# CSAMA 2014: RNA-Seq differential expression workflow

Michael Love<sup>1\*</sup>, Simon Anders<sup>2</sup>, Wolfgang Huber<sup>2</sup>

Department of Biostatistics, Dana Farber Cancer Institute and Harvard School of Public Health, Boston, US;
European Molecular Biology Laboratory (EMBL), Heidelberg, Germany
\*michaelisaiahlove (at) gmail.com

### June 30, 2014

#### **Abstract**

This lab will walk you through an end-to-end RNA-Seq differential expression workflow. We will start from the FASTQ files, align to the reference genome, prepare gene expression values as a count table by counting the sequenced fragments, perform differential gene expression analysis, and visually explore the results.

This lab covers the basic introduction. For a more in-depth explanation of the advanced details, we advise you to proceed to the vignette of the *DESeq2* package package, *Differential analysis of count data – the DESeq2 package*. For a treatment of exon-level differential expression, we refer to the vignette of the *DEXSeq* package, *Analyzing RNA-seq data for differential exon usage with the DEXSeq package*.

### Contents

1	Inpu	t data	1
	1.1	Preparing count matrices	2
		Aligning reads to a reference	
	1.3	Example BAM files	3
	1.4	Counting reads in genes	4
	1.5	The DESeqDataSet, column metadata, and the design formula	7
	1.6	Collapsing technical replicates	9
2	Run	ning the DESeq2 pipeline	11
	2.1	Preparing the data object for the analysis of interest	11
	2.2	Running the pipeline	12
	2.3	Inspecting the results table	13
	2.4	Other comparisons	14
	2.5	Adding gene names	15

3	Further points	18
	3.1 Multiple testing	18
	3.2 Diagnostic plots	19
	3.3 Independent filtering	22
	3.4 Exporting results	24
	3.5 Gene-set enrichment analysis	25
4	Working with rlog-transformed data	31
4	Working with rlog-transformed data 4.1 The rlog transform	
4		31
4	4.1 The rlog transform	31 33

# 1 Input data

As example data for this lab, we will use 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<sup>1</sup>.

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  $\beta$  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.

Part of the data from this experiment is provided in the Bioconductor data package *parathyroidSE*.

# 1.1 Preparing count matrices

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 have been mapped 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).

The count values must be raw counts of sequencing reads. This is important for *DESeq2*'s statistical model to hold, as only the actual counts allow assessing the measurement precision correctly. Hence, please do not supply other quantities, such as (rounded) normalized counts, or counts of covered base pairs – this will only lead to nonsensical results.

<sup>1</sup>http://www.ncbi.nlm.nih.gov/pubmed/23024189

### 1.2 Aligning reads to a reference

The computational analysis of an RNA-Seq experiment begins earlier: what we get from the sequencing machine is a set of FASTQ files that contain the nucleotide sequence of each read and a quality score at each position. These reads must first be aligned to a reference genome or transcriptome. It is important to know if the sequencing experiment was single-end or paired-end, as the alignment software will require the user to specify both FASTQ files for a paired-end experiment. The output of this alignment step is commonly stored in a file format called BAM.

A number of software programs exist to align reads to the reference genome, and the development is too rapid for this document to provide a current list. We recommend consulting benchmarking papers that discuss the advantages and disadvantages of each software, which include accuracy, ability to align reads over splice junctions, speed, memory footprint, and many other features.

Here we use the TopHat2 spliced alignment software<sup>2</sup> [?] in combination with the Bowtie index available at the Illumina iGenomes page<sup>3</sup>. For full details on this software and on the iGenomes, users should follow the links to the manual and information provided in the links in the footnotes. For example, the paired-end RNA-Seq reads for the *parathyroidSE* package were aligned using TopHat2 with 8 threads, with the call:

```
tophat2 -o file_tophat_out -p 8 genome file_1.fastq file_2.fastq
samtools sort -n file_tophat_out/accepted_hits.bam _sorted
```

The second line sorts the reads by name rather than by genomic position, which is necessary for counting paired-end reads within Bioconductor. This command uses the SAMtools software<sup>4</sup> [?]. The BAM files for a number of sequencing runs can then be used to generate count matrices, as described in the following section.

# 1.3 Example BAM files

Besides tha main count table, which we will use later, the *parathyroidSE* package also contains a small subset of the raw data from the Haglund et al. experiment, namely three BAM file each with a subset of the reads from three of the samples. We will use these files to demonstrate how a count table can be constructed from BAM files. Afterwards, we will load the full count table corresponding to all samples and all data, which is already provided in the same package, and will continue the analysis with that full table.

We load the data package with the example data

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

The R function system.file can be used to find out where on your computer the files from a package have been installed. Here we ask for the full path to the extdata directory, which is part of the

<sup>2</sup>http://tophat.cbcb.umd.edu/

<sup>3</sup>http://tophat.cbcb.umd.edu/igenomes.html

<sup>4</sup>http://samtools.sourceforge.net

#### parathyroidSE package:

```
extDataDir <- system.file("extdata", package = "parathyroidSE", mustWork = TRUE)
extDataDir
## [1] "/home/oles/R/R-3.1.0/library/parathyroidSE/extdata"</pre>
```

In this directory, we find the three BAM files (and some other files):

Typically, we have a table with experimental meta data for our samples. For these three files, it is as follows:

sampleName		fileName	treatment	time
	Ctrl_24h_1	SRR479052.bam	Control	24h
	$Ctrl_48h_1$	SRR479053.bam	Control	48h
	DPN_24h_1	SRR479054.bam	DPN	24h

To avoid mistakes, it is helpful to store such a sample table explicitly in a text file and load it.

#### Exercise 1

Construct a table with the content shown above on your disk using a spread sheet application such as Microsoft Excel and save the sheet in CSV (comma-separated values) format (or simply use a plain text editor).

Load it with:

```
sampleTable <- read.csv( "/path/to/your/sampleTable.csv", header=TRUE )</pre>
```

This is how the sample table should look like

```
sampleTable

## sampleName fileName treatment time
## 1 Ctrl_24h_1 SRR479052.bam Control 24h
## 2 Ctrl_48h_1 SRR479053.bam Control 48h
## 3 DPN_24h_1 SRR479054.bam DPN 24h
```

Using the fileName column in the table, we construct the full paths to the files we want to perform the counting operation on:

```
bamFiles <- file.path( extDataDir, sampleTable$fileName )
bamFiles

## [1] "/home/oles/R/R-3.1.0/library/parathyroidSE/extdata/SRR479052.bam"

## [2] "/home/oles/R/R-3.1.0/library/parathyroidSE/extdata/SRR479053.bam"

## [3] "/home/oles/R/R-3.1.0/library/parathyroidSE/extdata/SRR479054.bam"</pre>
```

## 1.4 Counting reads in genes

To count how many read map to each gene, we need transcript annotation. Download the current GTF file with human gene annotation from Ensembl. (In case the network is too slow for that, use the truncated version of this file, called Homo\_sapiens.GRCh37.75.subset.gtf.gz, which we have placed on the course server.)

From this file, the function makeTranscriptDbFromGFF from the *GenomicFeatures* constructs a database of all annotated transcripts.

```
library( "GenomicFeatures" )
hse <- makeTranscriptDbFromGFF( "/path/to/your/genemodel_file.GTF", format="gtf" )
exonsByGene <- exonsBy( hse, by="gene" )

## Warning: Infering Exon Rankings. If this is not what you expected, then please
be sure that you have provided a valid attribute for exonRankAttributeName
## Warning: None of the strings in your circ_seqs argument match your seqnames.</pre>
```

In the last step, we have used the exonsBy function to bring the transcriptome data base into the shape of a list of all genes,

```
exonsByGene
## GRangesList of length 100:
## $ENSG0000000003
## GRanges with 13 ranges and 2 metadata columns:
##
         seqnames
                                ranges strand
                                                    exon_id
                                                              exon_name
            <Rle>
##
                             <IRanges> <Rle>
                                                | <integer> <character>
##
      [1]
                X [99883667, 99884983]
                                                       3038
                                                                   <NA>
      [2]
                X [99885756, 99885863]
                                                                   < NA >
##
                                                       3039
                X [99887482, 99887565]
##
      [3]
                                                       3040
                                                                   <NA>
      [4]
                X [99887538, 99887565]
##
                                                       3041
                                                                   < NA >
##
     [5]
                X [99888402, 99888536]
                                                       3042
                                                                   < NA >
##
      . . .
                                                                    . . .
                X [99890555, 99890743]
##
     [9]
                                                       3046
                                                                   <NA>
     [10]
                X [99891188, 99891686]
##
                                                       3047
                                                                   < NA >
     [11]
                X [99891605, 99891803]
##
                                                       3048
                                                                   < NA >
                X [99891790, 99892101]
##
     [12]
                                                                   < NA >
                                                       3049
##
     [13]
                X [99894942, 99894988]
                                                       3050
                                                                   <NA>
##
## ...
## <99 more elements>
## ---
## seqlengths:
    7 12 2 6 16
                  4
                      3 1 17 8 19 X 11
```

#### Exercise 2

Note the warnings issued by makeTranscriptDbFromGFF. Can we safely ignore them?

#### Exercise 3

In exonsByGene, inspect the genomic intervals given for the exons of the first gene. Note that they are not disjunct (they overlap). How comes? Will this influence results in the following step?

After these preparations, the actual counting is easy. The function summerizeOverlap from the *GenomicAlignments* package will do this.

We use the counting mode "Union", which indicates that those reads which overlap any portion of exactly one feature are counted. For more information on the different counting modes, see the help page for summarizeOverlaps. As this experiment produced paired-end reads, we specify singleEnd = FALSE. As protocol was not strand-specific, we specify ignore.strand = TRUE. fragments = TRUE indicates that we also want to count reads with unmapped pairs. This last argument is only for use with paired-end experiments.

Remember that we have only used a small subset of reads from the original experiment: for 3 samples and for 100 genes. Nevertheless, we can still investigate the resulting *SummarizedExperiment* by looking at the counts in the assay slot, the phenotypic data about the samples in colData slot (in this case an empty *DataFrame*), and the data about the genes in the rowData slot. Figure 1 explains the basic structure of the *SummarizedExperiment* class.

```
## class: SummarizedExperiment
## dim: 100 3
## exptData(0):
## assays(1): counts
## rownames(100): ENSG0000000003 ENSG00000000005 ... ENSG00000005469
##
     ENSG00000005471
## rowData metadata column names(0):
## colnames(3): SRR479052.bam SRR479053.bam SRR479054.bam
## colData names(0):
head( assay(se) )
##
                   SRR479052.bam SRR479053.bam SRR479054.bam
                                0
                                              0
## ENSG0000000003
                                                             1
## ENSG0000000005
                                0
                                              0
                                                             ()
## ENSG0000000419
                                              0
                                                             0
                                0
## ENSG0000000457
                                0
                                              1
                                                             1
## ENSG0000000460
                                0
                                              0
                                                             0
                                                             0
## ENSG00000000938
```

```
colSums( assay(se) )
## SRR479052.bam SRR479053.bam SRR479054.bam
##
              31
                            21
                                          30
colData(se)
## DataFrame with 3 rows and 0 columns
rowData(se)
## GRangesList of length 100:
## $ENSG0000000003
## GRanges with 13 ranges and 2 metadata columns:
##
          seqnames
                                 ranges strand
                                                     exon_id
                                                               exon_name
             <Rle>
##
                              <IRanges>
                                         <Rle>
                                                   <integer> <character>
      [1]
                X [99883667, 99884983]
##
                                                        3038
                                                                    < NA >
##
      [2]
                X [99885756, 99885863]
                                                        3039
                                                                    < NA >
      [3]
                X [99887482, 99887565]
                                                        3040
##
                                                                    < NA >
      [4]
                X [99887538, 99887565]
##
                                                        3041
                                                                    < NA >
                X [99888402, 99888536]
##
      [5]
                                                        3042
                                                                    < NA >
##
      . . .
##
                X [99890555, 99890743]
      [9]
                                                        3046
                                                                    < NA >
                X [99891188, 99891686]
     [10]
                                                        3047
##
                                                                    < NA >
     [11]
                X [99891605, 99891803]
                                                                    <NA>
##
                                                        3048
##
     [12]
                X [99891790, 99892101]
                                                        3049
                                                                    < NA >
##
     [13]
                X [99894942, 99894988]
                                                        3050
                                                                    <NA>
##
## ...
## <99 more elements>
## ---
## seqlengths:
## 7 12 2 6 16 4 3 1 17 8 19 X 11 9 20
```

Note that the rowData slot is a *GRangesList*, which contains all the information about the exons for each gene, i.e., for each row of the count table.

The colData slot, so far empty, should contain all the meta data. We hence assign our sample table to it:

```
colData(se) <- DataFrame( sampleTable )</pre>
```

We can also use the sampleName table to name the columns of our data matrix:

```
colnames(se) <- sampleTable$sampleName
head( assay(se) )
## Ctrl_24h_1 Ctrl_48h_1 DPN_24h_1</pre>
```

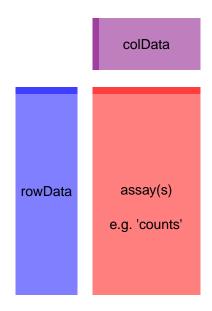


Figure 1: **Diagram of SummarizedExperiment** The component parts of a *SummarizedExperiment* object. The assay(s) (red block) contains the matrix (or matrices) of summarized values, the rowData (blue block) contains information about the genomic ranges, and the colData (purple block) contains information about the samples or experiments. The highlighted line in each block represents the first row (note that the first row of colData lines up with the first column of the assay.

##	ENSG0000000003	0	0	1
##	ENSG0000000005	0	0	0
##	ENSG00000000419	0	0	0
##	ENSG00000000457	0	1	1
##	ENSG00000000460	0	0	0
##	ENSG00000000938	0	0	0

This SummarizedExperiment object se is then all we need to start our analysis.

# 1.5 The DESeqDataSet, column metadata, and the design formula

Bioconductor software packages often have a special class of data object, which contains special slots and requirements. The data object class in *DESeq2* is the *DESeqDataSet*, which is built on top of the *SummarizedExperiment* class. One main differences is that the assay slot is instead accessed using the count accessor, and the values in this matrix must be non-negative integers.

A second difference is that the *DESeqDataSet* has an associated "design formula". The design is specified at the beginning of the analysis, as it will inform many of the *DESeq2* functions how to treat

the samples in the analysis (one exception is the size factor estimation, i. e., the adjustment for differing library sizes, which does not depend on the design formula). The design formula tells which variables in the column metadata table (colData) specify the experimental design and how these factors should be used in the analysis.

The simplest design formula for differential expression would be  $\sim$  condition, where condition is a column in colData(dds) which specifies which of two (or more groups) the samples belong to. For the parathyroid experiment, we will specify  $\sim$  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. Note that DESeq2 uses the same formula notation as, for instance, the 1m function of base R. If the question of interest is whether a fold change due to treatment is different across groups, for example across patients, "interaction terms" can be included using models such as  $\sim$  patient + treatment + patient:treatment. More complex designs such as these are covered in the other DESeq2 vignette.

We now use R's data command to load a prepared *SummarizedExperiment* that was generated from the publicly available sequencing data files associated with the Haglund et al. paper, described on page 1. The steps we used to produce this object were equivalent to those you worked through in Section 1.4, except that we used the complete set of samples and all reads.

```
data( "parathyroidGenesSE" )
se <- parathyroidGenesSE
colnames(se) <- se$run</pre>
```

Supposing we have constructed a *SummarizedExperiment* using one of the methods described in the previous section, we now need to make sure that the object contains all the necessary information about the samples, i.e., a table with metadata on the count table's columns stored in the colData slot:

```
colData(se)[1:5,1:4]
## DataFrame with 5 rows and 4 columns
##
                     run experiment patient treatment
                            <factor> <factor>
##
             <character>
                                               <factor>
## SRR479052
               SRR479052 SRX140503
                                            1
                                                 Control
## SRR479053
               SRR479053
                          SRX140504
                                             1
                                                Control
                          SRX140505
## SRR479054
                                             1
                                                     DPN
               SRR479054
## SRR479055
               SRR479055
                          SRX140506
                                             1
                                                     DPN
## SRR479056
               SRR479056 SRX140507
                                                     OHT
```

Here we see that this object already contains an informative colData slot – because we have already prepared it for you, as described in the *parathyroidSE* vignette. However, when you work with your own data, you will have to add the pertinent sample / phenotypic information for the experiment at this stage. We highly recommend keeping this information in a comma-separated value (CSV) or tabseparated value (TSV) file, which can be exported from an Excel spreadsheet, and the assign this to the colData slot, as shown in the previous section.

Make sure that the order of rows in your column data table matches the order of columns in the assay

data slot.

#### Exercise 4

How have we ensured this when building the se object in the previous section?

Once we have our fully annotated *SummerizedExperiment* object, we can construct a *DESeqDataSet* object from it, which will then form the staring point of the actual *DESeq2* package, described in the following sections. Here, we use the *SummerizedExperiment* object we got from the *parathyroidSE* package and augment it by specifying an appropriate design formula.

```
ddsFull <- DESeqDataSet( se, design = ~ patient + treatment )</pre>
```

Note that there are two alternative functions, DESeqDataSetFromMatrix and DESeqDataSetFromHTSeq, which allow you to get started in case you have your data not in the form of a *SummerizedExperiment* object, but either as a simple matrix of count values or a soutput files from the *htseq-count* script from the *HTSeq* Python package.

### 1.6 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( ddsFull )[ ,c("sample","patient","treatment","time") ] )
##
                 sample patient treatment time
## SRR479052 SRS308865
                                   Control
## SRR479053 SRS308866
                               1
                                   Control
                                            48h
                               1
## SRR479054 SRS308867
                                       DPN
                                            24h
## SRR479055 SRS308868
                               1
                                       DPN
                                            48h
## SRR479056 SRS308869
                               1
                                       OHT
                                            24h
## SRR479057 SRS308870
                               1
                                       OHT
                                            48h
## SRR479058 SRS308871
                               2
                                   Control
                                            24h
## SRR479059 SRS308872
                               2
                                   Control
                                            48h
                               2
## SRR479060 SRS308873
                                       DPN
                                            24h
## SRR479061 SRS308873
                               2
                                       DPN
                                            24h
## SRR479062 SRS308874
                               2
                                       DPN
                                            48h
                               2
## SRR479063 SRS308875
                                       OHT
                                            24h
## SRR479064 SRS308875
                               2
                                            24h
                                       OHT
                               2
## SRR479065 SRS308876
                                       OHT
                                            48h
## SRR479066 SRS308877
                               3
                                   Control
                                            24h
## SRR479067 SRS308878
                               3
                                   Control
                                            48h
## SRR479068 SRS308879
                               3
                                       DPN
                                            24h
## SRR479069 SRS308880
                               3
                                       DPN
                                            48h
                               3
## SRR479070 SRS308881
                                       OHT
                                            24h
                               3
## SRR479071 SRS308882
                                       OHT
                                            48h
```

```
## SRR479072 SRS308883
                               4
                                   Control
                                             48h
## SRR479073 SRS308884
                               4
                                       DPN
                                             24h
## SRR479074 SRS308885
                               4
                                       DPN
                                             48h
## SRR479075 SRS308885
                               4
                                       DPN
                                             48h
## SRR479076 SRS308886
                               4
                                             24h
                                       OHT
## SRR479077 SRS308887
                               4
                                       OHT
                                             48h
## SRR479078 SRS308887
                                       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. We have implemented a convenience function for this, which can take am object, either SummarizedExperiment or DESeqDataSet, and a grouping factor, in this case the sample name, and return the object with the counts summed up for each unique sample. This will also rename the columns of the object, such that they match the unique names which were used in the grouping factor. Optionally, we can provide a third argument, run, which can be used to paste together the names of the runs which were collapsed to create the new object. Note that dds\$variable is equivalent to colData(dds)\$variable.

```
ddsCollapsed <- collapseReplicates( ddsFull,</pre>
                                    groupby = ddsFull$sample,
                                    run = ddsFull$run )
head( as.data.frame( colData(ddsCollapsed)[ ,c("sample","runsCollapsed") ] ), 12 )
                              runsCollapsed
##
                 sample
## SRS308865 SRS308865
                                  SRR479052
## SRS308866 SRS308866
                                  SRR479053
## SRS308867 SRS308867
                                  SRR479054
## SRS308868 SRS308868
                                  SRR479055
## SRS308869 SRS308869
                                  SRR479056
## SRS308870 SRS308870
                                  SRR479057
## SRS308871 SRS308871
                                  SRR479058
## SRS308872 SRS308872
                                  SRR479059
## SRS308873 SRS308873 SRR479060, SRR479061
## SRS308874 SRS308874
                                  SRR479062
## SRS308875 SRS308875 SRR479063, SRR479064
## SRS308876 SRS308876
                                  SRR479065
```

We can confirm that the counts for the new object are equal to the summed up counts of the columns that had the same value for the grouping factor:

```
original <- rowSums( counts(ddsFull)[ , ddsFull$sample == "SRS308873" ] )
all( original == counts(ddsCollapsed)[ ,"SRS308873" ] )
## [1] TRUE</pre>
```

# 2 Running the DESeq2 pipeline

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

### 2.1 Preparing the data object for the analysis of interest

First we subset the relevant columns from the full dataset:

```
dds <- ddsCollapsed[ , ddsCollapsed$time == "48h" ]</pre>
```

Sometimes it is necessary to drop levels of the factors, in case that all the samples for one or more levels of a factor in the design have been removed. If time were included in the design formula, the following code could be used to take care of dropped levels in this column.

```
dds$time <- droplevels( dds$time )</pre>
```

It will be convenient to make sure that Control is the *first* level in the treatment factor, so that the default log2 fold changes are calculated as treatment over control and not the other way around. The function relevel achieves this:

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

A quick check whether we now have the right samples:

```
as.data.frame( colData(dds) )
##
                   run experiment patient treatment time submission
                                                                           study
## SRS308866 SRR479053
                        SRX140504
                                         1
                                              Control
                                                       48h
                                                            SRA051611 SRP012167
## SRS308868 SRR479055
                        SRX140506
                                         1
                                                  DPN
                                                       48h
                                                            SRA051611 SRP012167
## SRS308870 SRR479057
                                                  OHT
                        SRX140508
                                         1
                                                       48h
                                                            SRA051611 SRP012167
## SRS308872 SRR479059
                        SRX140510
                                         2
                                             Control
                                                       48h
                                                            SRA051611 SRP012167
## SRS308874 SRR479062
                                         2
                                                  DPN
                                                       48h
                                                            SRA051611 SRP012167
                        SRX140512
## SRS308876 SRR479065
                        SRX140514
                                         2
                                                  OHT
                                                       48h
                                                            SRA051611 SRP012167
## SRS308878 SRR479067
                        SRX140516
                                         3
                                             Control
                                                       48h
                                                            SRA051611 SRP012167
## SRS308880 SRR479069
                        SRX140518
                                         3
                                                  DPN
                                                       48h SRA051611 SRP012167
                                         3
                                                  OHT
## SRS308882 SRR479071
                        SRX140520
                                                       48h
                                                            SRA051611 SRP012167
## SRS308883 SRR479072
                        SRX140521
                                         4
                                             Control
                                                       48h
                                                            SRA051611 SRP012167
## SRS308885 SRR479074
                        SRX140523
                                         4
                                                  DPN
                                                       48h
                                                           SRA051611 SRP012167
## SRS308887 SRR479077
                        SRX140525
                                         4
                                                  OHT
                                                       48h
                                                            SRA051611 SRP012167
##
                              runsCollapsed
                sample
## SRS308866 SRS308866
                                  SRR479053
## SRS308868 SRS308868
                                  SRR479055
## SRS308870 SRS308870
                                  SRR479057
## SRS308872 SRS308872
                                  SRR479059
## SRS308874 SRS308874
                                  SRR479062
```

```
## SRS308876 SRS308876 SRR479065

## SRS308878 SRS308878 SRR479067

## SRS308880 SRS308880 SRR479069

## SRS308882 SRS308882 SRR479071

## SRS308883 SRS308883 SRR479072

## SRS308885 SRS308885 SRR479074,SRR479075

## SRS308887 SRS308887 SRR479077,SRR479078
```

### 2.2 Running the pipeline

Finally, we are ready to run the differential expression pipeline. With the data object prepared, the *DESeq2* analysis can now be run with a single call to the function DESeq:

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

This function will print out a message for the various steps it performs. These are described in more detail in the manual page for DESeq, which can be accessed by typing ?DESeq. Briefly these are: the estimation of size factors (which control for differences in the library size of the sequencing experiments), the estimation of dispersion for each gene, and fitting a generalized linear model.

A *DESeqDataSet* is returned which contains all the fitted information within it, and the following section describes how to extract out results tables of interest from this object.

## 2.3 Inspecting the results table

Calling results without any arguments will extract the estimated  $\log 2$  fold changes and p values for the last variable in the design formula. If there are more than 2 levels for this variable — as is the case in this analysis — results will extract the results table for a comparison of the last level over the first level. The following section describes how to extract other comparisons.

```
res <- results( dds )
res
## log2 fold change (MAP): treatment OHT vs Control
## Wald test p-value: treatment OHT vs Control
## DataFrame with 63193 rows and 6 columns
                     baseMean log2FoldChange
##
                                                  lfcSE
                                                             stat
                                                                      pvalue
##
                    <numeric>
                                   <numeric> <numeric> <numeric> <numeric>
                       613.82
## ENSG0000000003
                                    -0.04480
                                                 0.0879
                                                          -0.5098
                                                                    0.61017
## ENSG0000000005
                         0.55
                                    -0.56833
                                                 1.0875
                                                          -0.5226
                                                                    0.60127
## ENSG0000000419
                       304.05
                                     0.11612
                                                 0.0962
                                                           1.2067
                                                                    0.22755
## ENSG0000000457
                       183.52
                                     0.00744
                                                 0.1231
                                                           0.0604
                                                                    0.95182
## ENSG0000000460
                       207.43
                                     0.47084
                                                 0.1449
                                                           3.2487
                                                                     0.00116
## ...
```

##	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></numeric>					
##	ENSG0000000003	0.983					
##	ENSG0000000005	NA					
##	ENSG00000000419	NA					
##	ENSG00000000457	NA					
##	ENSG00000000460	NA					
##							
##	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, use.names=TRUE)
## DataFrame with 6 rows and 2 columns
##
                                                                      description
                           type
##
                   <character>
                                                                      <character>
## baseMean
                  intermediate
                                                     the base mean over all rows
## log2FoldChange
                       results log2 fold change (MAP): treatment OHT vs Control
## lfcSE
                        results
                                        standard error: treatment OHT vs Control
## stat
                        results
                                        Wald statistic: treatment OHT vs Control
                        results
## pvalue
                                     Wald test p-value: treatment OHT vs Control
## padj
                        results
                                                             BH adjusted p-values
```

The first column, baseMean, is a just the average of the normalized count values, dividing by size factors, 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 log2 fold change of 1.5 means that the gene's expression is increased by a multiplicative factor of  $2^{1.5}\approx 2.82$ .

Of course, this estimate has an uncertainty associated with it, which is available in the column 1fcSE, 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. 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.)

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. In addition, p values can be assigned NA if the gene was excluded from analysis because it contained an extreme count outlier. For more information, see the outlier detection section of the advanced vignette.

### 2.4 Other comparisons

In general, the results for a comparison of any two levels of a variable can be extracted using the contrast argument to results. The user should specify three values: the name of the variable, the name of the level in the numerator, and the name of the level in the denominator. Here we extract results for the log2 of the fold change of DPN / Control.

```
res <- results( dds, contrast = c("treatment", "DPN", "Control") )</pre>
res
## log2 fold change (MAP): treatment DPN vs Control
## Wald test p-value: treatment DPN vs Control
## DataFrame with 63193 rows and 6 columns
##
                     baseMean log2FoldChange
                                                    lfcSE
                                                                stat
                                                                        pvalue
##
                    <numeric>
                                    <numeric> <numeric> <numeric> <numeric>
## ENSG0000000003
                        613.82
                                                                        0.8425
                                       -0.0172
                                                   0.0867
                                                            -0.1987
                          0.55
                                                   1.0936
                                                            -0.0946
                                                                        0.9246
## ENSG0000000005
                                       -0.1034
## ENSG0000000419
                        304.05
                                       -0.0169
                                                   0.0952
                                                            -0.1781
                                                                        0.8587
## ENSG0000000457
                       183.52
                                       -0.0965
                                                   0.1214
                                                            -0.7953
                                                                        0.4264
## ENSG00000000460
                        207.43
                                        0.3500
                                                   0.1438
                                                             2.4350
                                                                        0.0149
## ...
                                           . . .
                                                                 . . .
                                                                            . . .
## LRG_94
                             0
                                            NA
                                                       NA
                                                                  NA
                                                                            NA
## LRG_96
                             0
                                            NA
                                                                  NA
                                                                            NA
                                                       NA
                             0
## LRG_97
                                            NA
                                                       NA
                                                                  NA
                                                                            NA
## LRG_98
                             0
                                            NA
                                                                  NA
                                                                            NA
                                                       NA
                             0
## LRG_99
                                            NA
                                                       NA
                                                                  NA
                                                                            NA
##
                          padj
##
                    <numeric>
## ENSG0000000003
                         0.976
## ENSG0000000005
                            NA
## ENSG00000000419
                         0.980
```

```
## ENSG0000000457
                         0.889
## ENSG0000000460
                         0.273
## ...
                           . . .
## LRG_94
                            NA
## LRG_96
                            NA
## LRG_97
                            NA
## LRG_98
                            NA
## LRG_99
                            NA
```

If results for an interaction term are desired, the name argument of results should be used. Please see the more advanced vignette for more details.

### 2.5 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

```
columns(org.Hs.eg.db)
##
    [1] "ENTREZID"
                        "PFAM"
                                        "IPI"
                                                         "PROSITE"
                                                                         "ACCNUM"
##
    [6] "ALIAS"
                        "CHR"
                                        "CHRLOC"
                                                         "CHRLOCEND"
                                                                         "ENZYME"
## [11] "MAP"
                        "PATH"
                                                         "REFSEQ"
                                                                         "SYMBOL"
                                        "PMID"
## [16] "UNIGENE"
                                                        "ENSEMBLTRANS" "GENENAME"
                        "ENSEMBL"
                                        "ENSEMBLPROT"
## [21] "UNIPROT"
                        "GO"
                                                         "ONTOLOGY"
                                                                         "GOALL"
                                        "EVIDENCE"
## [26] "EVIDENCEALL"
                        "ONTOLOGYALL"
                                        "MIMO"
                                                        "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, columns=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 ] )</pre>
```

```
}
```

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$hgnc_symbol <- convertIDs( row.names(res), "ENSEMBL", "SYMBOL", org.Hs.eg.db )</pre>
res$entrezgene <- convertIDs( row.names(res), "ENSEMBL", "ENTREZID", org.Hs.eg.db )
res
## log2 fold change (MAP): treatment DPN vs Control
## Wald test p-value: treatment DPN vs Control
## DataFrame with 63193 rows and 8 columns
##
                     baseMean log2FoldChange
                                                   lfcSE
                                                               stat
                                                                        pvalue
##
                    <numeric>
                                    <numeric> <numeric> <numeric> <numeric>
                       613.82
                                      -0.0172
                                                  0.0867
                                                            -0.1987
                                                                        0.8425
## ENSG00000000003
## ENSG0000000005
                         0.55
                                      -0.1034
                                                  1.0936
                                                            -0.0946
                                                                        0.9246
## ENSG0000000419
                       304.05
                                      -0.0169
                                                  0.0952
                                                            -0.1781
                                                                        0.8587
## ENSG0000000457
                       183.52
                                      -0.0965
                                                  0.1214
                                                            -0.7953
                                                                        0.4264
## ENSG0000000460
                       207.43
                                       0.3500
                                                  0.1438
                                                             2.4350
                                                                        0.0149
## ...
                                                      . . .
                                                                            . . .
                           . . .
                                           . . .
                                                                 . . .
## 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 hgnc_symbol entrezgene
##
                    <numeric> <character> <character>
                        0.976
                                    TSPAN6
## ENSG00000000003
                                                   7105
## ENSG0000000005
                           NA
                                      TNMD
                                                  64102
## ENSG00000000419
                        0.980
                                      DPM1
                                                   8813
## ENSG0000000457
                        0.889
                                     SCYL3
                                                  57147
## ENSG0000000460
                        0.273
                                  Clorf112
                                                  55732
## ...
## LRG_94
                           NA
                                        NA
                                                     NA
## LRG_96
                           NA
                                        NA
                                                     NA
## LRG_97
                            NA
                                        NA
                                                     NA
## LRG_98
                           NA
                                        NA
                                                     NA
## LRG_99
                                                     NA
                            NA
                                        NA
```

Now the results have the desired external gene ids:

```
head(res,4)
## log2 fold change (MAP): treatment DPN vs Control
## Wald test p-value: treatment DPN vs Control
## DataFrame with 4 rows and 8 columns
##
                     baseMean log2FoldChange
                                                  lfcSE
                                                                      pvalue
                                                             stat
##
                    <numeric>
                                   <numeric> <numeric> <numeric> <numeric>
## ENSG0000000003
                       613.82
                                     -0.0172
                                                 0.0867
                                                          -0.1987
                                                                       0.843
## ENSG0000000005
                         0.55
                                     -0.1034
                                                 1.0936
                                                          -0.0946
                                                                       0.925
## ENSG0000000419
                       304.05
                                     -0.0169
                                                 0.0952
                                                          -0.1781
                                                                       0.859
## ENSG0000000457
                                                          -0.7953
                       183.52
                                     -0.0965
                                                 0.1214
                                                                       0.426
##
                         padj hgnc_symbol entrezgene
##
                    <numeric> <character> <character>
                        0.976
                                   TSPAN6
## ENSG0000000003
                                                  7105
## ENSG0000000005
                           NA
                                     TNMD
                                                 64102
## ENSG00000000419
                        0.980
                                     DPM1
                                                  8813
## ENSG0000000457
                        0.889
                                    SCYL3
                                                 57147
```

#### Exercise 5

Go to the Ensembl web site, select the Biomart tab, and redo our BioMart query by manually inputting some of the Ensembl IDs and finding the HGNC names. What other data is available from this mart? Can you modify the code chunk above to add a column chrom to the res object that tells us for each gene which chromosome it resides on?

# 3 Further points

# 3.1 Multiple testing

Novices in high-throughput biology often assume that thresholding these p values at a low value, say 0.01, as is often done in other settings, would be appropriate – but it is not. We briefly explain why:

There are 495 genes with a p value below 0.01 among the 32082 genes, for which the test succeeded in reporting a p value:

```
sum( res$pvalue < 0.01, na.rm=TRUE )
## [1] 495

table( is.na(res$pvalue) )
##
## FALSE TRUE
## 32082 31111</pre>
```

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 1% of the genes to

have a p value below 0.01. This amounts to 321 genes. If we just considered the list of genes with a p value below 0.01 as differentially expressed, this list should therefore be expected to contain up to 321/495 = 65% 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] 248</pre>
```

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$padj < 0.1 ), ]</pre>
head( resSig[ order( resSig$log2FoldChange ), ] )
## log2 fold change (MAP): treatment DPN vs Control
## Wald test p-value: treatment DPN vs Control
## DataFrame with 6 rows and 8 columns
##
                    baseMean log2FoldChange
                                                  lfcSE
                                                                     pvalue
                                                             stat
##
                   <numeric>
                                   <numeric> <numeric> <numeric> <numeric>
                                                                   1.06e-03
## ENSG0000163631
                          233
                                      -0.931
                                                  0.284
                                                            -3.27
                                      -0.690
## ENSG00000119946
                                                  0.157
                                                            -4.41
                                                                   1.04e-05
                          152
## ENSG0000041982
                         1377
                                      -0.686
                                                  0.185
                                                            -3.72
                                                                   2.02e-04
## ENSG0000155111
                          531
                                      -0.676
                                                  0.211
                                                            -3.20
                                                                   1.36e-03
## ENSG00000233705
                          199
                                      -0.673
                                                  0.145
                                                            -4.65
                                                                   3.32e-06
## ENSG0000091137
                                                            -6.27
                                                                   3.70e-10
                         1144
                                      -0.654
                                                  0.104
##
                         padj hgnc_symbol entrezgene
##
                   <numeric> <character> <character>
## ENSG0000163631
                    5.69e-02
                                      ALB
                                                   213
## ENSG0000119946
                    2.42e-03
                                                 26507
                                    CNNM1
## ENSG00000041982
                    1.99e-02
                                      TNC
                                                  3371
## ENSG0000155111
                    6.69e-02
                                    CDK19
                                                 23097
## ENSG00000233705
                    1.06e-03 SLC26A4-AS1
                                                286002
## ENSG0000091137
                    8.56e-07
                                  SLC26A4
                                                  5172
```

#### and with the strongest upregulation

```
tail( resSig[ order( resSig$log2FoldChange ), ] )
## log2 fold change (MAP): treatment DPN vs Control
## Wald test p-value: treatment DPN vs Control
```

##	DataFrame with 6	rows and	8 columns				
##		baseMean	log2FoldChan	ıge	lfcSE	stat	pvalue
##		<numeric></numeric>	<numeri< td=""><td>ic&gt; ·</td><td><numeric></numeric></td><td><numeric></numeric></td><td><numeric></numeric></td></numeri<>	ic> ·	<numeric></numeric>	<numeric></numeric>	<numeric></numeric>
##	ENSG00000070669	282	0.5	547	0.162	3.39	7.08e-04
##	ENSG00000159307	236	0.6	521	0.126	4.94	7.90e-07
##	ENSG00000158457	282	0.6	529	0.162	3.88	1.06e-04
##	ENSG00000103257	168	0.8	326	0.164	5.02	5.06e-07
##	ENSG00000101255	255	0.8	386	0.169	5.23	1.71e-07
##	ENSG00000092621	559	0.9	900	0.126	7.16	8.19e-13
##		padj	hgnc_symbol	en <sup>-</sup>	trezgene		
##		<numeric></numeric>	<character></character>	<ch< td=""><td>aracter&gt;</td><td></td><td></td></ch<>	aracter>		
##	ENSG00000070669	4.68e-02	ASNS		440		
##	ENSG00000159307	4.57e-04	SCUBE1		80274		
##	ENSG00000158457	1.26e-02	TSPAN33		340348		
##	ENSG00000103257	3.12e-04	SLC7A5		8140		
##	ENSG00000101255	2.25e-04	TRIB3		57761		
##	ENSG00000092621	2.53e-09	PHGDH		26227		

### 3.2 Diagnostic plots

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

```
plotMA(res, ylim = c(-3, 3))
```

The plot (Fig. 2) represents each gene with a dot. The x axis is the average expression over all samples, the y axis the log2 fold change between treatment and control. Genes with an adjusted p value below a threshold (here 0.1, the default) are shown in red.

#### Exercise 6

Are the fold changes seen in this data set strong or weak compared to other gene expression data you may have seen? Can you find the names of the genes with the strongest differences?

This plot demonstrates that only genes with a large average normalized count contain sufficient information to yield a significant call.

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 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:

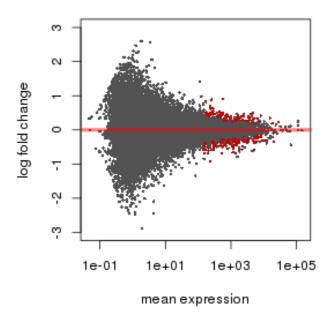


Figure 2: **MA-plot** The MA-plot shows the log2 fold changes from the treatment over the mean of normalized counts, i.e. the average of counts normalized by size factor. The *DESeq2* package incorporates a prior on log2 fold changes, resulting in moderated estimates from genes with low counts and highly variable counts, as can be seen by the narrowing of spread of points on the left side of the plot.

```
plotDispEsts( dds, ylim = c(1e-6, 1e1) )
```

The black points 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 points) that are then used in the hypothesis test. The blue circles above the main "cloud" of points are genes which have high gene-wise dispersion estimates which are labelled as dispersion outliers. These estimates are therefore not shrunk toward the fitted trend line.

#### Exercise 7

What can you learn from the dispersion plot about the typical within-group variability of gene-expression in the parathyphoid data set?

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

```
hist( res$pvalue, breaks=20, col="grey" )
```

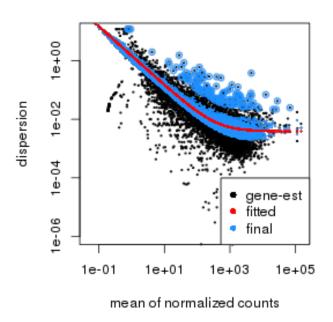


Figure 3: Plot of dispersion estimates See text for details

## Histogram of res\$pvalue

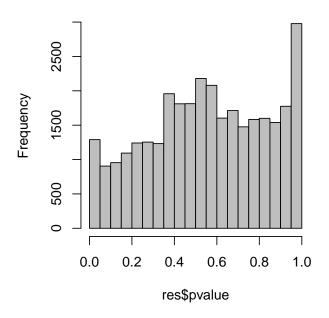


Figure 4: **Histogram** of the p values returned by the test for differential expression

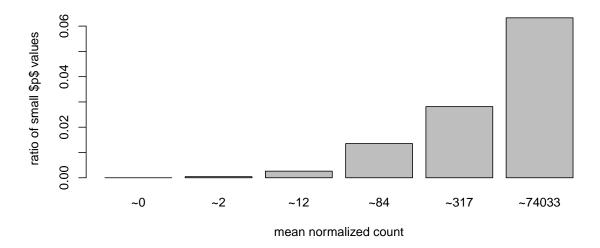


Figure 5: Ratio of small p values for groups of genes binned by mean normalized count

### 3.3 Independent filtering

The MA plot (Figure 2) 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. We can also show this by examining the ratio of small p values (say, less than, 0.01) for genes binned by mean normalized count:

```
# create bins using the quantile function
qs <- c( 0, quantile( res$baseMean[res$baseMean > 0], 0:7/7 ) )
# "cut" the genes into the bins
bins <- cut( res$baseMean, qs )
# rename the levels of the bins using the middle point
levels(bins) <- pasteO("~",round(.5*qs[-1] + .5*qs[-length(qs)]))
# calculate the ratio of £p£ values less than .01 for each bin
ratios <- tapply( res$pvalue, bins, function(p) mean( p < .01, na.rm=TRUE ) )
# plot these ratios
barplot(ratios, xlab="mean normalized count", ylab="ratio of small $p$ values")</pre>
```

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

The DESeq2 software automatically performs independent filtering which maximizes the number of genes which will have adjusted p value less than a critical value (by default, alpha is set to 0.1). This

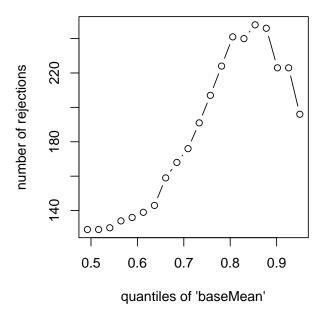


Figure 6: **Independent filtering.** DESeq2 automatically determines a threshold, filtering on mean normalized count, which maximizes the number of genes which will have an adjusted p value less than a critical value.

automatic independent filtering is performed by, and can be controlled by, the results function. We can observe how the number of rejections changes for various cutoffs based on mean normalized count. The following optimal threshold and table of possible values is stored as an *attribute* of the results object.

The term *independent* highlights an important caveat. Such filtering is permissible only if the filter criterion is independent of the actual test statistic [?]. 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.

### 3.4 Exporting results

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:

```
res[1:2,]
## log2 fold change (MAP): treatment DPN vs Control
## Wald test p-value: treatment DPN vs Control
## DataFrame with 2 rows and 8 columns
##
                    baseMean log2FoldChange
                                                 lfcSE
                                                             stat
                                                                     pvalue
##
                   <numeric>
                                   <numeric> <numeric> <numeric> <numeric>
## ENSG00000000003
                      613.82
                                     -0.0172
                                                0.0867
                                                          -0.1987
                                                                      0.843
## ENSG0000000005
                        0.55
                                     -0.1034
                                                          -0.0946
                                                1.0936
                                                                      0.925
                        padj hgnc_symbol entrezgene
##
##
                   <numeric> <character> <character>
                        0.976
                                   TSPAN6
## ENSG00000000003
                                                 7105
## ENSG0000000005
                          NA
                                     TNMD
                                                 64102
write.csv( as.data.frame(res), file="results.csv" )
```

# 3.5 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.

As noted in the lecture, gene-set enrichment analysis with RNA-Seq data entails some subtleties. Briefly, a number of different approaches have been proposed that each imply slightly different null hypotheses that are being tested against, and their biologically interpretation differs. This is a topic of ongoing research. We here present a relatively simplistic approach, to demonstrate the basic ideas, but note that a more careful treatment will be needed for more definitive results.

We use the gene sets in the Reactome database

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

This database works with Entrez IDs, so we will need the entrezgene column that we added earlier to the res object.

First, 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 DESeq2 test gave an adjusted p value that was not NA.

```
res2 <- res[ res$entrezgene %in% keys( reactome.db, "ENTREZID" ) &
  !is.na( res$padj ) , ]
head(res2)
## log2 fold change (MAP): treatment DPN vs Control</pre>
```

```
## Wald test p-value: treatment DPN vs Control
## DataFrame with 6 rows and 8 columns
##
                    baseMean log2FoldChange
                                                  lfcSE
                                                             stat
                                                                      pvalue
##
                    <numeric>
                                   <numeric> <numeric> <numeric> <numeric>
## ENSG0000000419
                          304
                                    -0.01695
                                                 0.0952
                                                           -0.178
                                                                       0.859
## ENSG0000001084
                          312
                                     0.03981
                                                 0.1023
                                                            0.389
                                                                       0.697
## ENSG0000001167
                          399
                                    -0.05600
                                                 0.0971
                                                           -0.577
                                                                       0.564
## ENSG0000001630
                          464
                                     0.07000
                                                 0.0968
                                                            0.723
                                                                       0.469
## ENSG00000002822
                                    -0.00193
                                                           -0.014
                                                                       0.989
                          170
                                                 0.1378
## ENSG0000003056
                          996
                                     0.00772
                                                 0.0678
                                                            0.114
                                                                       0.909
                        padj hgnc_symbol entrezgene
##
##
                    <numeric> <character> <character>
## ENSG00000000419
                        0.980
                                     DPM1
                                                  8813
## ENSG0000001084
                        0.951
                                     GCLC
                                                  2729
## ENSG0000001167
                        0.931
                                     NFYA
                                                  4800
## ENSG0000001630
                        0.907
                                  CYP51A1
                                                  1595
## ENSG00000002822
                        0.998
                                   MAD1L1
                                                  8379
## ENSG0000003056
                                                  4074
                        0.985
                                     M6PR
```

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,</pre>
   keys=as.character(res2$entrezgene), 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
                   446203
         8813
## 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$entrezgene %in% x ) ))
colnames(incm) <- res2$entrez</pre>
```

```
str(incm)
## logi [1:1454, 1:3296] FALSE FALSE FALSE FALSE FALSE FALSE FALSE ...
## - attr(*, "dimnames")=List of 2
## ..$ : chr [1:1454] "1059683" "109581" "109582" "109606" ...
## ..$ : chr [1:3296] "8813" "2729" "4800" "1595" ...
```

We remove all rows corresponding to Reactome Paths with less than 20 or more than 80 assigned genes.

```
within <- function(x, lower, upper) (x>=lower & x<=upper)
incm <- incm[ within(rowSums(incm), lower=20, upper=80), ]</pre>
```

To test whether the genes in a Reactome Path behave in a special way in our experiment, we calculate a number of statistics, including a t-statistic to see whether the average of the genes'  $\log_2$  fold change values in the gene set is different from zero. To facilitate the computations, we define a little helper function:

```
testCategory <- function( reactomeID ) {</pre>
  isMember <- incm[ reactomeID, ]</pre>
 data.frame(
    reactomeID = reactomeID,
    numGenes = sum( isMember ),
    avgLFC = mean( res2$log2FoldChange[isMember] ),
                = sd( res2$log2FoldChange[isMember] ),
    sdLFC
    zValue
              = mean( res2$log2FoldChange[isMember] ) /
                    sd( res2$log2FoldChange[isMember] ),
    strength = sum( res2$log2FoldChange[isMember] ) / sqrt(sum(isMember)),
                = t.test( res2$log2FoldChange[ isMember ] )$p.value,
    pvalue
    reactomeName = reactomePATHID2NAME[[reactomeID]],
    stringsAsFactors = FALSE ) }
```

The function can be called with a Reactome Path ID:

```
testCategory("109606")

## reactomeID numGenes avgLFC sdLFC zValue strength pvalue
## 1 109606  30 -0.0224 0.079 -0.283 -0.122 0.132
## reactomeName
## 1 Homo sapiens: Intrinsic Pathway for 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 ) )</pre>
```

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 sdLFC zValue strength
                                                             pvalue
                        79 -0.0413 0.0982 -0.421
## 377
            69242
                                                    -0.367 3.51e-04
## 417
            73894
                        67 -0.0442 0.1125 -0.393
                                                    -0.362 2.02e-03
## 366
            68877
                        64 -0.0434 0.0780 -0.557
                                                    -0.347 3.50e-05
## 230
          2500257
                        59 -0.0443 0.0752 -0.590
                                                    -0.341 3.00e-05
## 199
           212165
                        46 -0.0502 0.0943 -0.532
                                                    -0.341 7.67e-04
## 92
                        51 -0.0451 0.0816 -0.553
                                                    -0.322 2.46e-04
           167172
## 93
           167172
                        51 -0.0451 0.0816 -0.553
                                                    -0.322 2.46e-04
## 312
                        33 -0.0532 0.0756 -0.704
                                                    -0.306 3.11e-04
           390466
## 86
           167161
                        32 -0.0523 0.0744 -0.704
                                                    -0.296 3.86e-04
## 87
           167161
                        32 -0.0523 0.0744 -0.704
                                                    -0.296 3.86e-04
## 88
           167162
                        32 -0.0523 0.0744 -0.704
                                                    -0.296 3.86e-04
## 89
                        32 -0.0523 0.0744 -0.704
                                                    -0.296 3.86e-04
           167162
                        32 -0.0523 0.0744 -0.704
## 409
            73776
                                                    -0.296 3.86e-04
                        32 -0.0523 0.0744 -0.704
## 410
            73779
                                                    -0.296 3.86e-04
## 426
            75953
                        32 -0.0523 0.0744 -0.704
                                                    -0.296 3.86e-04
## 429
            76042
                        32 -0.0523 0.0744 -0.704
                                                    -0.296 3.86e-04
## 395
            72086
                        25 -0.0591 0.0857 -0.690
                                                    -0.296 2.09e-03
## 84
           167160
                        23 -0.0603 0.0879 -0.686
                                                    -0.289 3.34e-03
## 85
                        23 -0.0603 0.0879 -0.686
                                                    -0.289 3.34e-03
           167160
## 435
            77075
                        23 -0.0603 0.0879 -0.686
                                                    -0.289 3.34e-03
## 361
           674695
                        49 -0.0397 0.0780 -0.509
                                                    -0.278 8.49e-04
## 314
           391251
                        37 -0.0450 0.0771 -0.584
                                                    -0.274 1.09e-03
## 396
            72165
                        37 -0.0422 0.0742 -0.568
                                                    -0.256 1.42e-03
## 347
           450531
                        65 -0.0315 0.0763 -0.413
                                                    -0.254 1.44e-03
## 359
          5250941
                        34 -0.0428 0.0789 -0.542
                                                    -0.249 3.37e-03
## 402
            72689
                        64 0.0291 0.0728 0.399
                                                     0.232 2.21e-03
## 148
          1799339
                        75
                            0.0408 0.0779 0.524
                                                     0.354 2.15e-05
## 58
          1638091
                            0.0688 0.1074 0.640
                                                     0.377 1.50e-03
                        30
## 305
           381070
                        41
                            0.0706 0.1390 0.508
                                                     0.452 2.34e-03
           446203
                        79
                            0.0557 0.1137 0.490
                                                     0.495 4.02e-05
## 339
                        21 0.1081 0.0989 1.092
                                                     0.495 6.79e-05
## 172
           191273
                                                                                   reactomeName
##
## 377
                                                                          Homo sapiens: S Phase
## 417
                                                                       Homo sapiens: DNA Repair
## 366
                                                            Homo sapiens: Mitotic Prometaphase
## 230
                                         Homo sapiens: Resolution of Sister Chromatid Cohesion
## 199
                                        Homo sapiens: Epigenetic regulation of gene expression
## 92
                                                 Homo sapiens: Transcription of the HIV genome
## 93
                               Human immunodeficiency virus 1: Transcription of the HIV genome
```

```
## 312
                                            Homo sapiens: Chaperonin-mediated protein folding
## 86
                                                    Homo sapiens: HIV Transcription Initiation
## 87
                                 Human immunodeficiency virus 1: HIV Transcription Initiation
## 88
                                          Homo sapiens: RNA Polymerase II HIV Promoter Escape
## 89
                        Human immunodeficiency virus 1: RNA Polymerase II HIV Promoter Escape
                                               Homo sapiens: RNA Polymerase II Promoter Escape
## 409
## 410
            Homo sapiens: RNA Polymerase II Transcription Pre-Initiation And Promoter Opening
## 426
                                     Homo sapiens: RNA Polymerase II Transcription Initiation
## 429
              Homo sapiens: RNA Polymerase II Transcription Initiation And Promoter Clearance
## 395
                                                                    Homo sapiens: mRNA Capping
                         Homo sapiens: RNA Pol II CTD phosphorylation and interaction with CE
## 84
## 85
       Human immunodeficiency virus 1: RNA Pol II CTD phosphorylation and interaction with CE
## 435
                         Homo sapiens: RNA Pol II CTD phosphorylation and interaction with CE
## 361
                                     Homo sapiens: RNA Polymerase II Pre-transcription Events
## 314
                                                                 Homo sapiens: Protein folding
## 396
                                                   Homo sapiens: mRNA Splicing - Minor Pathway
## 347
            Homo sapiens: Regulation of mRNA stability by proteins that bind AU-rich elements
## 359
                              Homo sapiens: Negative epigenetic regulation of rRNA expression
## 402
                                       Homo sapiens: Formation of a pool of free 40S subunits
## 148
                    Homo sapiens: SRP-dependent cotranslational protein targeting to membrane
## 58
                                    Homo sapiens: Heparan sulfate/heparin (HS-GAG) metabolism
## 305
                                                 Homo sapiens: IRE1alpha activates chaperones
## 339
                                              Homo sapiens: Asparagine N-linked glycosylation
## 172
                                                        Homo sapiens: Cholesterol biosynthesis
##
       padjust
## 377 0.01039
## 417 0.03823
## 366 0.00460
## 230 0.00460
## 199 0.01952
## 92 0.01039
      0.01039
## 93
## 312 0.01039
## 86
      0.01039
## 87
       0.01039
## 88
      0.01039
## 89
      0.01039
## 409 0.01039
## 410 0.01039
## 426 0.01039
## 429 0.01039
## 395 0.03823
## 84
      0.04983
## 85 0.04983
## 435 0.04983
## 361 0.02047
## 314 0.02500
```

```
## 396 0.02984
## 347 0.02984
## 359 0.04983
## 402 0.03898
## 148 0.00460
## 58 0.02984
## 305 0.03971
## 339 0.00460
## 172 0.00622
```

However, as discussed in the lecture, it is highly questionable that a t test is appropriate here. After all, genes in a set are typically correlated, and this violates the assumption of independence that is at the core of a t test. Hence, should we really look at p values from t tests? A p value obtained by sample permutation would solve this issue as sample permutation preserves and so accounts for gene-gene correlation. However, with only four subjects, we do not have enough samples for this.

Hence, it may be more prudent to disregard these questionable p values altogether and instead look at a more primitive statistic, such as simply the average LFC within a path, perhaps divided by the standard deviation. We have stored this above as zValue.

```
head( reactomeResult[ order(reactomeResult$zValue), ] )
##
       reactomeID numGenes avgLFC sdLFC zValue strength
                                                             pvalue
                                                  -0.306 0.000311
## 312
           390466
                        33 -0.0532 0.0756 -0.704
## 86
           167161
                        32 -0.0523 0.0744 -0.704
                                                   -0.296 0.000386
## 87
           167161
                        32 -0.0523 0.0744 -0.704
                                                  -0.296 0.000386
## 88
           167162
                        32 -0.0523 0.0744 -0.704
                                                   -0.296 0.000386
                        32 -0.0523 0.0744 -0.704
                                                   -0.296 0.000386
## 89
           167162
                        32 -0.0523 0.0744 -0.704
                                                   -0.296 0.000386
## 409
            73776
##
                                                                 reactomeName
## 312
                           Homo sapiens: Chaperonin-mediated protein folding
## 86
                                  Homo sapiens: HIV Transcription Initiation
## 87
                Human immunodeficiency virus 1: HIV Transcription Initiation
                         Homo sapiens: RNA Polymerase II HIV Promoter Escape
## 88
## 89
       Human immunodeficiency virus 1: RNA Polymerase II HIV Promoter Escape
## 409
                             Homo sapiens: RNA Polymerase II Promoter Escape
##
       padjust
## 312 0.0104
## 86
        0.0104
## 87
        0.0104
## 88
        0.0104
## 89
        0.0104
## 409 0.0104
head( reactomeResult[ order(reactomeResult$zValue, decreasing=TRUE), ] )
##
       reactomeID numGenes avgLFC sdLFC zValue strength
## 172
           191273
                        21 0.1081 0.0989 1.092
                                                    0.495 6.79e-05
## 58
          1638091
                        30 0.0688 0.1074 0.640 0.377 1.50e-03
```

```
## 226
          2426168
                        26 0.0629 0.1163 0.541
                                                   0.321 1.07e-02
## 148
          1799339
                        75 0.0408 0.0779 0.524
                                                    0.354 2.15e-05
## 305
                        41 0.0706 0.1390 0.508
                                                    0.452 2.34e-03
           381070
## 44
           156580
                        31 0.0537 0.1084
                                          0.496
                                                    0.299 9.78e-03
##
                                                                     reactomeName
## 172
                                          Homo sapiens: Cholesterol biosynthesis
## 58
                       Homo sapiens: Heparan sulfate/heparin (HS-GAG) metabolism
                    Homo sapiens: Activation of gene expression by SREBF (SREBP)
## 226
## 148 Homo sapiens: SRP-dependent cotranslational protein targeting to membrane
## 305
                                    Homo sapiens: IRE1alpha activates chaperones
## 44
                                               Homo sapiens: Phase II conjugation
      padjust
##
## 172 0.00622
## 58 0.02984
## 226 0.08232
## 148 0.00460
## 305 0.03971
## 44 0.08142
```

If such an analysis is only considered exploratory, we may inspect various such tables and see whether the ranking of Paths helps us make sense of the data. Nevertheless, there is certainly room for improvement here.

#### Exercise 8

The first few hits in the ranking by z value all have exactly the same values. Why?

# 4 Working with rlog-transformed data

# 4.1 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 quantity (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 plus a small pseudocount; 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 log2 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.

Note that the rlog transformation is provided for applications other than differential testing. For differential testing we recommend the DESeq function applied to raw counts, as described earlier in this vignette, which also takes into account the dependence of the variance of counts on the mean value during the dispersion estimation step.

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

```
rld <- rlog( dds )</pre>
head( assay(rld) )
##
                    SRS308866 SRS308868 SRS308870 SRS308872 SRS308874 SRS308876
## ENSG0000000003
                                                         9.093
                        9.761
                                   9.728
                                              9.865
                                                                    9.16
                                                                               9.06
## ENSG00000000005
                       -0.981
                                  -0.741
                                                       -0.924
                                                                   -1.11
                                                                              -1.11
                                             -0.945
## ENSG0000000419
                                                        8.231
                                                                    8.29
                                                                               8.37
                        8.065
                                   8.077
                                              8.119
## ENSG0000000457
                                   7.250
                                              7.311
                                                        7.854
                                                                    7.73
                                                                               7.89
                        7.417
## ENSG0000000460
                        7.576
                                   7.691
                                              7.721
                                                        8.057
                                                                    8.26
                                                                               8.06
## ENSG0000000938
                                   2.936
                                              2.954
                                                        4.216
                                                                    3.64
                                                                               3.65
                        3.131
##
                    SRS308878 SRS308880 SRS308882 SRS308883 SRS308885 SRS308887
## ENSG00000000003
                         8.89
                                    8.79
                                               8.84
                                                          9.05
                                                                    9.07
                                                                               8.88
## ENSG00000000005
                        -1.11
                                   -1.12
                                              -1.11
                                                         -1.11
                                                                   -1.12
                                                                              -1.12
## ENSG0000000419
                                    8.31
                                                          8.26
                                                                    8.15
                                                                               8.32
                         8.31
                                               8.41
## ENSG0000000457
                                    7.30
                                               7.38
                                                          7.60
                                                                    7.48
                                                                               7.50
                         7.16
## ENSG0000000460
                         7.06
                                    7.44
                                               7.63
                                                          6.91
                                                                    7.34
                                                                               7.66
## ENSG0000000938
                         2.80
                                    3.33
                                               3.23
                                                          3.29
                                                                    3.45
                                                                               3.23
```

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) $^{5}$ .

In Figure 7, 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.

<sup>&</sup>lt;sup>5</sup>The function heatscatter from the package *LSD* offers a colorful alternative.

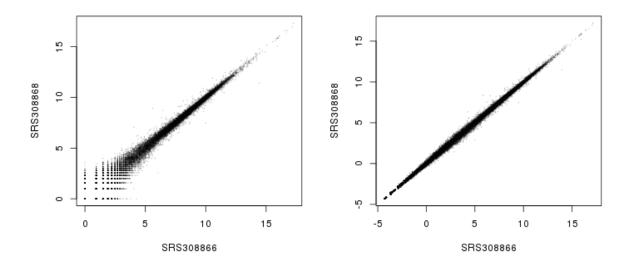


Figure 7: **Scatter plot of sample 2 vs sample 1.** Left: using an ordinary  $log_2$  transformation. Right: Using the rlog transformation.

### 4.2 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?

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:

	<pre>sampleDists &lt;- dist( t( assay(rld) ) ) sampleDists</pre>								
##		SRS308866	SRS308868	SRS308870	SRS308872	SRS308874	SRS308876	SRS308878	
##	SRS308868	32.9							
##	SRS308870	31.2	33.7						
##	SRS308872	83.8	86.6	85.6					
##	SRS308874	83.8	85.2	84.9	33.3				
##	SRS308876	82.8	85.6	84.6	32.0	33.0			
##	SRS308878	108.6	111.7	110.3	109.1	111.7	108.5		
##	SRS308880	108.7	111.1	110.1	110.5	112.3	110.1	31.7	
##	SRS308882	108.7	111.1	110.0	111.3	113.0	110.6	36.1	
##	SRS308883	95.1	98.4	97.5	91.6	93.9	90.5	81.6	
##	SRS308885	92.9	95.2	95.0	90.0	91.2	89.4	83.5	
##	SRS308887	95.9	98.7	98.0	95.0	96.6	93.7	82.9	
##		SRS308880	SRS308882	SRS308883	SRS308885				

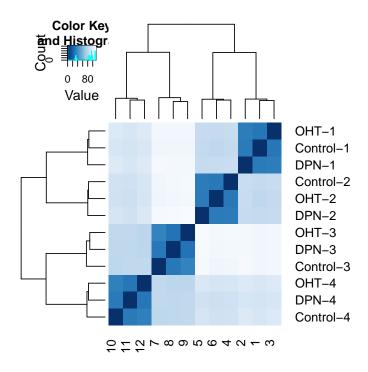


Figure 8: Heatmap of Euclidean sample distances after rlog transformation.

```
## SRS308868
## SRS308870
## SRS308872
## SRS308874
## SRS308876
## SRS308878
## SRS308880
## SRS308882
                   29.4
## SRS308883
                   82.1
                             83.0
## SRS308885
                   82.3
                             83.5
                                        32.7
## SRS308887
                   81.4
                             81.4
                                        34.9
                                                   31.2
```

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.

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

```
sampleDistMatrix <- as.matrix( sampleDists )
rownames(sampleDistMatrix) <- paste( rld$treatment,
    rld$patient, sep="-" )
colnames(sampleDistMatrix) <- NULL
library( "gplots" )
library( "RColorBrewer" )</pre>
```

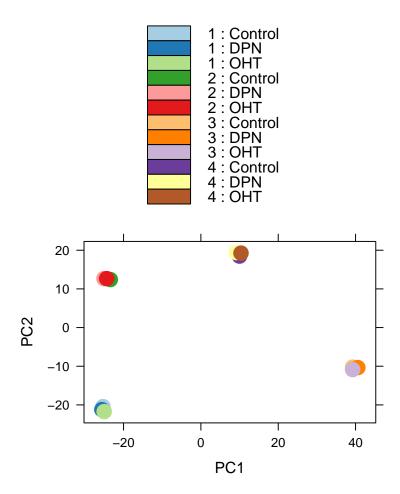


Figure 9: **Principal components analysis (PCA)** of samples after rlog transformation.

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

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. 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).

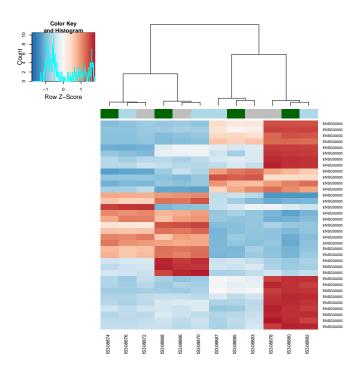


Figure 10: Heatmap with gene clustering.

```
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  $\sim$  patient + treatment when setting up the data object in the beginning. Had we used an un-paired analysis, by specifying only  $\sim$  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 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 package of the same name, useful.

# 4.3 Gene clustering

In the heatmap of Fig. 8, 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 )</pre>
```

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),
    ColSideColors = c( Control="gray", DPN="darkgreen", OHT="lightblue" )[
        colData(rld)$treatment ] )
```

#### Exercise 9

Which of the options to heatmap.2 cause the centering and scaling of the genes? How does the heatmap change if you do not scale it?

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 treatments from each patient.

#### Exercise 10

Still, this heatmap shows genes that vary strongest between patients. In our differential-expression analysis, we found a list of genes that varied significantly between treatment and control. Display these genes (or maybe only those showing up-regulation upon treatment) in a similar heatmap. Can you confirm the test result from visual inspection of the heatmap?

# 5 Session Info

As last part of this document, we call the function <code>sessionInfo</code>, 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. The session information should also always be included in any emails to the Bioconductor mailing list.

- R version 3.1.0 (2014-04-10), x86\_64-unknown-linux-gnu
- Locale: LC\_CTYPE=en\_US.UTF-8, LC\_NUMERIC=C, LC\_TIME=en\_US.UTF-8, LC\_COLLATE=en\_US.UTF-8, 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, graphics, grDevices, methods, parallel, stats, utils
- Other packages: AnnotationDbi 1.26.0, Biobase 2.24.0, BiocGenerics 0.10.0, Biostrings 2.32.0, BSgenome 1.32.0, DBI 0.2-7, DESeq2 1.4.5, genefilter 1.46.1, GenomeInfoDb 1.0.2,

- GenomicAlignments 1.0.1, GenomicFeatures 1.16.2, GenomicRanges 1.16.3, gplots 2.14.0, IRanges 1.22.9, knitr 1.6, org.Hs.eg.db 2.14.0, parathyroidSE 1.2.0, RColorBrewer 1.0-5, Rcpp 0.11.2, RcppArmadillo 0.4.300.8.0, reactome.db 1.48.0, Rsamtools 1.16.1, RSQLite 0.11.4, XVector 0.4.0
- Loaded via a namespace (and not attached): annotate 1.42.0, BatchJobs 1.2, BBmisc 1.7, BiocParallel 0.6.1, BiocStyle 1.2.0, biomaRt 2.20.0, bitops 1.0-6, brew 1.0-6, caTools 1.17, checkmate 1.1, codetools 0.2-8, digest 0.6.4, evaluate 0.5.5, fail 1.2, foreach 1.4.2, formatR 0.10, gdata 2.13.3, geneplotter 1.42.0, grid 3.1.0, gtools 3.4.1, highr 0.3, iterators 1.0.7, KernSmooth 2.23-12, lattice 0.20-29, locfit 1.5-9.1, plyr 1.8.1, RCurl 1.95-4.1, rtracklayer 1.24.2, sendmailR 1.1-2, splines 3.1.0, stats4 3.1.0, stringr 0.6.2, survival 2.37-7, tools 3.1.0, XML 3.98-1.1, xtable 1.7-3, zlibbioc 1.10.0