

# **tweeDEseq**: analysis of RNA-seq data using the Poisson-Tweedie family of distributions

Mikel Esnaola<sup>1</sup>                      Robert Castelo<sup>2,3</sup>  
mesnaola@creal.cat              robert.castelo@upf.edu

Juan Ramón González<sup>1,2,3</sup>  
jrgonzalez@creal.cat

March 2, 2012

<sup>1</sup>Centre for Research in Environmental Epidemiology (CREAL), Barcelona, Spain

<http://www.creal.cat/jrgonzalez/software.htm>

<sup>2</sup>Dept. of Experimental and Health Sciences, Research Program on Biomedical Informatics (GRIB), Universitat Pompeu Fabra, Barcelona, Spain

<sup>3</sup>Hospital del Mar Research Institute (IMIM), Barcelona, Spain

## 1 Getting started

This document gives an overview of the R package `tweeDEseq`, which provides statistical procedures for testing differential expression in RNA-seq count data. In fact, these procedures could be applied, in principle, to any kind of count data, other than RNA-seq. The `tweeDEseq` package offers a function for normalizing count data which actually calls other functions for that purpose from the `edgeR` package. For this reason, it is necessary to have installed the `edgeR` package as well, although it is not necessary to explicitly load it onto the session. Another package necessary for running this vignette is the `tweeDEseqCountData` package which contains data to illustrate the analyses and which is employed in the article introducing `tweeDEseq` by Esnaola et al. (2011). Therefore, we should start the R session with loading these libraries as follows:

```
> library(tweeDEseq)
> library(tweeDEseqCountData)
```

We will start loading into the workspace the data corresponding to the initial table of raw counts of the RNA-seq experiment from Pickrell et al. (2010) on lymphoblastoid cell lines derived from 69 unrelated nigerian individuals as well as a vector of gender labels for each sample matching the sample order in the table of counts:

```
> data(pickrell)
> countsNigerian <- exprs(pickrell.eset)
> dim(countsNigerian)
```

```
[1] 52580    69
```

```
> countsNigerian[1:5, 1:5]
```

	NA18486	NA18498	NA18499	NA18501	NA18502
ENSG00000000003	0	0	0	0	0
ENSG00000000005	0	0	0	0	0
ENSG000000000419	22	105	40	55	67
ENSG000000000457	22	100	107	53	72
ENSG000000000460	5	23	10	18	15

```
> genderNigerian <- pData(pickrell.eset)[,"gender"]
> head(genderNigerian)
```

```
[1] male male female male female male
Levels: female male
```

```
> table(genderNigerian)
```

```
genderNigerian
female male
    40    29
```

## 2 Normalization and filtering

We proceed to normalize this initial table of raw counts in order to try to remove any technical biases that might be affecting the data. The `tweeDEseq` package relies for this purpose on part of the functionality provided by the `edgeR` package (comprising RNA composition adjustment by TMM (Robinson and Oshlack, 2010) and quantile-to-quantile count adjustment) and offers one function (`normalizeCounts()`) that makes the appropriate calls to `edgeR` to normalize these data.

```
> countsNigerianNorm <- normalizeCounts(countsNigerian, genderNigerian)
```

```
> dim(countsNigerianNorm)
```

```
[1] 10231 69
```

If more control is needed in this step, the user should directly employ the corresponding `edgeR` functions. Next, we can filter out genes with very low expression using the function `filterCounts()` whose default parameters remove those genes with less than 5 counts per million in all samples but one.

```
> countsNigerianNorm <- filterCounts(countsNigerianNorm)
```

```
> dim(countsNigerianNorm)
```

```
[1] 10231 69
```

## 3 The Poisson-Tweedie family of distributions to model RNA-seq count data

The package `tweeDEseq` uses the Poisson-Tweedie (PT) family of distributions as the statistical model for count data. PT distributions have been studied by several authors (Hougaard et al., 1997; Gupta and Ong, 2004; Puig and Valero, 2006; El-Shaarawi et al., 2011) and unify several count data distributions (see **Fig. 1** in El-Shaarawi *et al.*, 2011) such as Poisson, negative

binomial, Poisson-Inverse Gaussian, Pólya-Aeppli or Neyman type A. These distributions can model different scenarios as, for instance, a RNA-seq expression profile with a wide dynamic range leading to a heavy tail in the distribution.

Following El-Shaarawi et al. (2011), let  $Y \sim \text{PT}(a, b, c)$  be a PT random variable with vector of parameters  $\theta = (a, b, c)^T$  defined in the domain

$$\Theta = (-\infty, 1] \times (0, +\infty) \times [0, 1). \quad (1)$$

For the sake of interpretability, we reparametrize  $\theta = (a, b, c)$  to  $\theta = (\mu, \phi, a)$ , where  $\mu$  denotes the mean,  $\phi = \sigma^2/\mu$  is the dispersion index ( $\sigma^2$  is the variance), and  $a$  the shape parameter that is employed to define some count data distributions that are particular cases of PT such as Poisson or Negative Binomial. The relationship between both parameterizations is the following:

$$c = \frac{\phi - 1}{\phi - a}, \quad b = \frac{\mu(1 - a)^{(1-a)}}{(\phi - 1)(d - a)^{-a}}. \quad (2)$$

Under this parametrization, the shape parameter determines the specific count data distribution being employed. For instance  $a = 1$  corresponds to Poisson and  $a = 0$  corresponds to negative binomial. We can estimate the parameter vector  $\theta$  by maximum likelihood through the function `mlePoissonTweedie()` as follows:

```
> set.seed(123)
> y <- rnbinom(1000, mu=8, size=1/0.2)
> thetahat <- mlePoissonTweedie(y)
> getParam(thetahat)
```

```
      mu      D      a
8.0680000 2.6620193 0.0405573
```

where here we have simulated 1000 random observations from a negative binomial distribution and the last call to `getParam()` allows us to extract the  $\hat{\theta}$  vector from the object return by `mlePoissonTweedie()`.

## 4 Goodness of fit to a count data distribution

The PT distribution allows one to test for the goodness of fit to a particular count data distribution defined by a specific value of the PT shape parameter  $a$ . For this purpose, the function `testShapePT()` allows us to test the goodness of fit to, for instance, the widely used negative binomial distribution (i.e.,  $H_0 : a = 0$ ) as illustrated here with the previously estimated vector  $\hat{\theta}$ :

```
> testShapePT(thetahat, a=0)
```

```
$statistic
[1] 0.01471032
```

```
$pvalue
[1] 0.9034644
```

These functions are called from another one called `gofTest()` which can perform for us a goodness-of-fit for every gene in the rows of a given matrix of counts, and will return the corresponding  $\chi^2_1$  statistics. Since calculating this for the entire gene set would take too long for building quickly

this vignette we are going to work on a subset of genes formed by human genes with documented sex-specific expression, a sampled subset of 25 human housekeeping genes and the secretin (*SCT*) gene which encodes for an endocrine hormone peptide in chromosome 11 that controls secretions in the duodenum. The gender-related and housekeeping gene lists form part from the previously loaded experimental data package `tweeDEseqCountData` and can be loaded as follows:

```
> data(genderGenes)
> data(hkGenes)
> length(XiEgenes)
```

```
[1] 63
```

```
> length(msYgenes)
```

```
[1] 32
```

```
> length(hkGenes)
```

```
[1] 669
```

The list of genes with documented sex-specific expression was built by first selecting genes in chromosome X that escape X-inactivation (Carrel and Willard, 2005) and genes in the male-specific region of the Y chromosome (Skaletsky et al., 2003), and then filtering out those that do not occur in the initial table of counts with 52580 Ensembl genes. The list of housekeeping genes was retrieved from the literature (Eisenberg and Levanon, 2003) and then also filtered to keep only those genes that form part of the initial table of counts. The selection is finally done as follows:

```
> set.seed(123)
> someHKgenes <- sample(hkGenes, size=25)
> geneSubset <- unique(c("ENSG00000070031",
+                       intersect(rownames(countsNigerianNorm),
+                                   c(someHKgenes, msYgenes, XiEgenes))))
> length(geneSubset)
```

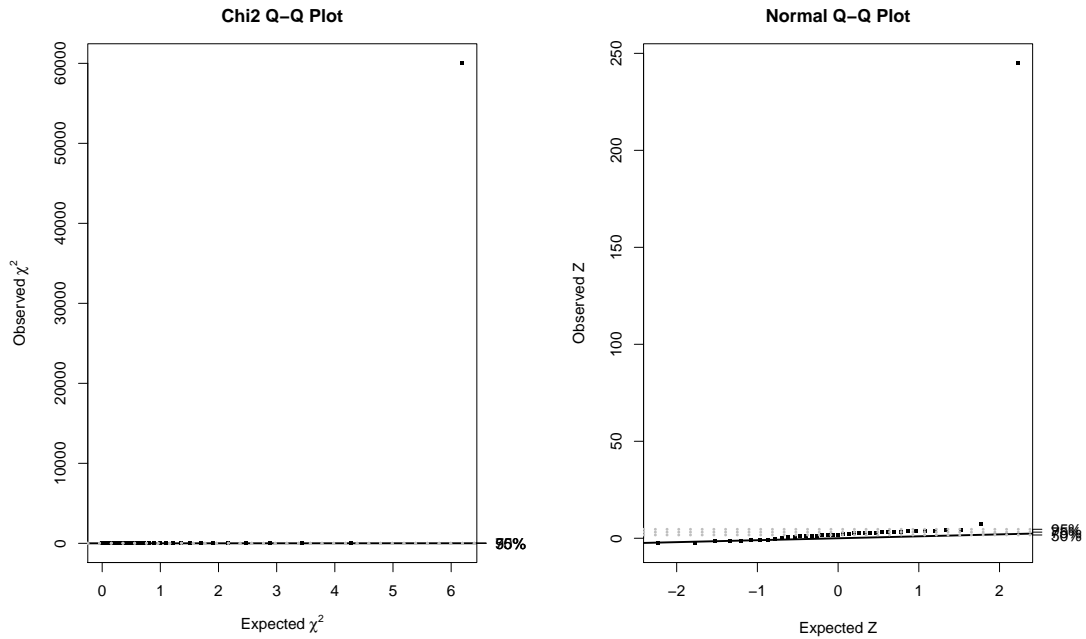
```
[1] 43
```

Now, we calculate the goodness of fit to a negative binomial distribution for each of these 43 genes using the function `gofTest()`:

```
> chi2gof <- gofTest(countsNigerianNorm[geneSubset, ], a=0)
```

and we can examine the result by means of the quantile-quantile plots produced with the function `qqchisq()` and shown in Figure 1, which indicates that about 50% of the genes show a substantial discrepancy with the respect to the negative binomial distribution.

```
> par(mfrow=c(1,2), mar=c(4, 5, 3, 4))
> qq <- qqchisq(chi2gof, main="Chi2 Q-Q Plot")
> qq <- qqchisq(chi2gof, normal=TRUE)
```



**Figure 1.** Goodness of fit to a binomial distribution. On the left a quantile-quantile plot of the  $\chi^2$  statistic employed to assess the goodness-of-fit of the RNA-seq data to a negative binomial distribution is shown. About 50% of the genes have expression profiles that depart substantially from the negative binomial distribution. On the right we have the same data but  $\chi^2$  statistics are transformed into standard normal z-scores to improve visibility of the lower quantiles.

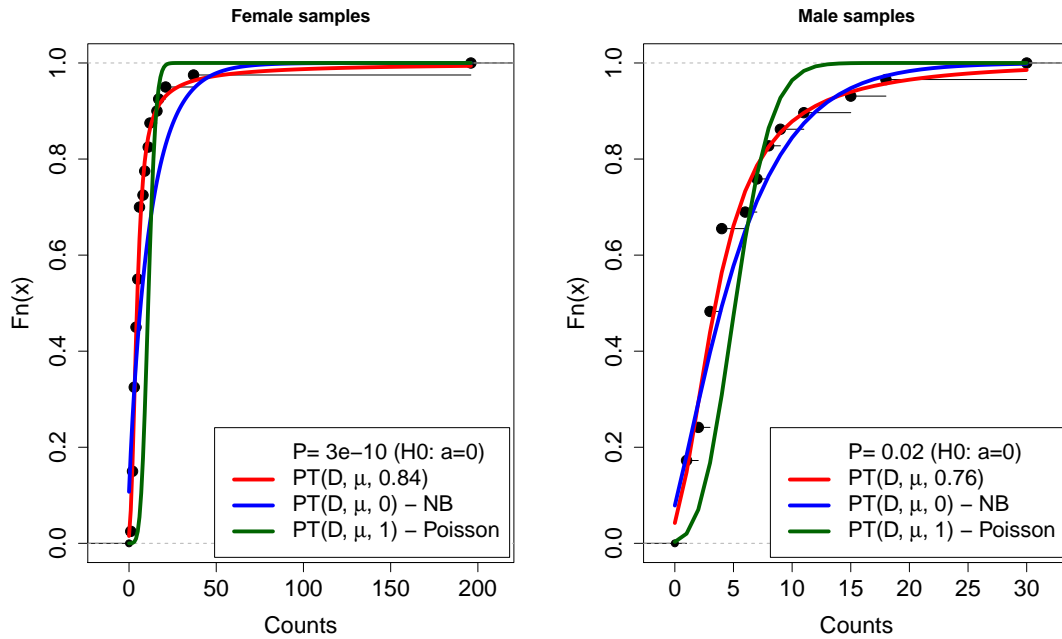
This indicates that different genes may require different count data distributions but, in fact, this can be also observed for different sample groups within the same gene. Figure 2 illustrates such a case with the secretin (*SCT*) gene (Ensembl ID ENSG00000070031) when looking separately to male and female samples. This figure is produced with the following code that calls the function `compareCountDist()` which helps in comparing an empirical distribution of counts against the Poisson, the negative binomial and the corresponding estimated PT distribution.

```
> par(mfrow=c(1,2), mar=c(4, 5, 3, 2))
> xf <- unlist(countsNigerianNorm["ENSG00000070031", genderNigerian=="female"])
> compareCountDist(xf, main="Female samples")
> xm <- unlist(countsNigerianNorm["ENSG00000070031", genderNigerian=="male"])
> compareCountDist(xm, main="Male samples")
```

What Figure 2 reflects can be also easily seen by just looking to the actual counts and identifying the female sample that produces the heavy tail on the distribution

```
> sort(xf)

NA19137 NA18855 NA19127 NA19159 NA19225 NA19257 NA18508 NA18511 NA18858 NA18912
      1      2      2      2      2      2      3      3      3      3
NA19102 NA19114 NA19116 NA18505 NA18909 NA19099 NA19140 NA19222 NA18517 NA18520
```



**Figure 2.** Empirical cumulative distribution function (CDF) of counts (black dots), calculated separately from male and female samples, for the secretin gene (*SCT*) and estimated CDF (solid lines) of Poisson-Tweedie distributions with shape parameter fixed to  $a = 1$  corresponding to a Poisson (green),  $a = 0$  corresponding to a negative binomial (blue) and with the value of  $a$  estimated from data too (red). The legend shows the values of the  $a$  parameter and the  $P$  value of the likelihood ratio test on whether the expression profile follows a negative binomial distribution ( $H_0 : a = 0$ ). We can observe that for both, male and female samples, the Poisson distribution is not adequate and that the negative binomial distribution is not adequate for female samples.

```

      3      3      3      4      4      4      4      4      5      5
NA19131 NA19143 NA18523 NA18861 NA18870 NA19108 NA19152 NA19201 NA18499 NA19093
      5      5      6      6      6      6      6      6      8      9
NA19147 NA18852 NA19209 NA18502 NA19238 NA19190 NA19204 NA18916 NA19172 NA19193
      9      11     11     12     12     16     17     21     37     196

```

```
> sort(xm)
```

```

NA18507 NA19128 NA19138 NA19160 NA19210 NA18510 NA19192 NA18498 NA18519 NA18522
      1      1      1      1      1      2      2      3      3      3
NA18856 NA18862 NA18871 NA19239 NA18486 NA18853 NA18913 NA19153 NA19171 NA19098
      3      3      3      3      4      4      4      4      4      6
NA18516 NA19101 NA18501 NA19130 NA18504 NA19200 NA19203 NA19119 NA19144
      7      7      8      8      9      11     15     18     30

```

and realize that by just removing that sample, the large overexpression in females just vanishes:

```
> xf[which.max(xf)]
```

```
NA19193
```

```
196
```

```
> 2^{log2(mean(xf))-log2(mean(xm))}
```

```
[1] 2.003402
```

```
> 2^{log2(mean(xf[-which.max(xf)])) - log2(mean(xm))}
```

```
[1] 1.192384
```

This illustrates a case in which Poisson and negative binomial distributions may be too restrictive to account for the biological variability that extensively-replicated RNA-seq experiments can reveal in count data.

## 5 Testing for differential expression

In order to illustrate the accuracy of `tweeDEseq` for detecting DE genes in a extensively-replicated RNA-seq experiment we have compared the expression profiles between males and females from the population of 69 unrelated Nigerian individuals (Pickrell et al., 2010).

The `tweeDEseq` package contains a function to test for differential expression between two different conditions using a score based test: the `tweeDE()` function. This function takes as input a matrix of counts with genes on the rows and samples on the columns.

An important feature of the `tweeDE()` function is that it allows to use multiple processors in the computing process. This is done by loading first the `multicore` package and specifying the number of cores to be used with the `mc.cores` argument in the call to `tweeDE()`.

```
> resPT <- tweeDE(countsNigerianNorm[geneSubset, ], group = genderNigerian)
```

The function `tweeDE()` returns a `data.frame` object of class `tweeDE` which can be examined with the `print()`:

```
> print(resPT)
```

```
Comparison of groups: male - female
Showing top 6 genes ranked by P-value
Minimum absolute log2 fold-change of 0
Maximum adjusted P-value of 1
```

	overallMean	female	male	log2fc	stat	pval
ENSG00000129824	186.942	3.275	440.276	7.070769	-41.385650	0.000000e+00
ENSG00000099749	31.464	3.550	69.966	4.300753	-22.924421	2.652314e-116
ENSG00000154620	16.043	2.900	34.172	3.558707	-18.401267	1.283279e-75
ENSG00000157828	7.507	1.875	15.276	3.026291	-9.993958	1.619816e-23
ENSG00000198692	18.551	3.775	38.931	3.366372	-7.769401	7.885790e-15
ENSG00000006757	152.101	196.175	91.310	-1.103291	6.767661	1.308812e-11
	pval.adjust					
ENSG00000129824	0.000000e+00					
ENSG00000099749	5.702475e-115					
ENSG00000154620	1.839366e-74					
ENSG00000157828	1.741302e-22					
ENSG00000198692	6.781780e-14					
ENSG00000006757	9.379820e-11					

which will show us by default the top 6 genes ranked by  $P$ -value including information on the magnitude of the fold-change in log2 scale (`log2fc`), overall mean expression in counts (`overallMean`), mean expression in counts for each sample group, raw  $P$ -value (`pval`) and the Benjamini-Hochberg (FDR) adjusted  $P$ -value (`pval.adjust`).

The same `print()` function allows us to call differentially expressed a subset of gene meeting cutoffs on the minimum magnitude of the fold-change and maximum FDR and store that subset in a `data.frame` object by using the appropriate arguments as follows:

```
> deGenes <- print(resPT, n=Inf, log2fc=log2(1.5), pval.adjust=0.05, print=FALSE)
> dim(deGenes)

[1] 9 7
```

We can further enrich the output with information like the symbol and description of the gene by using the annotation information stored as a `data.frame` in the experimental data package `tweeDEseqCountData` as follows:

```
> data(annotEnsembl63)
> head(annotEnsembl63)
```

	Symbol	Chr	Start	End	EntrezID	Description	Length
ENSG00000252775	U7	5	133913821	133913880	<NA>	[Source:RFAM;Acc:RF00066]	NA
ENSG00000207459	U6	5	133970529	133970635	<NA>	[Source:RFAM;Acc:RF00026]	NA
ENSG00000252899	U7	5	133997420	133997479	<NA>	[Source:RFAM;Acc:RF00066]	NA
ENSG00000201298	U6	5	134036862	134036968	<NA>	[Source:RFAM;Acc:RF00026]	NA
ENSG00000222266	U6	5	134051173	134051272	<NA>	[Source:RFAM;Acc:RF00026]	NA
ENSG00000222924	U6	5	137405044	137405147	<NA>	[Source:RFAM;Acc:RF00026]	NA

```

GCcontent
ENSG00000252775      NA
ENSG00000207459      NA
ENSG00000252899      NA
ENSG00000201298      NA
ENSG00000222266      NA
ENSG00000222924      NA

> deGenes <- merge(deGenes, annotEnsembl63, by="row.names", sort=FALSE)
```

and select certain columns to build Table 1 using the `xtable` package (code not shown but available in the source of the vignette).

## 6 Visualizing the results

An informative way to visualize the results of a differential expression analysis is by means of MA and Volcano plots, which we can easily obtain through the `tweeDEseq` package functions `MAPlot()` and `Vpplot()`, respectively as follows. Their result is shown in Figure 3.



**Table 1.** Differentially expressed genes between female and male Nigerian individuals found by tweedEseq.

Symbol	Chr	log2fc	pval.adjust	Description
RPS4Y1	Y	7.07	0.00E+00	ribosomal protein S4, Y-linked 1
CYorf15A	Y	4.30	5.70E-115	chromosome Y open reading frame 15A
TMSB4Y	Y	3.56	1.84E-74	thymosin beta 4, Y-linked
RPS4Y2	Y	3.03	1.74E-22	ribosomal protein S4, Y-linked 2
EIF1AY	Y	3.37	6.78E-14	eukaryotic translation initiation factor 1A, Y-linked
PNPLA4	X	-1.10	9.38E-11	patatin-like phospholipase domain containing 4
STS	X	-0.99	2.14E-05	steroid sulfatase (microsomal), isozyme S
HDHD1	X	-0.77	7.61E-04	haloacid dehalogenase-like hydrolase domain containing 1
UTY	Y	0.63	4.21E-02	ubiquitously transcribed tetratricopeptide repeat gene, Y-linked

```
> deGenes <- rownames(print(resPT, n=Inf, log2fc=log2(1.5), pval.adjust=0.05, print=FALSE))
> length(deGenes)

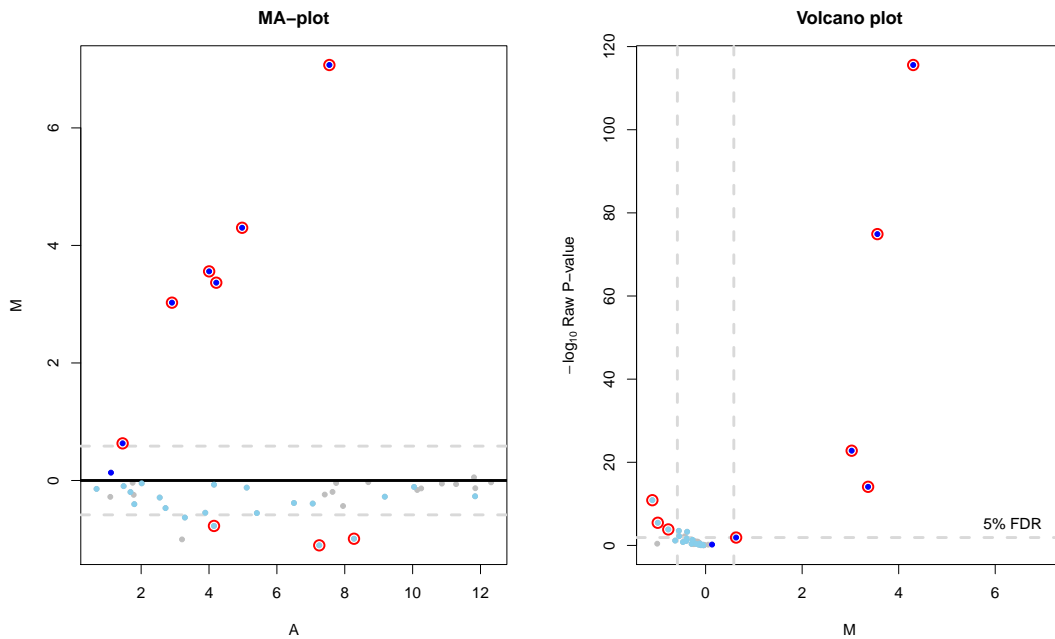
[1] 9

> hl <- list(list(genes=msYgenes, pch=21, col="blue", bg="blue", cex=0.7),
+           list(genes=XiEgenes, pch=21, col="skyblue", bg="skyblue", cex=0.7),
+           list(genes=deGenes, pch=1, col="red", lwd=2, cex=1.5)
+           )
> par(mfrow=c(1,2), mar=c(4, 5, 3, 2))
> MAplot(resPT, cex=0.7, log2fc.cutoff=log2(1.5), highlight=hl, main="MA-plot")
> Vplot(resPT, cex=0.7, highlight=hl, log2fc.cutoff=log2(1.5),
+       pval.adjust.cutoff=0.05, main="Volcano plot")
```

## 7 Assessing differential expression calling accuracy

Finally, the accuracy of the differential expression analysis illustrated here can be assessed by comparing our list of differentially expressed genes with the list of genes with documented sex-specific expression by means of a Fisher's exact test.

```
> genderGenes <- c(msYgenes[msYgenes %in% rownames(resPT)],
+                 XiEgenes[XiEgenes %in% rownames(resPT)])
> N <- nrow(resPT)
> m <- length(genderGenes)
> n <- length(deGenes)
> k <- length(intersect(deGenes, genderGenes))
> t <- array(c(k,n-k,m-k,N+k-n-m), dim=c(2,2), dimnames=list(SEX=c("in", "out"), DE=c("yes", "no")))
> t
```



**Figure 3.** Differential expression analysis for a subset of genes between male and female lymphoblastoid cell lines. On the left a MA-plot shows the magnitude of the fold-change of every gene as function of its average normalized expression level. No expression-level dependent biases can be observed in the data. On the right a volcano plot shows the raw  $P$  value of every gene for the null hypothesis of no differential expression, calculated by tweedEseq, as function of its fold-change. In both plots, red circles indicate differentially expressed genes defined by the cutoffs depicted with horizontal and vertical dashed lines. Light blue dots denote genes from the male-specific region (Skaletsky et al., 2003) of chromosome Y (MSY) and dark blue dots denote genes from Xi that escape chromosome inactivation (XiE) in female samples (Carrel and Willard, 2005).

```

DE
SEX  yes no
in   9 18
out  0 16

> fisher.test(t, alternative="greater")

Fisher's Exact Test for Count Data

data:  t
p-value = 0.008311
alternative hypothesis: true odds ratio is greater than 1
95 percent confidence interval:
 1.884678      Inf
sample estimates:
odds ratio
      Inf

```

## 8 Session info

```
> sessionInfo()

R version 2.14.2 (2012-02-29)
Platform: x86_64-unknown-linux-gnu (64-bit)

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

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

other attached packages:
[1] xtable_1.7-0          tweedEseqCountData_1.0.5 Biobase_2.14.0
[4] tweedEseq_1.0.14

loaded via a namespace (and not attached):
[1] MASS_7.3-17  edgeR_2.4.4  limma_3.10.3 tools_2.14.2
```

## References

- Carrel, L. and Willard, H. (2005). X-inactivation profile reveals extensive variability in X-linked gene expression in females. *Nature*, 434:400–404.
- Eisenberg, E. and Levanon, E. Y. (2003). Human housekeeping genes are compact. *Trends Genet*, 19(7):362–5.
- El-Shaarawi, A., Zhu, R., and Joe, H. (2011). Modelling species abundance using the Poisson-Tweedie family. *Environmetrics*, 22:152–164.
- Eснаоla, M., Puig, P., Gonzalez, D., Castelo, R., and Gonzalez, J. R. (2011). A flexible count data model to fit the wide diversity of expression profiles arising from extensively replicated rna-seq experiments. *Submitted*.
- Gupta, R. and Ong, S. (2004). A new generalization of the negative binomial distribution. *Compu Stat Data An*, 45:287–300.
- Hougaard, P., Lee, M.-L., and Whitmore, G. (1997). Analysis of overdispersed count data by mixtures of Poisson variables and Poisson processes. *Biometrics*, 53:1225–1238.
- Pickrell, J., Marioni, J., Pai, A., Degner, J., Engelhardt, B., Nkadori, E., Veyrieras, J., Stephens, M., Gilad, Y., and Pritchard, J. (2010). Understanding mechanisms underlying human gene expression variation with RNA sequencing. *Nature*, 464:768–772.
- Puig, P. and Valero, J. (2006). Count data distributions: Some characterizations with applications. *J Am Stat Assoc*, 101:332–340.

Robinson, M. and Oshlack, A. (2010). A scaling normalization method for differential expression analysis of RNA-seq data. *Genome Biol*, page 11:R25.

Skaletsky, H., Kuroda-Kawaguchi, T., Minx, P., Cordum, H., Hillier, L., Brown, L., Repping, S., Pyntikova, T., Ali, J., Bieri, T., Chinwalla, A., Delehaunty, A., Delehaunty, K., Du, H., Fewell, G., Fulton, L., Fulton, R., Graves, T., Hou, S.-F., Latrielle, P., Leonard, S., Mardis, E., Maupin, R., McPherson, J., Miner, T., Nash, W., Nguyen, C., Ozersky, P., Pepin, K., Rock, S., Rohlfing, T., Scott, K., Schultz, B., Strong, C., Tin-Wollam, A., Yang, S.-P., Waterston, R., Wilson, R., Rozen, S., and Page, D. (2003). The male-specific region of the human y chromosome is a mosaic of discrete sequence classes. *Nature*, 423:825–837.