



Max Planck Institute for Molecular Genetics
Computational Diagnostics Group @ Dept. Vingron
Innestrasse 63-73, D-14195 Berlin, Germany
<http://compdiag.molgen.mpg.de/>



Estimation of local false discovery rate
User's guide to the Bioconductor package *twilight*
1.1.0

Stefanie Scheid¹ and Rainer Spang

Max Planck Institute for Molecular Genetics
Computational Diagnostics
Department for Computation Molecular Biology
Innestr. 63-73, D-14195 Berlin, Germany

Technical Report
Nr. 2004/01

Abstract

¹Corresponding author: stefanie.scheid@molgen.mpg.de

This is the vignette of the Bioconductor compliant package *twilight*. We describe our implementation of a stochastic search algorithm to estimate the local false discovery rate. In addition, the package provides functions to test for differential gene expression in the common two-condition setting.

Contents

1	Introduction	2
2	Implemented methods	3
2.1	Function <code>twilight.pval</code> : Testing effect sizes	3
2.2	Function <code>twilight.pval</code> : Testing correlation	10
2.3	Function <code>twilight</code>	13
2.4	Function <code>twilight.combi</code>	18
3	Differences to earlier versions	21
3.1	Changes in version 1.1.0	21
3.2	Changes in version 1.0.3	22
3.3	Changes in version 1.0.2	22
4	Bibliography	23

Chapter 1

Introduction

In a typical microarray setting with gene expression data observed under two conditions, the local false discovery rate describes the probability that a gene is not differentially expressed between the two conditions given its corresponding observed score or p -value level. The resulting curve of p -values versus local false discovery rate offers an insight into the **twilight zone** between clear differential and clear non-differential gene expression. The Bioconductor compliant package *twilight* contains two main functions: Function `twilight.pval` performs a two-condition test on differences in means for a given input matrix or expression set (*exprSet*) and computes permutation based p -values. Function `twilight` performs the successive exclusion procedure described in Scheid and Spang (2004) [4] to estimate local false discovery rates and effect size distributions.

Chapter 2

Implemented methods

2.1 Function `twilight.pval`: Testing effect sizes

```
twilight.pval(xin, yin, method="fc", paired=FALSE, B=1000, yperm=NULL,
balance=FALSE, quant.ci=0.95, s0=NULL, verbose=TRUE)
```

The input object `xin` is either a pre-processed gene expression set of class *exprSet* or any data matrix where rows correspond to genes and columns to samples. Each sample was taken under one of two distinct conditions, for example under treatment A or treatment B. The functions in package *twilight* are not limited to microarray data only but can be applied to any two-sample data matrix. However, it is necessary for both expression set or numerical matrix that values are on **additive scale** like log or arsinh scale. The function does not check or transform the data to additive scale. The input vector `yin` contains condition labels of the samples. Vector `yin` has to be numeric and dichotomous. Note that in terms of *under*- and *over*-expression, the samples of the higher labeled condition are compared to the samples of the lower labeled condition.

We are given a pre-processed matrix for samples belonging to two distinct conditions A and B, and gene expression values on additive scale. For gene i in the experiment ($i = 1, \dots, N$), $\bar{\alpha}_i$ is the mean expression under condition A and $\bar{\beta}_i$ is the mean expression under condition B. To test the null hypothesis of no differential gene expression, function `twilight.pval` compares the mean expression levels $\bar{\alpha}_i$ and $\bar{\beta}_i$. The current version offers three test variants: The classical t -test uses score T_i with

$$T_i = \frac{\bar{\alpha}_i - \bar{\beta}_i}{s_i}, \tag{2.1}$$

where s_i denotes the root pooled variance. The t -test is called with `method="t"`.

The t -test score can be misleadingly high if s_i is very small. To overcome this problem,

the Z -test enlarges the denominator by a fudge factor s_0 [6], [1]:

$$Z_i = \frac{\bar{\alpha}_i - \bar{\beta}_i}{s_i + s_0}. \quad (2.2)$$

The Z -test is called with `method="z"`. Fudge factor s_0 is set to `s0=NULL` by default and is only evaluated if `method="z"`. In that case, it is the median of all root pooled variances s_1, \dots, s_N . However, the fudge factor can be chosen manually. Note that if `method="z"` is chosen with `s0=0`, the test call is altered to `method="t"`, the t -test as described above.

The third variant is based on log ratios only with score

$$F_i = \bar{\alpha}_i - \bar{\beta}_i. \quad (2.3)$$

The distribution of scores F_i under the alternative is called *effect size distribution*. With expression values on log or arsinh scale, $\exp(|F_i|)$ is an estimator for the fold change. We call $\exp(|F_i|)$ the *fold change equivalent score* [4]. Note that the package contains a function for plotting the effect size distribution which is only available if function `twilight.pval` was run with `method="fc"`, the fold change test.

Function `twilight.pval` handles paired and unpaired data. In the unpaired case (`paired=FALSE`), only one microarray was hybridized for each patient, like in a treatment and control group setting. In the paired case (`paired=TRUE`), we observed expression values of the same patient under both conditions. The typical example are before and after treatment experiments, where each patient's expression was measured twice. The input arguments `xin` and `yin` do not need to be ordered in a specific manner. It is only necessary that samples within each group have the same order, such that the first samples of the two groups represent the first pair and so on. As an example, we apply function `twilight.pval` on the leukemia data set of Golub et al. (1999) [3] as given in `library(golubEsets)`. For normalization, apply the variance-stabilizing method `vsn` in `library(vsn)` [2].

```
> data(golubMerge)
> golubNorm <- vsn(exprs(golubMerge))
> id <- as.numeric(golubMerge$ALL.AML)
```

There are 72 samples either expressing acute lymphoblastic leukemia (ALL) or acute myeloid leukemia (AML). As the AML patients are labeled with "2" and ALL with "1", we compare AML to ALL expression.

```
> golubMerge$ALL.AML
```

```
[1] ALL ALL
[17] ALL ALL ALL ALL AML AML
[33] AML AML ALL ALL
```

```
[49] ALL AML AML AML
```

```
[65] AML AML AML AML AML AML AML AML
```

```
Levels: ALL AML
```

```
> id
```

```
[1] 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 2 2 2 2 2 2 2 2 2 2 2 2
```

```
[34] 2 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 2 2 2 2 2
```

```
[67] 2 2 2 2 2 2
```

Additionally to computation of scores, empirical p -values are calculated. Argument `B` specifies the number of permutations with default set to `B=1000`. The distribution of scores under the null hypothesis is estimated by computing test scores from the same input matrix with randomly permuted class labels. These permutations are either balanced or unbalanced, with default `balance=FALSE`. The permutation options are described in detail in section 2.4. For computing empirical p -values, we count for each gene how many of *all* absolute permutation scores exceed the absolute observed score, and divide by `B`·(number of genes).

Permutation scores are also used to compute expected scores as described in Tusher et al. (2001) [6]. In addition, we compute confidence bounds for the maximum absolute difference of each set of permutation scores to expected scores. The width of the confidence bound is chosen with `quant.ci`. With default `quant.ci=0.95`, the maximum absolute difference of permutation to expected scores exceeded the confidence bound in only 5% of all permutations.

Using the optional argument `yperm`, a user-specified permutation matrix can be passed to the function. In that case, `yperm` has to be a *binary* matrix where each row is one vector of permuted class labels. The label "1" in `yperm` corresponds to the higher labeled original class. If the permutation matrix is specified, no other permutation is done and argument `B` will be ignored. Besides `set.seed`, argument `yperm` can be used to reproduce results by fixing the matrix of random permutations.

Continuing the example above, we perform a fold change test on the expression data in `golubNorm` which was transformed to `arsinh` scale by normalization with `vsu`.

```
> library(twilight)
```

```
> pval <- twilight.pval(golubNorm, id)
```

```
No complete enumeration. Prepare permutation matrix.
```

```
Compute vector of observed statistics.
```

```
Compute expected scores and p-values. This will take approx. 50 seconds.
```

```
Compute q-values.
```

```
Compute values for confidence lines.
```

The function checks whether complete enumeration of all permutations is possible. Complete enumeration is performed as long as the number of permutations does not exceed 10 000. The admissible sample sizes depend on whether the experiment was paired or unpaired and whether balanced or unbalanced permutations are used. Details are given in section 2.4. The values in the accompanying data set `expval` were computed in the same manner as in the example above.

```
> data(expval)
```

```
> expval
```

```
Twilight object with
  7129 transcripts
  observed and expected test statistics
  p- and q-values
```

```
Estimated percentage of non-induced genes:
```

```
  pi0
0.6198193
```

```
Function call:
```

```
Test: fc. Paired: FALSE. Number of permutations: 1000. Balanced: FALSE.
```

The output object of function `twilight.pval` is of class `twilight` with several elements stored in a list.

```
> class(expval)
```

```
[1] "twilight"
```

```
> names(expval)
```

```
[1] "result"  "ci.line" "quant.ci" "lambda"  "pi0"     "boot.pi0"
[7] "boot.ci" "effect"  "call"
```

The element `quant.ci` contains the corresponding input value which is passed to the plotting function. Element `ci.line` is used for plotting confidence bounds and contains the computed quantile of maximum absolute differences. The output dataframe `result` contains a matrix with several columns.

```
> names(expval$result)
```

```
[1] "observed" "expected" "candidate" "pvalue"   "qvalue"
[6] "fdr"      "mean.fdr" "lower.fdr" "upper.fdr" "index"
```

The dataframe stores observed and expected scores and corresponding empirical p -values. The genes are ordered by absolute observed test scores. Genes with observed score exceeding the confidence bounds are marked as “1” in the binary vector `result$candidate`. The output object is passed to function `plot.twilight` to produce a plot as in Tusher et al. (2001) [6] with additional confidence lines and genes marked as candidates, see Figure 2.1.

```
> expval$result[1:7, 1:5]
```

	observed	expected	candidate	pvalue
M84526_at	3.990578	1.1115133	1	1.402721e-07
M27891_at	3.669657	0.9761424	1	2.805443e-07
M89957_at	-3.153319	-1.0965293	1	4.208164e-07
X82240_rna1_at	-3.111376	-0.9599907	1	5.610885e-07
U89922_s_at	-2.954233	-0.8874557	1	7.013606e-07
M19507_at	2.925666	0.9013118	1	8.416328e-07
M11722_at	-2.689999	-0.8352241	1	9.819049e-07
	qvalue			
M84526_at	0.0006198193			
M27891_at	0.0006198193			
M89957_at	0.0006198193			
X82240_rna1_at	0.0006198193			
U89922_s_at	0.0006198193			
M19507_at	0.0006198193			
M11722_at	0.0006198193			

In addition, q -values and the estimated percentage of non-induced genes π_0 are computed as described in Remark B of Storey and Tibshirani (2003) [5]. These are stored in `result$qvalue` (see above) and `pi0`. The remaining output elements of `expval` are left free to be filled by function `twilight`. With “`qvalues`”, Figure 2.2 shows the plot of q -values against the corresponding number of rejected hypotheses.

```
> expval$pi0
```

```
[1] 0.6198193
```

Column `result$index` contains the original gene ordering of the input object. With these numbers, resorting of the `result` table is possible without knowing the original order of the row names.




Figure 2.1: Expected versus observed test scores. Deviation from the diagonal line gives evidence for differential expression. The red lines mark the 95% confidence interval on the absolute difference between observed and expected scores. The plotting call is `plot(expval,which="scores",grayscale=F,legend=F)`.




Figure 2.2: Stairplot of q -values against the resulting size of list of significant genes. A list containing all genes with $q \leq q_0$ has an estimated global false discovery rate of q_0 . The plotting call is `plot(expval, which="qvalues")`.

2.2 Function `twilight.pval`: Testing correlation

```
twilight.pval(xin, yin, method="fc", B=1000, quant.ci=0.95, ver-
bose=TRUE)
```

From version 1.1.0 on, function `twilight.pval` offers the computation of correlation scores instead of effect size scores. Now, vector `yin` can be any clinical parameter consisting of numerical values and having length equal to the number of samples. With `method="pearson"`, Pearson's coefficient of correlation to `yin` is computed for every gene in `xin`. With `method="spearman"`, `yin` and the rows of `xin` are converted into ranks and Spearman's rank correlation is computed.

Note that most input arguments of `twilight.pval` will be ignored. Only `B` takes effect and causes the computation of correlation scores based on `B` random permutations of `yin`. All successive analyses like expected scores, p - and q -values are kept as before. As an illustration, we search for genes with high correlation to the highest scoring gene found in the effect size test. Figure 2.3 displays the resulting scores.

```
> gene <- exprs(golubNorm)[expval$result$index[1], ]
> corr <- twilight.pval(golubNorm, gene, method = "spearman",
+   quant.ci = 0.99)
```

Compute vector of observed statistics.

Compute expected scores and p-values. This will take approx. 50 seconds.

Compute q-values.

Compute values for confidence lines.

```
> corr
```

```
Twilight object with
  7129 transcripts
  observed and expected test statistics
  p- and q-values
```

```
Estimated percentage of non-induced genes:
```

```
  pi0
0.6926677
```

```
Function call:
```

```
Test: spearman. Number of permutations: 1000.
```

Note that the overall percentage of non-induced genes π_0 is now interpreted as the overall percentage of genes not correlated to the clinical parameter under the null hypothesis.

```
> corr$result[1:10, 1:5]
```

	observed	expected	candidate	pvalue
M84526_at	1.0000000	0.4271397	1	1.402721e-07
X95735_at	0.7518490	0.4011408	1	2.805443e-07
M83667_rna1_s_at	0.7416554	0.3864695	1	4.208164e-07
X52056_at	0.7362531	0.3769718	1	5.610885e-07
X62654_rna1_at	0.7274744	0.3691550	1	7.013606e-07
X62320_at	0.7260274	0.3630023	1	8.416328e-07
X61587_at	0.7236478	0.3576861	1	9.819049e-07
HG3494-HT3688_at	0.7194353	0.3532224	1	1.122177e-06
L09209_s_at	0.7188565	0.3491847	1	1.262449e-06
M27891_at	0.7095633	0.3456352	1	1.402721e-06
	qvalue			
M84526_at	0.0006926677			
X95735_at	0.0006926677			
M83667_rna1_s_at	0.0006926677			
X52056_at	0.0006926677			
X62654_rna1_at	0.0006926677			
X62320_at	0.0006926677			
X61587_at	0.0006926677			
HG3494-HT3688_at	0.0006926677			
L09209_s_at	0.0006926677			
M27891_at	0.0006926677			




Figure 2.3: Expected versus observed Spearman correlation scores. Deviation from the diagonal line gives evidence for significant correlation. The red lines mark the 99% confidence interval on the absolute difference between observed and expected scores. The plotting call is `plot(corr,which="scores",grayscale=F,legend=F)`.

2.3 Function twilight

```
twilight(xin, lambda=NULL, B=0, boot.ci=0.95, clus=NULL, verbose=TRUE)
```

Local false discovery rates (fdr) are estimated from a simple mixture model given the density $f(t)$ of observed scores $T = t$:

$$f(t) = \pi_0 f_0(t) + (1 - \pi_0) f_1(t) \quad \Rightarrow \quad \text{fdr}(t) = \pi_0 \frac{f_0(t)}{f(t)}, \quad (2.4)$$

where $\pi_0 \in [0, 1]$ is the overall percentage of non-induced genes. Terms f_0 and f_1 are score densities under no induction and under induction respectively. Assume that there exists a transformation W such that $U = W(T)$ is uniformly distributed in $[0, 1]$ for all genes not differentially expressed. In a multiple testing scenario these u -values are p -values corresponding to the set of observed scores. However, we do not regard the local false discovery rate as a multiple error rate but as an exploratory tool to describe a microarray experiment over the whole range of significance, thus speaking of u -values instead of p -values.

Mapping scores to u -values allows to assume $f_0(u)$ to be the uniform density instead of specifying the null density $f_0(t)$ with respect to a chosen scoring method. The implemented successive exclusion procedure (SEP) splits any vector of u -values into a uniformly distributed null part and an alternative part. The uniform part represents genes that are not differentially expressed. Based on histogram counts of the null and the alternative part, we estimate local false discovery rates. These denote the conditional probability of being non-induced given we observe a certain u -value level. The proportion of the uniform part to the total number of genes in the experiment is a natural estimator for percentage π_0 . The successive exclusion procedure is described in detail in Scheid and Spang (2004) [4]. The functionality of `twilight` is not limited to microarray experiments. In principle, any vector of u -values can be passed to `twilight` as long as the assumption of uniformity under the null hypothesis is valid.

The objective function in `twilight` includes a penalty term that is controlled by regularization parameter $\lambda \geq 0$. The regularization ensures that we find a separation such that the uniform part contains as many u -values as possible. As percentage π_0 is often underestimated, the inclusion of a penalty term results in a more “conservative” estimate that is usually less biased. If not specified (`lambda=NULL`), function `twilight.getlambda` finds a suitable λ .

The estimates for probability π_0 and local false discovery rate are averaged over 10 runs of SEP. In addition, bootstrapping can be performed to give bootstrap estimates and bootstrap percentile confidence intervals on both π_0 and local false discovery rate. The number of bootstrap samples is set by argument `B`, and the width of the bootstrap confidence interval is set by argument `boot.ci`.

Function `twilight` takes *twilight* objects or any vector of u -values as input and returns a *twilight* object. If the input is of class *twilight*, the function works on the set of empirical p -values and fills in the remaining output elements. Note that the estimate for π_0 is replaced, and q -values are recalculated with the new estimate π_0 . As an example, we run SEP with 1000 bootstrap samples and 95% bootstrap confidence intervals: `twilight(xin=expval, B=1000, boot.ci=0.95)`, as was done for data set `exfdr`.

```
> data(exfdr)
> exfdr
```

```
Twilight object with
  7129 transcripts
  observed and expected test statistics
  p- and q-values
  local FDR
  bootstrap estimates of local FDR
```

```
Bootstrap estimate of percentage of non-induced
genes with lower and upper 95% CI:
  pi0 lower.pi0 upper.pi0
0.6282096 0.5981204 0.6575992
```

```
Function call:
Test: fc. Paired: FALSE. Number of permutations: 1000. Balanced: FALSE.
Function twilight used lambda = 0.03
```

```
> exfdr$result[1:5, 6:9]
```

	fdr	mean.fdr	lower.fdr	upper.fdr
M84526_at	0.08605050	0.1024667	0.08325147	0.1342973
M27891_at	0.08605105	0.1024672	0.08325201	0.1342977
M89957_at	0.08605159	0.1024677	0.08325255	0.1342982
X82240_rna1_at	0.08605214	0.1024682	0.08325308	0.1342986
U89922_s_at	0.08605268	0.1024687	0.08325362	0.1342990

The output elements `result$fdr`, `result$mean.fdr`, `result$lower.fdr` and `result$upper.fdr` contain the estimated local false discovery rate, the bootstrap average and upper and lower bootstrap confidence bounds. These values are used to produce the following plots which are only available after application of function `twilight`. First, we plot u -values against the corresponding conditional probabilities of being induced given the u -value level, that is $1 - \text{fdr}$, see Figure 2.4. Going back to observed scores, we produce a *volcano plot*, that is observed scores versus local false discovery rate, see Figure 2.5.




Figure 2.4: Curve of estimated local false discovery over u -values. The red lines denote the bootstrap mean (solid line) and the 95% bootstrap confidence interval on the local false discovery rate (dashed lines). The bottom ticks are 1% quantiles of u -values. The plotting call is `plot(exfdr, which="fdr", grayscale=F, legend=T)`.

Output element `effect` contains histogram information about the effect size distribution, that is log ratio under the alternative. One run of the successive exclusion procedure results in a split of the input u -value vector into a null and an alternative part. We estimate the effect size distribution from the distribution of log ratio scores corresponding to u -values in the alternative part. Again, this estimate is averaged over 10 runs of the procedure. Argument `which="effectsize"` produces the histogram of all observed log ratios overlaid with the averaged histogram of log ratios in the alternative, see Figure 2.6. The x-axis is changed to fold change equivalent scores or rather to increase in effect size. Given an observed log ratio F , the increase in effect size is $(\exp(|F|) - 1) \cdot \text{sign}(F) \cdot 100\%$. A value of 0% corresponds to no change (fold change of 1), a value of 50% to fold change 1.5 and so on. A value of -100% corresponds to a 2-fold down-regulation, that is fold change of 0.5.

The last plotting argument `which="table"` tabulates the histogram information in terms of fold change equivalent scores and log ratios.




Figure 2.5: Volcano plot of observed test scores versus local false discovery rate. The bottom ticks are 1% quantiles of observed scores. The plotting call is `plot(exfdr, which="volcano")`.

```
> tab <- plot(exfdr, which = "table")
> tab[1:8, ]
```

	LogRatio	Mixture	Alternative
-2234%	-3.15	2	2.0
-2012%	-3.05	0	0.0
-1811%	-2.95	1	0.6
-1629%	-2.85	0	0.0
-1464%	-2.75	0	0.0
-1315%	-2.65	1	1.0
-1181%	-2.55	0	0.0
-1059%	-2.45	1	0.9

The input argument `clus` of function `twilight` is used to perform parallel computation within `twilight`. Parallelizing saves computation time which is especially useful if the number of bootstrap samples `B` is large. With default `clus=NULL`, no parallelizing is done. If specified, `clus` is passed as input argument to `makeCluster` in `library(snow)`. Please make sure that `makeCluster(clus)` works properly in your environment.




Figure 2.6: Observed effect size distribution (gray histogram) overlaid with the estimated effect size distribution under the null hypothesis (black histogram). The plotting call is `plot(exfdr, which="effectsize", legend=T)`.

2.4 Function `twilight.combi`

```
twilight.combi(xin, pin, bin)
```

Function `twilight.combi` is used within `twilight.pval` to completely enumerate all permutations of a *binary* input vector `xin`. Argument `pin` specifies whether the input vector corresponds to paired or unpaired data. Argument `bin` specifies whether permutations are balanced or unbalanced. Note that the resulting permutations are always “as balanced as possible”: The balancing is done for the smaller subsample. If its sample size is odd, say 7, `twilight.combi` computes all permutations with 3 and 4 samples unchanged.

As first example, compute all unbalanced permutations of an unpaired binary vector of length 5 with two zeros and three ones. The number of rows are

$$m = \frac{5!}{2! \cdot 3!} = 10. \quad (2.5)$$

```
> x <- c(rep(0, 2), rep(1, 3))
> x

[1] 0 0 1 1 1

> twilight.combi(x, pin = FALSE, bin = FALSE)

      [,1] [,2] [,3] [,4] [,5]
[1,]    0    0    1    1    1
[2,]    0    1    0    1    1
[3,]    0    1    1    0    1
[4,]    0    1    1    1    0
[5,]    1    0    0    1    1
[6,]    1    0    1    0    1
[7,]    1    0    1    1    0
[8,]    1    1    0    0    1
[9,]    1    1    0    1    0
[10,]   1    1    1    0    0
```

Each row contains one permutation. The first row contains the input vector. In balanced permutations, we omit those rows where both original zeros have been shifted to the last three columns. The number of balanced rows is

$$m = \binom{2}{1} \cdot \frac{3!}{1! \cdot 2!} = 6. \quad (2.6)$$

```
> twilight.combi(x, pin = FALSE, bin = TRUE)
```

	[,1]	[,2]	[,3]	[,4]	[,5]
[1,]	0	0	1	1	1
[2,]	0	1	0	1	1
[3,]	0	1	1	0	1
[4,]	0	1	1	1	0
[5,]	1	0	0	1	1
[6,]	1	0	1	0	1
[7,]	1	0	1	1	0

Note that the function returns six balanced rows *and* the original input vector although it is not balanced.

Next, consider a paired input vector with four pairs. The first zero and the first one are the first pair and so on. In paired settings, values are flipped only within a pair. The number of rows is

$$m = \frac{1}{2} \cdot 2^4 = 2^3 = 8. \quad (2.7)$$

```
> y <- c(rep(0, 4), rep(1, 4))
> y
```

```
[1] 0 0 0 0 1 1 1 1
```

```
> twilight.combi(y, pin = TRUE, bin = FALSE)
```

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

The matrix above contains only half of all possible $2^4 = 16$ permutations. The reversed case `1 - twilight.combi(y, pin=TRUE, bin=FALSE)` is omitted as this will lead to the same absolute test scores as `twilight.combi(y, pin=TRUE, bin=FALSE)`. The same concept applies to balanced paired permutations. Now, two pairs are kept fixed and two pairs are flipped in each row. The number of balanced rows is

$$m = \frac{1}{2} \cdot \binom{4}{2} = 3. \quad (2.8)$$

```
> twilight.combi(y, pin = TRUE, bin = TRUE)
```

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

Again, the input vector is part of the output.

The complete enumeration of `twilight.combi` is limited by the sample sizes. The function returns `NULL` if the resulting number of rows exceeds 10000. If `NULL` is returned, function `twilight.pval` uses the functions `twilight.permute.unpair` and `twilight.permute.pair` which return a matrix of random permutations. For example, use the latter function to compute 7 balanced permutations of the paired vector `y`. Similar to `twilight.combi`, these two functions return the input vector in the first row of their output matrices.

```
> twilight.permute.pair(y, 7, bal = TRUE)
```

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

Acknowledgements

This work was done within the context of the Berlin Center for Genome Based Bioinformatics (BCB), part of the German National Genome Network (NGFN), and supported by BMBF grants 031U109C and 03U117 of the German Federal Ministry of Education.

Chapter 3

Differences to earlier versions

3.1 Changes in version 1.1.0

The computation of p -values in `twilight.pval` changed from gene-wise to pooled p -values. For the computation of a gene-wise p -value for gene i , only the permutation scores of gene i are taken into account. For pooled p -values, all permutation scores of all genes are taken as null distribution. This change has several advantages: First, gene-wise p -values were not monotonically increasing with scores because each gene had its own null distribution. Thus, two genes with almost equal scores might get quite different p -values. Now, the null distribution is the same for all genes, that is the union set of all permutation scores. Second, pooled p -values are less granular than gene-wise p -values. Gene-wise p -values are computed from B permutation scores whereas pooled p -values are computed from $B \cdot$ (number of genes) scores.

These two important features gave rise to further changes: The ordering of the `result` table is now more intuitive because the most significant genes on top have the highest scores, the lowest p - and q -values and are candidates (if there are any). In addition, the default value of the number of permutations B is lowered to 1000 permutations. Computation of pooled p -values is slower than for gene-wise p -values. On the other hand, changing to pooled p -values increases the number of values in the null distribution by the factor of number of genes. Hence, even with less permutations, the number of null values is larger than before.

We integrated Pearson and Spearman correlation coefficients into `twilight.pval`. Each gene is correlated to an numerical input vector. Expected scores are computed from random permutations of the input vector.

The `result` table contains an additional `index` column with genes indices which comes in handy for sorting back to original ordering.

All output matrices of the permutation functions `twilight.combi`, `twilight.permute.pair` and `twilight.permute.unpair` have the original labeling vector as first row. This is also

the case if balanced permutations are wanted, although the input vector is not balanced. Hence, the permutation matrix within `twilight.pval` now includes the original labeling even for balanced permutations implying that the smallest *possible* p -value is $1/(\text{number of permutations})$.

3.2 Changes in version 1.0.3

A `print.twilight` function was added which produces a short information about the contents stored in the `twilight` object.

3.3 Changes in version 1.0.2

The `which` argument of the `plot` command changed from `plot1` style to more intuitive labels like `scores` or `fdr`.

Chapter 4

Bibliography

- [1] B. Efron, R. Tibshirani, J.D. Storey and V.G. Tusher, "Empirical Bayes Analysis of a Microarray Experiment", *J. Am. Stat. Assoc.*, vol. 96, no. 456, pp. 1151-1160, 2001.
- [2] W. Huber, A. von Heydebreck, H. Sültmann, A. Poustka and M. Vingron, "Variance stabilization applied to microarray data calibration and to the quantification of differential expression", *Bioinformatics*, vol. 18, suppl. 1, pp. S96-S104, 2002.
- [3] T.R. Golub, D.K. Slonim, P. Tamayo, C. Huard, M. Gaasenbeek, J.P. Mesirov, H. Coller, M.L. Loh, J.R. Downing, M.A. Caligiuri, C.D. Bloomfield and E.S. Lander, "Molecular Classification of Cancer: Class Discovery and Class Prediction by Gene Expression Monitoring", *Science*, vol. 286, pp. 531-537, 1999.
- [4] S. Scheid and R. Spang, "A stochastic downhill search algorithm for estimating the local false discovery rate", *IEEE Transactions on Computational Biology and Bioinformatics*, vol. 1, no. 3, pp. 98-108, 2004.
- [5] J.D. Storey and R. Tibshirani, "Statistical significance for genomewide studies", *Proc. Natl. Acad. Sci.*, vol. 100, no. 16, pp. 9440-9445, 2003.
- [6] V. Tusher, R. Tibshirani and C. Chu, "Significance analysis of microarrays applied to ionizing radiation response", *Proc. Natl. Acad. Sci.*, vol. 98, no. 9, pp. 5116-5121, 2001.