...      ...       ...      ...        ...
## 23    SRR479074  SRX140523        4       DPN      48h  SRA051611
## 24    SRR479075  SRX140523        4       DPN      48h  SRA051611
## 25    SRR479076  SRX140524        4       OHT      24h  SRA051611
## 26    SRR479077  SRX140525        4       OHT      48h  SRA051611
## 27    SRR479078  SRX140525        4       OHT      48h  SRA051611
##         study    sample
##      <factor>  <factor>
## 1   SRP012167 SRS308865
```

```
## 2   SRP012167 SRS308866
## 3   SRP012167 SRS308867
## 4   SRP012167 SRS308868
## 5   SRP012167 SRS308869
## ...       ...       ...
## 23  SRP012167 SRS308885
## 24  SRP012167 SRS308885
## 25  SRP012167 SRS308886
## 26  SRP012167 SRS308887
## 27  SRP012167 SRS308887
```

Question 2: What are the metadata for the genes (the "rows" of the count table)?
Answer 2:

```
rowData(parathyroidGenesSE)
```

```
## GRangesList of length 63193:
## $ENSG00000000003
## GRanges with 17 ranges and 2 metadata columns:
##        seqnames              ranges strand   |   exon_id       exon_name
##           <Rle>           <IRanges>  <Rle>   | <integer>     <character>
##    [1]        X [99883667, 99884983]      -   |    664095 ENSE00001459322
##    [2]        X [99885756, 99885863]      -   |    664096 ENSE00000868868
##    [3]        X [99887482, 99887565]      -   |    664097 ENSE00000401072
##    [4]        X [99887538, 99887565]      -   |    664098 ENSE00001849132
##    [5]        X [99888402, 99888536]      -   |    664099 ENSE00003554016
##    ...      ...                  ...    ... ...       ...             ...
##   [13]        X [99890555, 99890743]      -   |    664106 ENSE00003512331
##   [14]        X [99891188, 99891686]      -   |    664108 ENSE00001886883
##   [15]        X [99891605, 99891803]      -   |    664109 ENSE00001855382
##   [16]        X [99891790, 99892101]      -   |    664110 ENSE00001863395
##   [17]        X [99894942, 99894988]      -   |    664111 ENSE00001828996
##
## ...
## <63192 more elements>
## ---
## seqlengths:
##                  1                  2 ...            LRG_99
##          249250621          243199373 ...             13294
```

2.2 Collapsing technical replicates

There are a number of samples which were sequenced in multiple runs. For example, sample `SRS308873` was sequenced twice. To see, we list the respective columns of the `colData`. (The use of `as.data.frame` forces R to show us the full list, not just the beginning and the end as before.)

```
as.data.frame(colData(parathyroidGenesSE)[, c("sample", "patient", "treatment",
    "time")])
```

```
##       sample patient treatment time
## 1  SRS308865       1   Control  24h
## 2  SRS308866       1   Control  48h
## 3  SRS308867       1       DPN  24h
## 4  SRS308868       1       DPN  48h
## 5  SRS308869       1       OHT  24h
## 6  SRS308870       1       OHT  48h
## 7  SRS308871       2   Control  24h
## 8  SRS308872       2   Control  48h
## 9  SRS308873       2       DPN  24h
## 10 SRS308873       2       DPN  24h
## 11 SRS308874       2       DPN  48h
## 12 SRS308875       2       OHT  24h
## 13 SRS308875       2       OHT  24h
## 14 SRS308876       2       OHT  48h
## 15 SRS308877       3   Control  24h
## 16 SRS308878       3   Control  48h
## 17 SRS308879       3       DPN  24h
## 18 SRS308880       3       DPN  48h
## 19 SRS308881       3       OHT  24h
## 20 SRS308882       3       OHT  48h
## 21 SRS308883       4   Control  48h
## 22 SRS308884       4       DPN  24h
## 23 SRS308885       4       DPN  48h
## 24 SRS308885       4       DPN  48h
## 25 SRS308886       4       OHT  24h
## 26 SRS308887       4       OHT  48h
## 27 SRS308887       4       OHT  48h
```

We recommend to first add together technical replicates (i.e., libraries derived from the same samples), such that we have one column per sample.

As is often the case, this preparatory step looks more complicated than the subsequent actual analysis. In fact, the following operations are not specific to *DESeq2*, but are specific preparations needed for this data set. To understand the general ideas of *DESeq2*, you could now skip to Section 3. What you will learn in the rest of this section is an example of a typical preparatory data manipulation task done with elementary R functions. Details on these can be found in general textbooks on R; also consider reading the help pages of the functions used.

We first use the function `split` to see which columns need to be collapsed.

```
allColSamples <- colData(parathyroidGenesSE)$sample
sp <- split(seq(along = allColSamples), allColSamples)
```

Using `sapply`, we loop over the elements of `sp`, which correspond to the distinct samples, construct subtables of the count table (i.e., `assay(parathyroidGenesSE)`) corresponding only to the current sample considered, and add up across rows if there is more than one column. The result of the `sapply` call is a new table, in which each column now corresponds to a different sample.

```
countdata <- sapply(sp, function(columns) rowSums(assay(parathyroidGenesSE)[,
    columns, drop = FALSE]))
head(countdata)
```

```
##                 SRS308865 SRS308866 SRS308867 SRS308868 SRS308869
## ENSG00000000003       792      1064       444       953       519
## ENSG00000000005         4         1         2         3         3
## ENSG00000000419       294       282       164       263       179
## ENSG00000000457       156       184        93       145        75
## ENSG00000000460       396       207       210       212       221
## ENSG00000000938         3         8         2         5         0
##                 SRS308870 SRS308871 SRS308872 SRS308873 SRS308874
## ENSG00000000003       855       413       365      1451       463
## ENSG00000000005         1         0         1         0         0
## ENSG00000000419       217       277       204       790       257
## ENSG00000000457       122       228       171       538       182
## ENSG00000000460       173       611       199      1817       286
## ENSG00000000938         4        13        22        41        13
##                 SRS308875 SRS308876 SRS308877 SRS308878 SRS308879
## ENSG00000000003      1303       424       305       391       586
## ENSG00000000005         0         0         0         0         0
## ENSG00000000419       771       275       263       281       406
## ENSG00000000457       563       211       131       115       196
## ENSG00000000460      1869       238       188       102       389
## ENSG00000000938        36        13         7         3        10
##                 SRS308880 SRS308881 SRS308882 SRS308883 SRS308884
## ENSG00000000003       714       957       346       433       402
## ENSG00000000005         0         1         0         0         0
## ENSG00000000419       568       764       288       259       250
## ENSG00000000457       266       347       133       168       148
## ENSG00000000460       294       778       162        85       339
## ENSG00000000938        18        15         7         8         7
##                 SRS308885 SRS308886 SRS308887
## ENSG00000000003       788       366       763
## ENSG00000000005         0         0         0
## ENSG00000000419       418       227       560
## ENSG00000000457       267       136       313
## ENSG00000000460       229       314       350
## ENSG00000000938        18         8        15
```

Novice users might find the preceding two code chunks difficult. Of course, there is a much easier way to add up the columns, namely by explicitly specifying the indices of the columns we want to use as is and the columns we want to add up, and using `cbind` to bind all the columns to a matrix:

```
a <- assay(parathyroidGenesSE)
countdata2 <- cbind(a[, 1:8], a[, 9] + a[, 10], a[, 11], a[, 12] + a[, 13],
    a[, 14:22], a[, 23] + a[, 24], a[, 25], a[, 26] + a[, 27])
all(countdata == countdata2)
```

```
## [1] TRUE
```

While this is simpler to understand, it is more error-prone. Mistakes can easily happen when determining the column indices, and it is tedious to update the code if the input data changes, for instance, if at a later time you would like to add more replicates to your data set. Hence, if you are a beginner in R and want to improve your R skills, try to understand how the `split` and the `sapply` calls above work, because only learning to master such expressions will give you the skills to make full use of R.

Having reduced our count data table to only one column per sample, we next need to subset the column metadata accordingly, as we now have less columns. We also now use the sample names as names for the column data rows:

```
coldata <- colData(parathyroidGenesSE)[sapply(sp, `[`, 1), ]
rownames(coldata) <- coldata$sample
coldata
```

```
## DataFrame with 23 rows and 8 columns
##                    run experiment  patient treatment     time submission
##            <character>   <factor> <factor>  <factor> <factor>   <factor>
## SRS308865   SRR479052  SRX140503        1   Control      24h  SRA051611
## SRS308866   SRR479053  SRX140504        1   Control      48h  SRA051611
## SRS308867   SRR479054  SRX140505        1       DPN      24h  SRA051611
## SRS308868   SRR479055  SRX140506        1       DPN      48h  SRA051611
## SRS308869   SRR479056  SRX140507        1       OHT      24h  SRA051611
## ...               ...        ...      ...       ...      ...        ...
## SRS308883   SRR479072  SRX140521        4   Control      48h  SRA051611
## SRS308884   SRR479073  SRX140522        4       DPN      24h  SRA051611
## SRS308885   SRR479074  SRX140523        4       DPN      48h  SRA051611
## SRS308886   SRR479076  SRX140524        4       OHT      24h  SRA051611
## SRS308887   SRR479077  SRX140525        4       OHT      48h  SRA051611
##                study    sample
##             <factor>  <factor>
## SRS308865 SRP012167 SRS308865
## SRS308866 SRP012167 SRS308866
## SRS308867 SRP012167 SRS308867
## SRS308868 SRP012167 SRS308868
## SRS308869 SRP012167 SRS308869
## ...              ...       ...
## SRS308883 SRP012167 SRS308883
## SRS308884 SRP012167 SRS308884
## SRS308885 SRP012167 SRS308885
## SRS308886 SRP012167 SRS308886
## SRS308887 SRP012167 SRS308887
```

Question 3: What do the quotation marks in the expression ‘`[`‘ do? What happens if you omit them?

Answer 3:

The function `sapply` expects an R function as its second argument. Here, we want to provide it with the function for vector subsetting (as in `a[1]`), and the name of this function is `[`. However, if we provide that name without the quotation marks, the R interpreter gets confused and complains about the unexpected symbol (try this out). Hence we need to quote the function name

in our call to `sapply`.

To unclutter the output in the subsequent steps, we only keep the column data columns that we actually need for our analysis.

```
coldata <- coldata[, c("patient", "treatment", "time")]
head(coldata)
```

```
## DataFrame with 6 rows and 3 columns
##            patient treatment     time
##           <factor>  <factor> <factor>
## SRS308865        1   Control      24h
## SRS308866        1   Control      48h
## SRS308867        1       DPN      24h
## SRS308868        1       DPN      48h
## SRS308869        1       OHT      24h
## SRS308870        1       OHT      48h
```

Our *SummarizedExperiment* object also contains metadata on the rows, which we can simply keep unchanged:

```
rowdata <- rowData(parathyroidGenesSE)
rowdata
```

```
## GRangesList of length 63193:
## $ENSG00000000003
## GRanges with 17 ranges and 2 metadata columns:
##        seqnames              ranges strand   |   exon_id       exon_name
##           <Rle>           <IRanges>  <Rle>   | <integer>     <character>
##    [1]        X [99883667, 99884983]      -   |    664095 ENSE00001459322
##    [2]        X [99885756, 99885863]      -   |    664096 ENSE00000868868
##    [3]        X [99887482, 99887565]      -   |    664097 ENSE00000401072
##    [4]        X [99887538, 99887565]      -   |    664098 ENSE00001849132
##    [5]        X [99888402, 99888536]      -   |    664099 ENSE00003554016
##    ...      ...                 ...    ... ...       ...             ...
##   [13]        X [99890555, 99890743]      -   |    664106 ENSE00003512331
##   [14]        X [99891188, 99891686]      -   |    664108 ENSE00001886883
##   [15]        X [99891605, 99891803]      -   |    664109 ENSE00001855382
##   [16]        X [99891790, 99892101]      -   |    664110 ENSE00001863395
##   [17]        X [99894942, 99894988]      -   |    664111 ENSE00001828996
##
## ...
## <63192 more elements>
## ---
## seqlengths:
##                  1                  2 ...            LRG_99
##          249250621          243199373 ...             13294
```

We now have all the ingredients to prepare our data object in a form that is suitable for analysis, namely:

- `countdata`: a table with the read counts, with technical replicates summed up,

- `coldata`: a table with metadata on the count table's columns, i.e., on the samples,
- `rowdata`: a table with metadata on the count table's rows, i.e., on the genes, and
- a design formula, which tells which factors in the column metadata table specify the experimental design and how these factors should be used in the analysis. We specify `~ patient + treatment`, which means that we want to test for the effect of treatment (the last factor), controlling for the effect of patient (the first factor). You can use R's formula notation to express any experimental design that can be described within an ANOVA-like framework.

To now construct the data object from the matrix of counts and the metadata table, we use:

```
ddsFull <- DESeqDataSetFromMatrix(countData = countdata, colData = coldata,
    design = ~patient + treatment, rowData = rowdata)


## Usage note:  the following factors have 3 or more levels:
##
## patient, treatment
##
## For DESeq2 versions < 1.3, if you plan on extracting results for
## these factors, we recommend using betaPrior=FALSE as an argument
## when calling DESeq().
## As currently implemented in version 1.2, the log2 fold changes can
## vary if the base level is changed, when extracting results for a
## factor with 3 or more levels.  A solution will be implemented in
## version 1.3 which allows for the use of a beta prior and symmetric
## log2 fold change estimates regardless of the selection of base level.

ddsFull


## class: DESeqDataSet
## dim: 63193 23
## exptData(0):
## assays(1): counts
## rownames(63193): ENSG00000000003 ENSG00000000005 ... LRG_98 LRG_99
## rowData metadata column names(0):
## colnames(23): SRS308865 SRS308866 ... SRS308886 SRS308887
## colData names(3): patient treatment time
```

3 Running the DESeq2 pipeline

Here we will analyze a subset of the samples, namely those taken after 48 hours, with either control or DPN treatment, taking into account the multifactor design.

3.1 Preparing the data object for the analysis of interest

First we subset the relevant columns from the full dataset:

```
dds <- ddsFull[, colData(ddsFull)$treatment %in% c("Control", "DPN") & colData(ddsFull)$time ==
    "48h"]
```

Sometimes it is necessary to "refactor" the factors, in case that levels have been dropped. (Here, for example, the treatment factor still contains the level "OHT", but no sample to this level.)

```
dds$patient <- factor(dds$patient)
dds$treatment <- factor(dds$treatment)
```

It will be convenient to make sure that `Control` is the *first* level in the treatment factor, so that the $\log_2$ fold changes are calculated as treatment over control. The function `relevel` achieves this:

```
dds$treatment <- relevel(dds$treatment, "Control")
```

A quick check whether we now have the right samples:

```
colData(dds)
```

```
## DataFrame with 8 rows and 3 columns
##            patient treatment     time
##           <factor>  <factor> <factor>
## SRS308866        1   Control      48h
## SRS308868        1       DPN      48h
## SRS308872        2   Control      48h
## SRS308874        2       DPN      48h
## SRS308878        3   Control      48h
## SRS308880        3       DPN      48h
## SRS308883        4   Control      48h
## SRS308885        4       DPN      48h
```

3.2 Running the pipeline

With the data object prepared, the *DESeq2* analysis can now be run with a single call to the function `DESeq`:

```
dds <- DESeq(dds)
```

```
## estimating size factors
## estimating dispersions
```

```
## gene-wise dispersion estimates
## mean-dispersion relationship
## final dispersion estimates
## 1 rows did not converge in dispersion, labelled in mcols(object)$dispConv.  Use larger maxit argument
with estimateDispersions
## fitting model and testing
```

3.3 Inspecting the results table

The results for the last variable in the design formula, in our case the `treatment` variable, can be extracted using the `results` function.

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

```
## DataFrame with 63193 rows and 6 columns
##                  baseMean log2FoldChange     lfcSE      stat    pvalue
##                 <numeric>      <numeric> <numeric> <numeric> <numeric>
## ENSG00000000003    623.42       -0.02045   0.05850   -0.3496  0.726644
## ENSG00000000005      0.68        0.02101   0.20529    0.1023  0.918485
## ENSG00000000419    299.75       -0.01483   0.07417   -0.2000  0.841482
## ENSG00000000457    183.50       -0.08919   0.10654   -0.8371  0.402521
## ENSG00000000460    200.46        0.35388   0.11086    3.1922  0.001412
## ...                   ...            ...       ...       ...       ...
## LRG_94                  0             NA        NA        NA        NA
## LRG_96                  0             NA        NA        NA        NA
## LRG_97                  0             NA        NA        NA        NA
## LRG_98                  0             NA        NA        NA        NA
## LRG_99                  0             NA        NA        NA        NA
##                     padj
##                <numeric>
## ENSG00000000003   0.90568
## ENSG00000000005        NA
## ENSG00000000419   0.94831
## ENSG00000000457   0.74551
## ENSG00000000460   0.02827
## ...                   ...
## LRG_94                 NA
## LRG_96                 NA
## LRG_97                 NA
## LRG_98                 NA
## LRG_99                 NA
```

As `res` is a `DataFrame` object, it carries metadata with information on the meaning of the columns:

```
mcols(res)
```

```
## DataFrame with 6 rows and 2 columns
```

```
##           type                                        description
##    <character>                                        <character>
## 1 intermediate                          the base mean over all rows
## 2      results log2 fold change (MAP): treatment DPN vs Control
## 3      results         standard error: treatment DPN vs Control
## 4      results         Wald statistic: treatment DPN vs Control
## 5      results      Wald test p-value: treatment DPN vs Control
## 6      results                                BH adjusted p-values
```

The first column, `baseMean`, is a just the average of the normalized count values, taken over all samples. The remaining four columns refer to a specific *contrast*, namely the comparison of the levels *DPN* versus *Control* of the factor variable *treatment*. See the help page for `results` (by typing `?results`) for information on how to obtain other contrasts.

The column `log2FoldChange` is the effect size estimate. It tells us how much the gene's expression seems to have changed due to treatment with DPN in comparison to control. This value is reported on a logarithmic scale to base 2: for example, a $\log_2$ fold change of 1.5 means that the gene's expression is increased by a factor of $2^{1.5} \approx 2.82$.

Of course, this estimate has an uncertainty associated with it, which is available in the column `lfcSE`, the standard error estimate for the log2 fold change estimate. We can also express the uncertainty of a particular effect size estimate as the result of a statistical test. The purpose of a test for differential expression is to test whether the data provides sufficient evidence to conclude that this value is really different from zero (and that the sign is correct). DESeq2 performs for each gene a *hypothesis test* to see whether evidence is sufficient to decide against the *null hypothesis* that there is no effect of the treatment on the gene and that the observed difference between treatment and control was merely caused by experimental variability (i. e., the type of variability that you can just as well expect between different samples in the same treatment group). As usual in statistics, the result of this test is reported as a *p value*, and it is found in the column `pvalue`. (Remember that a p value indicates the probability that a fold change as strong as the observed one, or even stronger, would be seen under the situation described by the null hypothesis.)

Finally, we note that a subset of the p values in `res` are `NA` ("not available"). This is `DESeq`'s way of reporting that all counts for this gene were zero, and hence not test was applied.

Question 4: How could you check to see if the `baseMean` is the mean of raw counts or the mean of normalized counts?

Answer 4: The raw counts and normalized counts of a *DESeqDataSet* object are available via the accessor function `counts`, which has an argument `normalized`, which defaults to `FALSE`.

```
all.equal(res$baseMean, rowMeans(counts(dds)))
```

```
## [1] "names for current but not for target"
## [2] "Mean relative difference: 0.05815"
```

```
all.equal(res$baseMean, rowMeans(counts(dds, normalized = TRUE)))
```

```
## [1] "names for current but not for target"
```

3.4 Multiple testing

Novices in high-throughput biology often assume that thresholding these p values at 0.05, as is often done in other settings, would be appropriate – but it is not. We briefly explain why:

There are 2115 genes with a p value below 0.05 among the 30434 genes, for which the test succeeded in reporting a p value:

```
sum(res$pvalue < 0.05, na.rm = TRUE)
```

```
## [1] 2115
```

```
table(is.na(res$pvalue))
```

```
##
## FALSE  TRUE
## 30434 32759
```

Now, assume for a moment that the null hypothesis is true for all genes, i.e., no gene is affected by the treatment with DPN. Then, by the definition of *p value*, we expect up to 5% of the genes to have a p value below 0.05. This amounts to 1522 genes. If we just considered the list of genes with a p value below 0.05 as differentially expressed, this list should therefore be expected to contain up to $1522/2115 = 72\%$ false positives!

DESeq2 uses the so-called Benjamini-Hochberg (BH) adjustment; in brief, this method calculates for each gene an *adjusted p value* which answers the following question: if one called significant all genes with a p value less than or equal to this gene's p value threshold, what would be the fraction of false positives (the *false discovery rate*, FDR) among them (in the sense of the calculation outlined above)? These values, called the BH-adjusted p values, are given in the column `padj` of the `results` object.

Hence, if we consider a fraction of 10% false positives acceptable, we can consider all genes with an *adjusted* p value below 10%=0.1 as significant. How many such genes are there?

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

```
## [1] 882
```

We subset the results table to these genes and then sort it by the log2-fold-change estimate to get the significant genes with the strongest down-regulation

```
resSig <- res[which(res$padj < 0.1), ]
head(resSig[order(resSig$log2FoldChange), ])
```

```
## DataFrame with 6 rows and 6 columns
##                  baseMean log2FoldChange     lfcSE      stat    pvalue
##                 <numeric>      <numeric> <numeric> <numeric> <numeric>
```

```
## ENSG00000163631     268.8        -0.9417   0.10335    -9.112 8.086e-20
## ENSG00000169239    1547.6        -0.7531   0.10645    -7.075 1.497e-12
## ENSG00000041982    1493.3        -0.6864   0.09224    -7.441 9.964e-14
## ENSG00000145244     173.3        -0.6764   0.25564    -2.646 8.146e-03
## ENSG00000119946     183.5        -0.6745   0.10754    -6.272 3.565e-10
## ENSG00000155111     587.9        -0.6647   0.08574    -7.753 8.964e-15
##                      padj
##                 <numeric>
## ENSG00000163631 6.388e-17
## ENSG00000169239 5.914e-10
## ENSG00000041982 4.971e-11
## ENSG00000145244 9.213e-02
## ENSG00000119946 7.679e-08
## ENSG00000155111 4.998e-12
```

and with the strongest upregulation

```
tail(resSig[order(resSig$log2FoldChange), ])
```

```
## DataFrame with 6 rows and 6 columns
##                  baseMean log2FoldChange     lfcSE      stat    pvalue
##                 <numeric>      <numeric> <numeric> <numeric> <numeric>
## ENSG00000005189     226.8         0.6411    0.2293     2.796 5.176e-03
## ENSG00000156414     136.9         0.7274    0.1382     5.265 1.400e-07
## ENSG00000103257     168.2         0.7802    0.1513     5.157 2.514e-07
## ENSG00000101255     285.0         0.7970    0.1856     4.293 1.759e-05
## ENSG00000135069     136.1         0.8427    0.3210     2.625 8.661e-03
## ENSG00000092621     594.2         0.8686    0.1508     5.762 8.327e-09
##                      padj
##                 <numeric>
## ENSG00000005189 6.881e-02
## ENSG00000156414 1.528e-05
## ENSG00000103257 2.408e-05
## ENSG00000101255 9.984e-04
## ENSG00000135069 9.603e-02
## ENSG00000092621 1.316e-06
```

Question 5: What is the proportion of down- and up-regulation among the genes with adjusted p value less than 0.1?

Answer 5:

```
table(sign(resSig$log2FoldChange))
```

```
##
##  -1   1
## 459 423
```

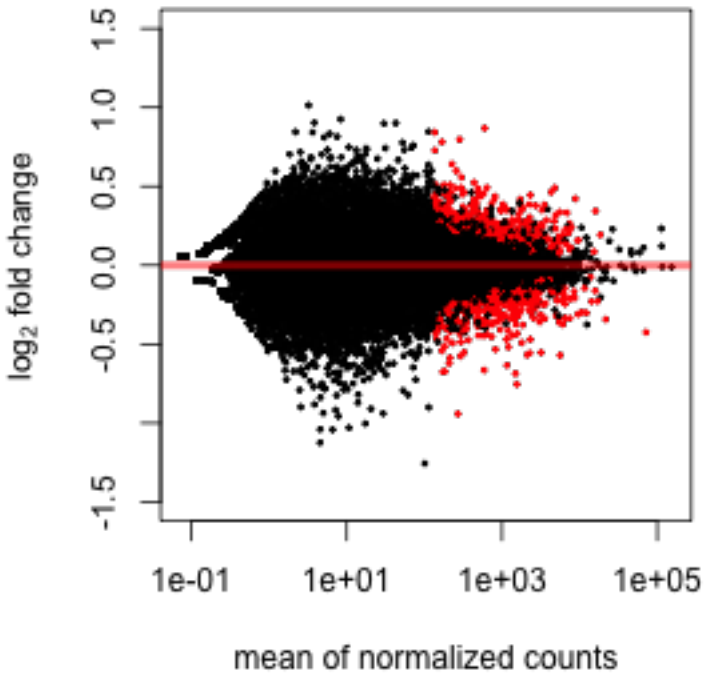

Figure 1: The MA-plot shows the $\log_2$ fold changes from the treatment over the mean of normalized counts, i.e. the average of counts normalized by size factor. The *DESeq2* package incorporates a prior on $\log_2$ fold changes, resulting in moderated estimates from genes with low counts and highly variable counts, as can be seen by the narrowing of spread of points on the left side of the plot.

3.5 Diagnostic plots

A so-called MA plot provides a useful overview for an experiment with a two-group comparison:

```
plotMA(dds, ylim = c(-1.5, 1.5))
```

The plot (Fig. 1) represents each gene with a dot. The x axis is the average expression over all samples, the y axis the $\log_2$ fold change between treatment and control. Genes with an adjusted p value below a threshold (here 0.1, the default) are shown in red.

This plot demonstrates that only genes with an average normalized count above 10 contain sufficient information to yield a significant call, and only above about 300 counts can smaller fold-changes become significant.

Also note *DESeq2*'s shrinkage estimation of log fold changes (LFCs): When count values are too low to allow an accurate estimate of the LFC, the value is "shrunken" towards zero to avoid

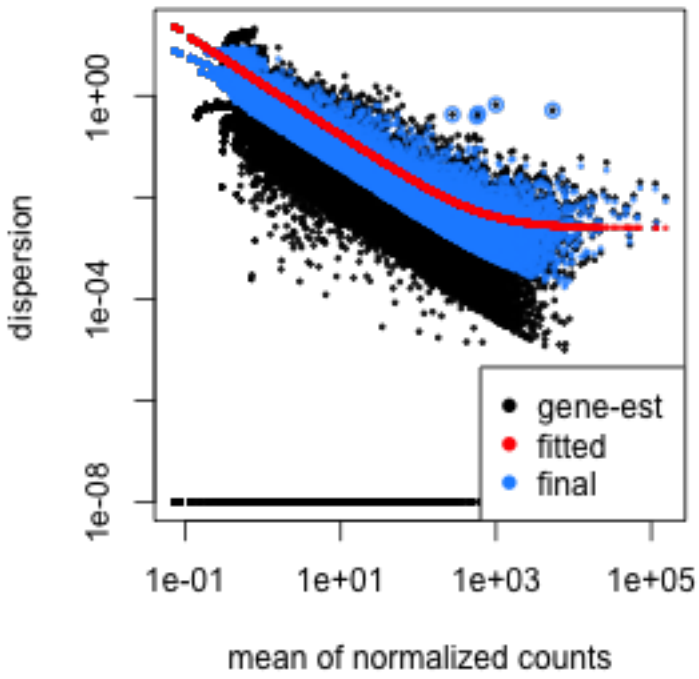

Figure 2: Plot of dispersion estimates. See text for details.

that these values, which otherwise would frequently be unrealistically large, dominate the top-ranked log fold changes.

Whether a gene is called significant depends not only on its LFC but also on its within-group variability, which *DESeq2* quantifies as the *dispersion*. For strongly expressed genes, the dispersion can be understood as a squared coefficient of variation: a dispersion value of 0.01 means that the gene's expression tends to differ by typically $\sqrt{0.01} = 10\%$ between samples of the same treatment group. For weak genes, the Poisson noise is an additional source of noise, which is added to the dispersion.

The function `plotDispEsts` visualizes *DESeq2*'s dispersion estimates:

```
plotDispEsts(dds)
```

The black dots are the dispersion estimates for each gene as obtained by considering the information from each gene separately. Unless one has many samples, these values fluctuate strongly around their true values. Therefore, we fit the red trend line, which shows the dispersions' dependence on the mean, and then shrink each gene's estimate towards the red line to obtain the final estimates (blue circles) that are then used in the hypothesis test.

Question 6: How could you change the MA-plot so as to color those genes with adjusted

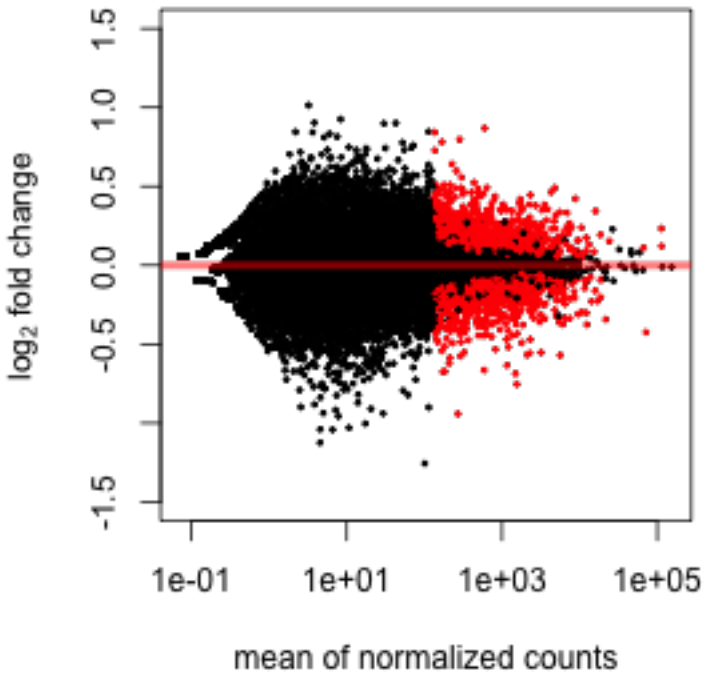

Figure 3: The MA-plot with red points indicating adjusted p value less than 0.5.

p-value less than 0.5 instead of 0.1?

Answer 6:

```
plotMA(dds, pvalCutoff = 0.5, ylim = c(-1.5, 1.5))
```

See Figure 3.

Another useful diagnostic plot is the histogram of the p values (Fig. 4).

```
hist(res$pvalue, breaks = 100)
```

Question 7: Revisit the discussion about p values and multiple testing in the previous section. Which part of the histogram is caused by genes that are called significant? And which part is caused by those that are truly significant? Why are there "spikes" at intermediate values?

Answer 7: Genes that are not differentially expressed have p values that are approximately uniformly distributed between 0 and 1. This gives rise to the floor of bars of equal heights. The truly differentially expressed genes give rise to the tall bar(s) at the very left – but only to that part of the bars that raises above the uniform floor. Of course, we cannot know which of the genes in these tall bars are true ones and which are not. When only looking at the bars to the

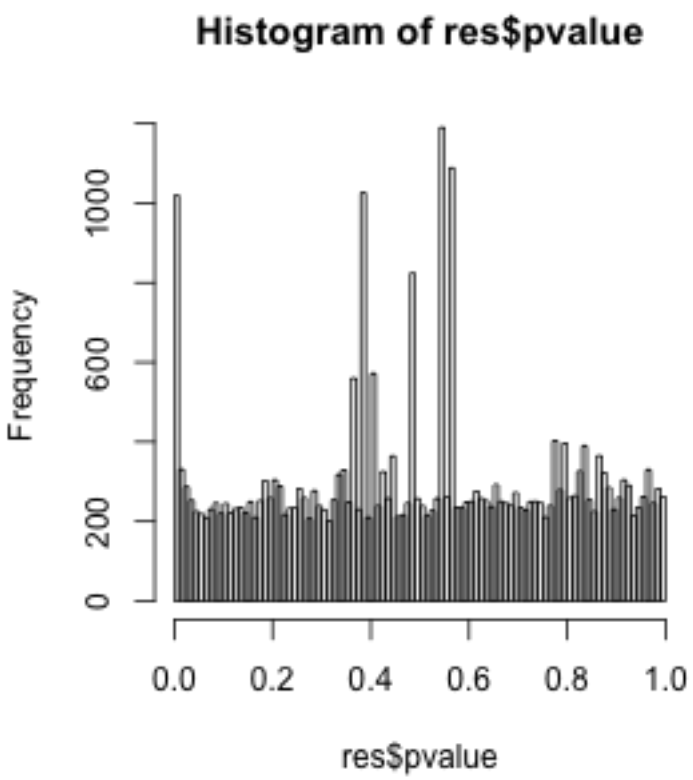

Figure 4: Histogram of the p values returned by the test for differential expression.

left of our chosen p value cut-off, the ratio of “floor” area to total area provides an estimate of the false discovery rate. This is a graphical way of understanding FDR.

The rule that p values from null cases are uniform is true only for continuous test statistics. However, for genes with low counts, the fact that we are working with integer counts becomes noticeable, and gives rise to the spikes at intermediate p values.

4 Independent filtering

The MA plot (Figure 1) highlights an important property of RNA-Seq data. For weakly expressed genes, we have no chance of seeing differential expression, because the low read counts suffer from so high Poisson noise that any biological effect is drowned in the uncertainties from the read counting. The MA plot suggests that for genes with less than one or two counts per sample, averaged over all samples, there is no real inferential power. We loose little if we filter out these genes:

```
filterThreshold <- 2
keep <- rowMeans(counts(dds, normalized = TRUE)) > filterThreshold
```

```
table(keep)
```

```
## keep
## FALSE  TRUE
## 43939 19254
```

Note that none of the genes below the threshold had a significant adjusted p value

```
min(res$padj[!keep], na.rm = TRUE)
```

```
## Warning:  no non-missing arguments to min; returning Inf
```

```
## [1] Inf
```

At first sight, there may seem to be little benefit in filtering out these genes. After all, the test found them to be non-significant anyway. However, these genes have an influence on the multiple testing adjustment, whose performance improves if such genes are removed. Compare:

```
table(p.adjust(res$pvalue, method = "BH") < 0.1)
```

```
##
## FALSE  TRUE
## 29863   571
```

```
table(p.adjust(res$pvalue[keep], method = "BH") < 0.1)
```

```
##
## FALSE  TRUE
## 18545   709
```

By removing the weakly-expressed genes from the input to the FDR procedure, we have found more genes to be significant among those which we kept, and so improved the power of our test. This approach is known as *independent filtering*.

The term *independent* highlights an important caveat. Such filtering is permissible only if the filter criterion is independent of the actual test statistic [1]. Otherwise, the filtering would invalidate the test and consequently the assumptions of the BH procedure. This is why we filtered on the average over *all* samples: this filter is blind to the assignment of samples to the treatment and control group and hence independent.

Question 8: Redo the histogram as in Figure 4, now only using the genes that passed the filtering. What happened to the spikes at intermediate values?

Answer 8: Run

```
hist(res$pvalue[keep], breaks = 100)
```

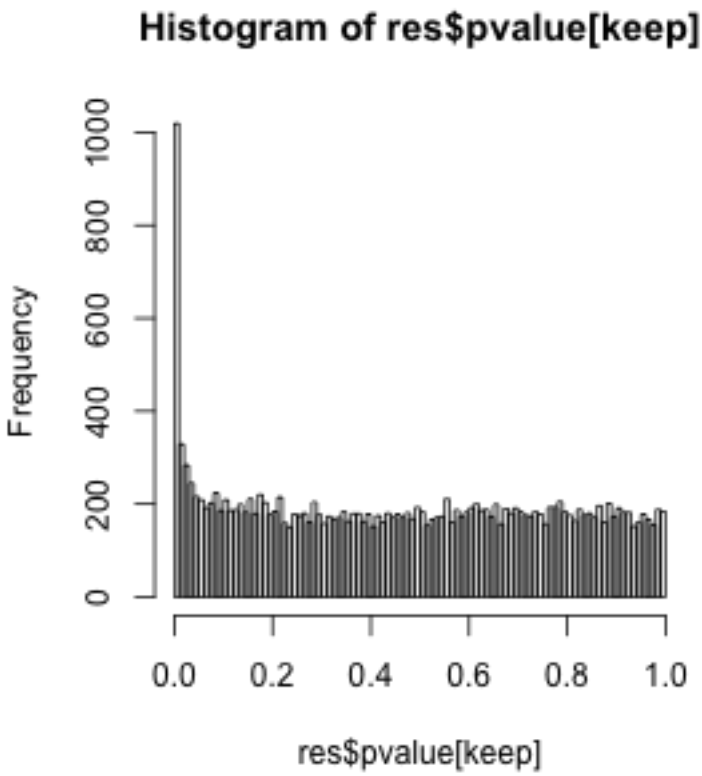

Figure 5: Histogram of the p values returned by the test for differential expression.

See Figure 5. As explained before, the spikes were caused by genes with low counts. Having removed these, our p value histogram now looks smoother.

In this vignette, we have determined the value for `filterThreshold`, 2, by looking at Figure 1. More formal, automatable ways exist; if you are interested, please have a look at the vignette *Diagnostics for independent filtering* in the *genefilter* package.

4.1 Adding gene names

Our result table only uses Ensembl gene IDs, but gene names may be more informative. Bioconductor's annotation packages help with mapping various ID schemes to each other.

We load the annotation package *org.Hs.eg.db*:

```
library("org.Hs.eg.db")
```

This is the organism annotation package ("org") for *Homo sapiens* ("Hs"), organized as an *AnnotationDbi* package ("db"), using Entrez Gene IDs ("eg") as primary key.

To get a list of all available key types, use

```
cols(org.Hs.eg.db)
```

```
## Warning:  'cols' has been deprecated and replaced by 'columns' for versions
##  of Bioc that are higher than 2.13.  Please use 'columns'
##  anywhere that you previously used 'cols'
```

```
##  [1] "ENTREZID"     "PFAM"         "IPI"          "PROSITE"
##  [5] "ACCNUM"       "ALIAS"        "CHR"          "CHRLOC"
##  [9] "CHRLOCEND"    "ENZYME"       "MAP"          "PATH"
## [13] "PMID"         "REFSEQ"       "SYMBOL"       "UNIGENE"
## [17] "ENSEMBL"      "ENSEMBLPROT"  "ENSEMBLTRANS" "GENENAME"
## [21] "UNIPROT"      "GO"           "EVIDENCE"     "ONTOLOGY"
## [25] "GOALL"        "EVIDENCEALL"  "ONTOLOGYALL"  "OMIM"
## [29] "UCSCKG"
```

Converting IDs with the native functions from the *AnnotationDbi* package is currently a bit cumbersome, so we provide the following convenience function (without explaining how exactly it works):

```
convertIDs <- function(ids, fromKey, toKey, db, ifMultiple = c("putNA", "useFirst")) {
    stopifnot(inherits(db, "AnnotationDb"))
    ifMultiple <- match.arg(ifMultiple)
    suppressWarnings(selRes <- AnnotationDbi::select(db, keys = ids, keytype = fromKey,
        cols = c(fromKey, toKey)))
    if (ifMultiple == "putNA") {
        duplicatedIds <- selRes[duplicated(selRes[, 1]), 1]
        selRes <- selRes[!selRes[, 1] %in% duplicatedIds, ]
    }
    return(selRes[match(ids, selRes[, 1]), 2])
}
```

This function takes a list of IDs as first argument and their key type as the second argument. The third argument is the key type we want to convert to, the fourth is the *AnnotationDb* object to use. Finally, the last argument specifies what to do if one source ID maps to several target IDs: should the function return an NA or simply the first of the multiple IDs?

To convert the Ensembl IDs in the rownames of `res` to gene symbols and add them as a new column, we use:

```
res$symbol <- convertIDs(row.names(res), "ENSEMBL", "SYMBOL", org.Hs.eg.db)
res
```

```
## DataFrame with 63193 rows and 7 columns
##                   baseMean log2FoldChange     lfcSE      stat    pvalue
##                  <numeric>      <numeric> <numeric> <numeric> <numeric>
## ENSG00000000003    623.42       -0.02045   0.05850   -0.3496  0.726644
## ENSG00000000005      0.68        0.02101   0.20529    0.1023  0.918485
## ENSG00000000419    299.75       -0.01483   0.07417   -0.2000  0.841482
## ENSG00000000457    183.50       -0.08919   0.10654   -0.8371  0.402521
## ENSG00000000460    200.46        0.35388   0.11086    3.1922  0.001412
```

```
## ...                    ...             ...        ...        ...        ...
## LRG_94                   0              NA         NA         NA         NA
## LRG_96                   0              NA         NA         NA         NA
## LRG_97                   0              NA         NA         NA         NA
## LRG_98                   0              NA         NA         NA         NA
## LRG_99                   0              NA         NA         NA         NA
##                       padj      symbol
##                  <numeric> <character>
## ENSG00000000003   0.90568      TSPAN6
## ENSG00000000005        NA        TNMD
## ENSG00000000419   0.94831        DPM1
## ENSG00000000457   0.74551       SCYL3
## ENSG00000000460   0.02827    C1orf112
## ...                   ...         ...
## LRG_94                 NA          NA
## LRG_96                 NA          NA
## LRG_97                 NA          NA
## LRG_98                 NA          NA
## LRG_99                 NA          NA
```

Finally, we note that you can easily save the results table in a CSV file, which you can then load with a spreadsheet program such as Excel:

```
write.csv(as.data.frame(res), file = "results.csv")
```

5 Downstream analyses

A list of gene names is no final result. We demonstrate two possible further analysis steps.

5.1 Gene set enrichment analysis

Do the genes with a strong up- or down-regulation have something in common? We perform next a gene-set enrichment analysis (GSEA) to examine this question.

We use the gene sets in the Reactome database

```
library("reactome.db")
```

This database works with Entrez IDs, so we add a column with such IDs, using our `convertIDs` function:

```
res$entrez <- convertIDs(row.names(res), "ENSEMBL", "ENTREZID", org.Hs.eg.db)
```

Next, we subset the results table, `res`, to only those genes for which the Reactome database has data (i.e, whose Entrez ID we find in the respective key column of `reactome.db`) and for which the test gave a p value that was not NA.

```
res2 <- res[res$entrez %in% keys(reactome.db, "ENTREZID") & !is.na(res$pvalue),
    ]
head(res2)
```

```
## DataFrame with 6 rows and 8 columns
##                  baseMean log2FoldChange     lfcSE      stat    pvalue
##                 <numeric>      <numeric> <numeric> <numeric> <numeric>
## ENSG00000000419    299.75      -0.014833   0.07417  -0.20000    0.8415
## ENSG00000000938     12.14      -0.004828   0.32050  -0.01506    0.9880
## ENSG00000000971     19.70      -0.385546   0.26481  -1.45594    0.1454
## ENSG00000001084    323.54       0.045042   0.09453   0.47646    0.6337
## ENSG00000001167    412.38      -0.053865   0.10309  -0.52253    0.6013
## ENSG00000001626     10.99       0.152666   0.29882   0.51090    0.6094
##                      padj      symbol      entrez
##                 <numeric> <character> <character>
## ENSG00000000419    0.9483        DPM1        8813
## ENSG00000000938        NA         FGR        2268
## ENSG00000000971        NA         CFH        3075
## ENSG00000001084    0.8702        GCLC        2729
## ENSG00000001167    0.8590        NFYA        4800
## ENSG00000001626        NA        CFTR        1080
```

Using `select`, a function from *AnnotationDbi* for querying database objects, we get a table with the mapping from Entrez IDs to Reactome Path IDs

```
reactomeTable <- AnnotationDbi::select(reactome.db, keys = res2$entrez, keytype = "ENTREZID",
    columns = c("ENTREZID", "REACTOMEID"))
```

```
## Warning:  'select' and duplicate query keys resulted in 1:many mapping
## between keys and return rows
```

```
head(reactomeTable)
```

```
##   ENTREZID REACTOMEID
## 1     8813     162699
## 2     8813     163125
## 3     8813     392499
## 4     8813     446193
## 5     8813     446203
## 6     8813     446219
```

The next code chunk transforms this table into an *incidence matrix*. This is a boolean matrix with one row for each Reactome Path and one column for each gene in `res2`, which tells us which genes are members of which Reactome Paths. (If you want to understand how this chunk exactly works, read up about the `tapply` function.)

```
incm <- do.call(rbind, with(reactomeTable, tapply(ENTREZID, factor(REACTOMEID),
    function(x) res2$entrez %in% x)))
colnames(incm) <- res2$entrez
str(incm)
```

```
##  logi [1:1435, 1:5494] FALSE FALSE FALSE FALSE FALSE FALSE ...
##  - attr(*, "dimnames")=List of 2
##   ..$ : chr [1:1435] "1059683" "109581" "109582" "109606" ...
##   ..$ : chr [1:5494] "8813" "2268" "3075" "2729" ...
```

We remove all rows corresponding to Reactome Paths with less than 5 assigned genes.

```
incm <- incm[rowSums(incm) >= 5, ]
```

To test whether the genes in a Reactome Path behave in a special way in our experiment, we perform t-tests to see whether the average of the genes' $\log_2$ fold change values are different from zero. If so, we can say that our treatment tends to upregulate (or downregulate) the genes in the category. To facilitate the computations, we define a little helper function:

```
testCategory <- function(reactomeID) {
    isMember <- incm[reactomeID, ]
    data.frame(reactomeID = reactomeID, numGenes = sum(isMember), avgLFC = mean(res2$log2FoldChange[isMemb
        strength = sum(res2$log2FoldChange[isMember])/sqrt(sum(isMember)), pvalue = t.test(res2$log2FoldCh
        reactomeName = reactomePATHID2NAME[[reactomeID]])
}
```

The function can be called with a Reactome Path ID:

```
testCategory("109581")
```

```
##   reactomeID numGenes    avgLFC strength pvalue            reactomeName
## 1     109581      148 -0.008654  -0.1053 0.3563 Homo sapiens: Apoptosis
```

As you can see the function not only performs the t test and returns the p value but also lists other useful information such as the number of genes in the category, the average log fold change, a "strength" measure (see below) and the name with which Reactome describes the Path.

We call the function for all Paths in our incidence matrix and collect the results in a data frame:

```
reactomeResult <- do.call(rbind, lapply(rownames(incm), testCategory))
```

As we performed many tests, we should again use a multiple testing adjustment.

```
reactomeResult$padjust <- p.adjust(reactomeResult$pvalue, "BH")
```

This is a list of Reactome Paths which are significantly differentially expressed in our comparison of DPN treatment with control, sorted according to sign and strength of the signal:

```
reactomeResultSignif <- reactomeResult[reactomeResult$padjust < 0.05, ]
reactomeResultSignif[order(reactomeResultSignif$strength), ]
```

```
##      reactomeID numGenes   avgLFC strength    pvalue
## 878      445355       20 -0.14415  -0.6447 4.545e-04
## 1137     913531      140 -0.04617  -0.5462 2.049e-04
## 104     1280215      246 -0.03135  -0.4918 3.845e-04
## 69      1169408       64 -0.05716  -0.4573 6.472e-04
## 70      1169410       64 -0.05716  -0.4573 6.472e-04
## 552     2032785       10 -0.11754  -0.3717 2.341e-04
## 1066      74159      137 -0.03069  -0.3592 2.851e-04
## 1044      73857       96 -0.03649  -0.3575 2.558e-05
## 651     2468052       11 -0.10114  -0.3354 2.869e-04
## 1079      75071      124 -0.03010  -0.3352 3.226e-05
## 1013      72163      101 -0.03236  -0.3252 2.331e-05
## 1015      72172      101 -0.03236  -0.3252 2.331e-05
## 1018      72203      105 -0.03165  -0.3243 2.723e-05
## 439      189445       14  0.09715   0.3635 2.264e-04
## 467      191273       22  0.10121   0.4747 1.418e-04
## 928      597592      176  0.03666   0.4863 7.299e-04
## 169      156590       20  0.11118   0.4972 6.374e-04
## 539     2024096       20  0.11150   0.4986 1.376e-04
## 165      156580       49  0.07301   0.5111 5.042e-04
## 505      196849       53  0.07370   0.5365 6.849e-04
## 506      196854       53  0.07370   0.5365 6.849e-04
## 731      381070       46  0.08294   0.5625 2.505e-04
## 130     1430728     1319  0.02010   0.7300 2.110e-05
##                                                    reactomeName
## 878                          Homo sapiens: Smooth Muscle Contraction
## 1137                              Homo sapiens: Interferon Signaling
## 104                 Homo sapiens: Cytokine Signaling in Immune system
## 69                            Homo sapiens: ISG15 antiviral mechanism
## 70           Homo sapiens: Antiviral mechanism by IFN-stimulated genes
## 552    Homo sapiens: YAP1- and WWTR1 (TAZ)-stimulated gene expression
## 1066                                    Homo sapiens: Transcription
## 1044                    Homo sapiens: RNA Polymerase II Transcription
## 651           Homo sapiens: Establishment of Sister Chromatid Cohesion
## 1079                                  Homo sapiens: mRNA Processing
## 1013                      Homo sapiens: mRNA Splicing - Major Pathway
## 1015                                    Homo sapiens: mRNA Splicing
## 1018    Homo sapiens: Processing of Capped Intron-Containing Pre-mRNA
## 439                             Homo sapiens: Metabolism of porphyrins
## 467                             Homo sapiens: Cholesterol biosynthesis
## 928               Homo sapiens: Post-translational protein modification
## 169                               Homo sapiens: Glutathione conjugation
## 539                                    Homo sapiens: HS-GAG degradation
## 165                                 Homo sapiens: Phase II conjugation
## 505  Homo sapiens: Metabolism of water-soluble vitamins and cofactors
## 506                 Homo sapiens: Metabolism of vitamins and cofactors
## 731               Homo sapiens: Activation of Chaperones by IRE1alpha
## 130                                            Homo sapiens: Metabolism
##       padjust
## 878  0.033857
```

```
## 1137 0.024432
## 104  0.030555
## 69   0.037107
## 70   0.037107
## 552  0.024432
## 1066 0.024432
## 1044 0.006409
## 651  0.024432
## 1079 0.006409
## 1013 0.006409
## 1015 0.006409
## 1018 0.006409
## 439  0.024432
## 467  0.021128
## 928  0.037826
## 169  0.037107
## 539  0.021128
## 165  0.035356
## 505  0.037107
## 506  0.037107
## 731  0.024432
## 130  0.006409
```

Note that such lists need to be interpreted with care, and a grain of salt. Which of these categories make sense, given the biology of the experiment?

6 Working with rlog-transformed data

6.1 The rlog transform

Many common statistical methods for exploratory analysis of multidimensional data, especially methods for clustering and ordination (e. g., principal-component analysis and the like), work best for (at least approximately) homoskedastic data; this means that the variance of an observable (i.e., here, the expression strength of a gene) does not depend on the mean. In RNA-Seq data, however, variance grows with the mean. For example, if one performs PCA directly on a matrix of normalized read counts, the result typically depends only on the few most strongly expressed genes because they show the largest absolute differences between samples. A simple and often used strategy to avoid this is to take the logarithm of the normalized count values; however, now the genes with low counts tend to dominate the results because, due to the strong Poisson noise inherent to small count values, they show the strongest relative differences between samples.

As a solution, *DESeq2* offers the *regularized-logarithm transformation*, or *rlog* for short. For genes with high counts, the rlog transformation differs not much from an ordinary $\log_2$ transformation. For genes with lower counts, however, the values are shrunken towards the genes' averages across all samples. Using an empirical Bayesian prior in the form of a *ridge penality*, this is done such that the rlog-transformed data are approximately homoskedastic.

The function `rlogTransform` returns a *SummarizedExperiment* object which contains the rlog-transformed values in its *assay* slot:

```
rld <- rlogTransformation(dds)


## you had estimated gene-wise dispersions, removing these
## you had estimated fitted dispersions, removing these


head(assay(rld))


##                 SRS308866 SRS308868 SRS308872 SRS308874 SRS308878
## ENSG00000000003    9.7165    9.6866    9.1335    9.1904    8.9599
## ENSG00000000005   -0.6789   -0.4751   -0.6294   -0.7818   -0.7834
## ENSG00000000419    8.1004    8.1104    8.2398    8.2902    8.3018
## ENSG00000000457    7.4441    7.3024    7.8164    7.7088    7.2215
## ENSG00000000460    7.5731    7.6711    7.9871    8.1711    7.1350
## ENSG00000000938    3.2347    3.0754    4.1804    3.6720    2.9665
##                 SRS308880 SRS308883 SRS308885
## ENSG00000000003    8.8707    9.0940    9.1140
## ENSG00000000005   -0.7952   -0.7825   -0.7928
## ENSG00000000419    8.3071    8.2645    8.1716
## ENSG00000000457    7.3426    7.6011    7.4953
## ENSG00000000460    7.4533    7.0171    7.3745
## ENSG00000000938    3.4031    3.3676    3.4999
```

To show the effect of the transformation, we plot the first sample against the second, first simply using the `log2` function (after adding 1, to avoid taking the log of zero), and then using the rlog-transformed values.

```
par(mfrow = c(1, 2))
plot(log2(1 + counts(dds, normalized = TRUE)[, 1:2]), col = "#00000020", pch = 20,
    cex = 0.3)
plot(assay(rld)[, 1:2], col = "#00000020", pch = 20, cex = 0.3)
```

Note that, in order to make it easier to see where several points are plotted on top of each other, we set the plotting color to a semi-transparent black (encoded as `#00000020`) and changed the points to solid disks (`pch=20`) with reduced size (`cex=0.3`)[1].

In Figure 6, we can see how genes with low counts seem to be excessively variable on the ordinary logarithmic scale, while the rlog transform compresses differences for genes for which the data cannot provide good information anyway.

6.2 Sample distances

A useful first step in an RNA-Seq analysis is often to assess overall similarity between samples: Which samples are similar to each other, which are different? Does this fit to the expectation from the experiment's design?

[1] The function `heatscatter` from the package *LSD* offers a colourful alternative.

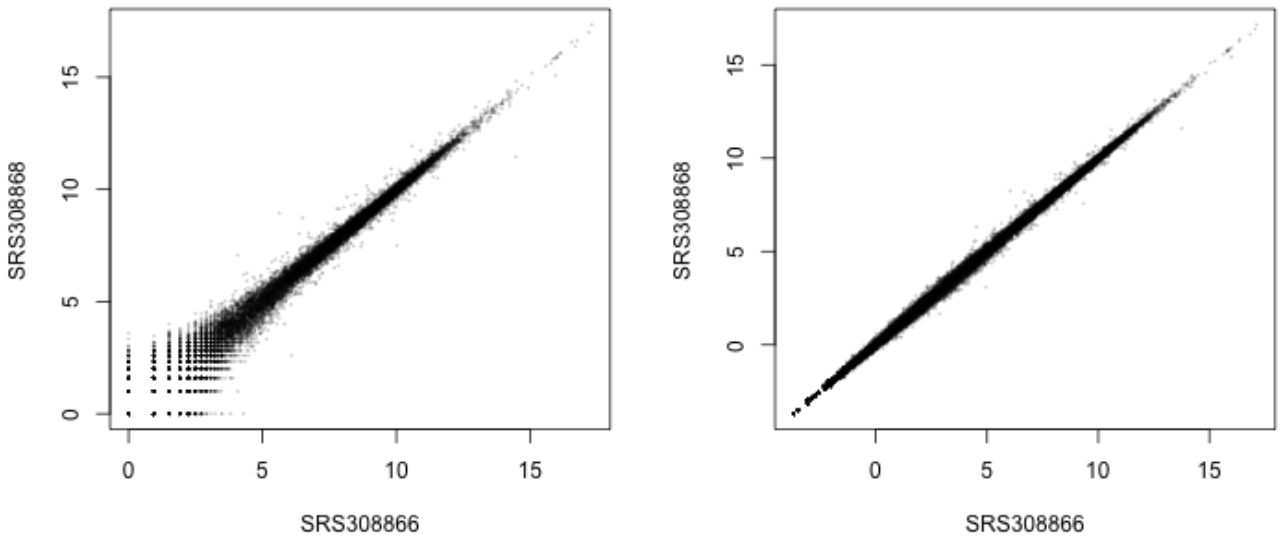

Figure 6: Scatter plot of sample 2 versus sample 1. Left: using an ordinary $\log_2$ transformation. Right: Using the rlog transformation.

We use the R function `dist` to calculate the Euclidean distance between samples. To avoid that the distance measure is dominated by a few highly variable genes, and have a roughly equal contribution from all genes, we use it on the rlog-transformed data:

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

```
##           SRS308866 SRS308868 SRS308872 SRS308874 SRS308878 SRS308880
## SRS308868     27.59
## SRS308872     70.37     72.71
## SRS308874     70.31     71.47     27.82
## SRS308878     91.81     94.38     91.83     93.86
## SRS308880     91.81     93.79     92.94     94.32     26.75
## SRS308883     79.93     82.72     76.56     78.43     69.42     69.88
## SRS308885     78.03     79.84     75.21     76.12     71.01     69.95
##           SRS308883
## SRS308868
## SRS308872
## SRS308874
## SRS308878
## SRS308880
## SRS308883
## SRS308885     27.40
```

Note the use of the function `t` to transpose the data matrix. We need this because `dist` calculates distances between data *rows* and our samples constitute the columns.

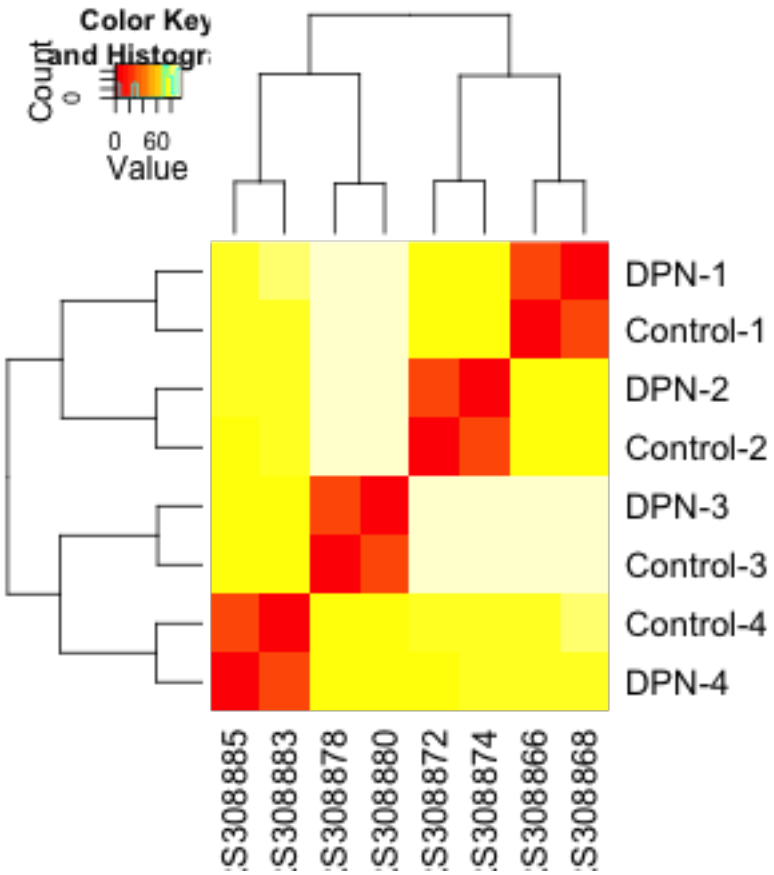

Figure 7: Heatmap of Euclidean sample distances after rlog transformation.

We visualize the distances in a heatmap, using the function `heatmap.2` from the *gplots* package.

```
sampleDistMatrix <- as.matrix(sampleDists)
rownames(sampleDistMatrix) <- paste(colData(rld)$treatment, colData(rld)$patient,
    sep = "-")
library("gplots")
heatmap.2(sampleDistMatrix, trace = "none")
```

Note that we have changed the row names of the distance matrix to contain treatment type and patient number instead of sample ID, so that we have all this information in view when looking at the heatmap (Fig. 7).

Question 9: Some people find the colour scheme used in Figure 7 ugly. Make a better version. *Hint:* Look at the sequential colour schemes in the *RColorBrewer* package and at the `colorRampPalette` function.

Answer 9:

```
library("RColorBrewer")
colours = colorRampPalette(rev(brewer.pal(9, "Blues")))(255)
heatmap.2(sampleDistMatrix, trace = "none", col = colours)
```

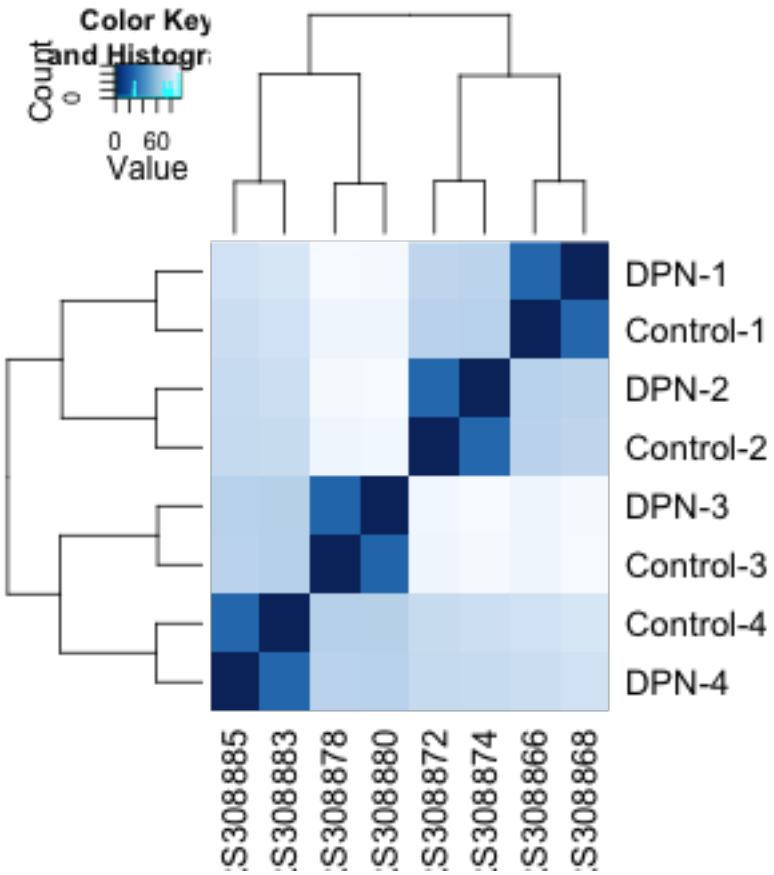

Figure 8: Using RColorBrewer to improve our heatmap

See Figure 8.

Another way to visualize sample-to-sample distances is a principal-components analysis (PCA). In this ordination method, the data points (i.e., here, the samples) are projected onto the 2D plane such that they spread out optimally (Fig. 9).

```
print(plotPCA(rld, intgroup = c("patient", "treatment")))
```

Here, we have used the function `plotPCA` which comes with *DESeq2*. The two terms specified as `intgroup` are column names from our sample data; they tell the function to use them to choose colours.

From both visualizations, we see that the differences between patients is much larger than the difference between treatment and control samples of the same patient. This shows why it was important to account for this paired design ("paired", because each treated sample is paired with one control sample from the *same* patient). We did so by using the design formula `!~ patient  treatment!` when setting up the data object in the beginning. Had we used an unpaired analysis, by specifying only `~ treatment`, we would not have found many hits, because then, the patient-to-patient differences would have drowned out any treatment effects.

Here, we have performed this sample distance analysis towards the end of our analysis. In

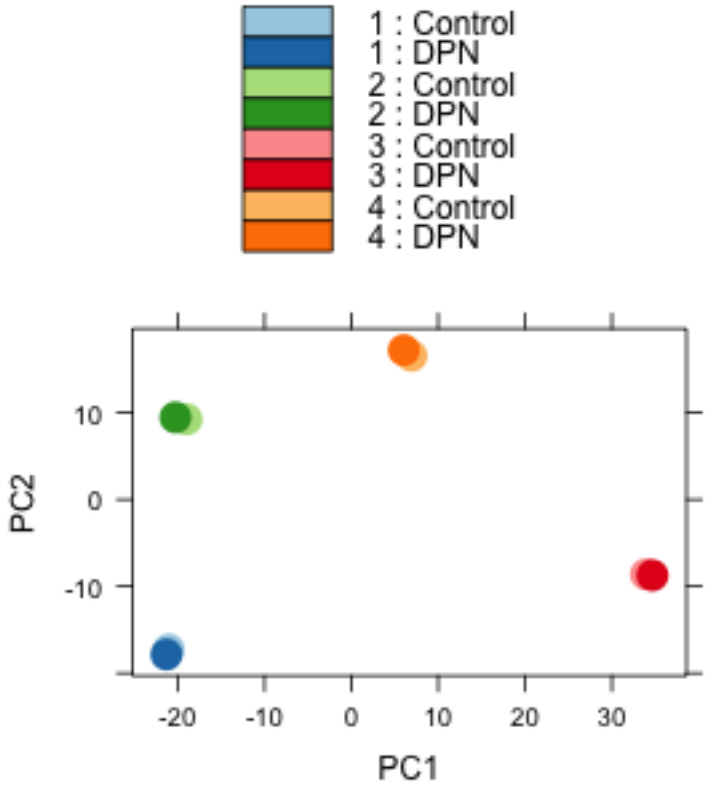

Figure 9: The Principal Components Analysis plot shows the sample relationships in a different way than the heatmap. In this plot, one can quickly notice that the difference between patients is much geater than the difference between treatments.

practice, however, this is a step suitable to give a first overview on the data. Hence, one will typically carry out this analysis as one of the first steps in an analysis. To this end, you may also find the function `arrayQualityMetrics`, from the equinymous package, useful.

6.3 Gene clustering

In the heatmap of Fig. 7, the dendrogram at the side shows us a hierarchical clustering of the samples. Such a clustering can also be performed for the genes.

Since the clustering is only relevant for genes that actually carry signal, one usually carries it out only for a subset of most highly variable genes. Here, for demonstration, let us select the 35 genes with the highest variance across samples:

```
library("genefilter")
topVarGenes <- head(order(rowVars(assay(rld)), decreasing = TRUE), 35)
```

The heatmap becomes more interesting if we do not look at absolute expression strength but

rather at the amount by which each gene deviates in a specific sample from the gene's average across all samples. Hence, we center and scale each genes' values across samples, and plot a heatmap.

```
heatmap.2(assay(rld)[topVarGenes, ], scale = "row", trace = "none", dendrogram = "column",
    col = colorRampPalette(rev(brewer.pal(9, "RdBu")))(255))
```

We can now see (Fig. 10) blocks of genes which covary across patients. Often, such a heatmap is insightful, even though here, seeing these variations across patients is of limited value because we are rather interested in the effects between the two samples from each patient.

7 Advanced Questions

For these questions, we provide (and probably have) no solutions, advanced readers are encouraged to explore them.

1. *DESeq2* performs the shrinkage of the dispersion estimates by fitting a parametric curve on the mean of normalized counts (cf. Figure 2). However, one could argue that the biological variability of genes should not be a function of counts, but of counts per gene length (i. e., expression level), and that regression on that covariate should lead to a better fit. Write your own version of the `estimateDispersions` function to explore this question.

2. What is the contribution of UTR length variations to the between-replicates variability modelled by *DESeq2*? The read counting script (available in the vignette of *parathyroidSE*) uses all exons of the genes, which includes UTRs. Would detection power be increased –or would we preferentially detect different phenomena– if we left out UTRs from the counting (i. e. count reads that fall on coding exons only); or indeed, if we looked only at UTRs?

References

[1] Richard Bourgon, Robert Gentleman, and Wolfgang Huber. Independent filtering increases detection power for high-throughput experiments. *PNAS*, 107(21):9546–9551, 2010.

8 Session Info

As last part of this document, we call the function `sessionInfo`, which reports the version numbers of R and all the packages used in this session. It is good practice to always keep such a record as it will help to trace down what has happened in case that an R script ceases to work because a package has been changed in a newer version.

```
## R version 3.0.2 Patched (2014-01-22 r64855)
## Platform: x86_64-apple-darwin10.8.0 (64-bit)
##
## locale:
## [1] C
##
## attached base packages:
## [1] grid      parallel  stats     graphics  grDevices utils     datasets
## [8] methods   base
##
## other attached packages:
##  [1] rtracklayer_1.21.12
##  [2] reactome.db_1.44.0
##  [3] randomForest_4.6-7
##  [4] pheatmap_0.7.7
##  [5] org.Hs.eg.db_2.10.1
##  [6] gplots_2.11.3
##  [7] gtools_3.1.0
##  [8] ggplot2_0.9.3.1
##  [9] genefilter_1.43.0
## [10] gdata_2.13.2
## [11] devtools_1.3
## [12] caTools_1.14
## [13] VariantAnnotation_1.7.47
## [14] TxDb.Hsapiens.UCSC.hg19.knownGene_2.10.1
## [15] ShortRead_1.19.13
## [16] SRAdb_1.15.0
## [17] graph_1.39.3
## [18] SNPlocs.Hsapiens.dbSNP.20120608_0.99.9
## [19] Rsamtools_1.13.46
## [20] RSQLite_0.11.4
## [21] RCurl_1.95-4.1
## [22] bitops_1.0-6
## [23] RColorBrewer_1.0-5
## [24] MASS_7.3-29
## [25] KernSmooth_2.23-10
## [26] Gviz_1.5.15
## [27] GenomicFeatures_1.13.43
## [28] GEOquery_2.28.0
## [29] DBI_0.2-7
## [30] BiocBrazil2014_1.0
## [31] parathyroidSE_1.0.4
## [32] DESeq2_1.2.9
## [33] RcppArmadillo_0.3.910.0
## [34] Rcpp_0.10.4
```

```
## [35] minfiData_0.4.2
## [36] IlluminaHumanMethylation450kanno.ilmn12.hg19_0.2.1
## [37] IlluminaHumanMethylation450kmanifest_0.4.0
## [38] minfi_1.8.9
## [39] bumphunter_1.1.17
## [40] locfit_1.5-9.1
## [41] iterators_1.0.6
## [42] foreach_1.4.1
## [43] reshape_0.8.4
## [44] plyr_1.8
## [45] lattice_0.20-24
## [46] BSgenome.Hsapiens.UCSC.hg19_1.3.19
## [47] BSgenome_1.29.1
## [48] Biostrings_2.29.19
## [49] GenomicRanges_1.13.45
## [50] XVector_0.1.4
## [51] IRanges_1.19.38
## [52] AnnotationDbi_1.24.0
## [53] ALL_1.4.14
## [54] Biobase_2.21.7
## [55] BiocGenerics_0.7.5
## [56] knitr_1.5
## [57] BiocInstaller_1.12.0
##
## loaded via a namespace (and not attached):
##  [1] Hmisc_3.12-2          R.methodsS3_1.5.1     XML_3.95-0.2
##  [4] annotate_1.39.0       base64_1.1            beanplot_1.1
##  [7] biomaRt_2.17.3        biovizBase_1.9.4      cluster_1.14.4
## [10] codetools_0.2-8       colorspace_1.2-3      compiler_3.0.2
## [13] dichromat_2.0-0       digest_0.6.3          doRNG_1.5.5
## [16] evaluate_0.5          formatR_0.9           gtable_0.1.2
## [19] highr_0.2.1           httr_0.2              hwriter_1.3
## [22] illuminaio_0.3.11     itertools_0.1-1       labeling_0.2
## [25] latticeExtra_0.6-26   limma_3.17.25         matrixStats_0.8.12
## [28] mclust_4.2            memoise_0.1           multtest_2.17.0
## [31] munsell_0.4.2         nlme_3.1-113          nor1mix_1.1-4
## [34] pkgmaker_0.17.4       preprocessCore_1.23.0 proto_0.3-10
## [37] registry_0.2          reshape2_1.2.2        rngtools_1.2.3
## [40] rpart_4.1-4           scales_0.2.3          siggenes_1.35.0
## [43] splines_3.0.2         stats4_3.0.2          stringr_0.6.2
## [46] survival_2.37-7       tools_3.0.2           whisker_0.3-2
## [49] xtable_1.7-1          zlibbioc_1.7.0
```

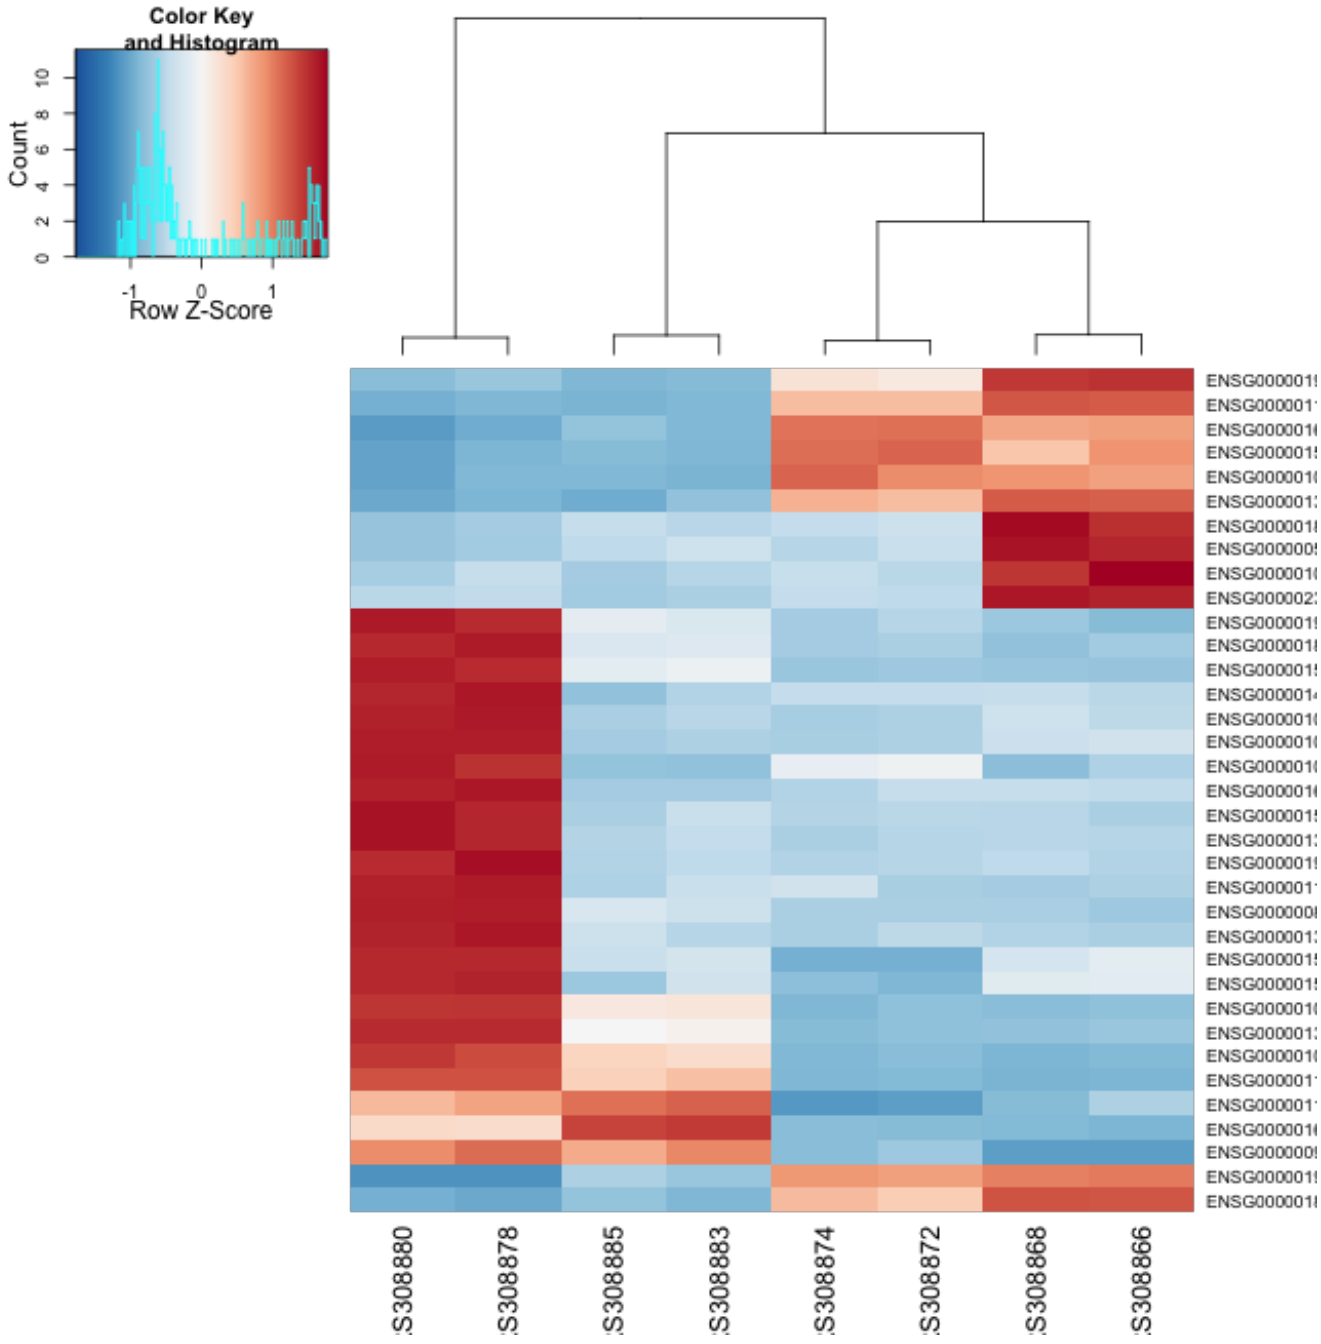

Figure 10: Heatmap with gene clustering.