Test association between phenotype and gene expression

Evarist Planet

Bioinformatics & Biostatistics Unit

IRB Barcelona

Contents

1 Introduction

Imagine a situation where we have gene expression and phenotype variables and we want to test the association of each gene with phenotype. We would probably be interested in testing association of groups of genes (or gene sets) with phenotype. This library provides the tools to do both things in a way that is efficient, structured, fast and scalable. We also provide tools to do *GSEA* (Gene set enrichment analysis) of all phenotype variables at once.

The functions and methods presented on this `vignette` provide tools to easily test association between gene expression levels of individual genes or gene sets of genes and the selected phenotypes of a given gene expression dataset. These can be particularly useful for datasets arising from RNAseq or microarray gene expression studies.

We will load the `ExpressionSet` of a cohort (*GSE2034*) we downloaded from GEO.

```
> options(width = 100)
> library(phenoTest)
> data(eset)
> eset


ExpressionSet (storageMode: lockedEnvironment)
assayData: 1000 features, 100 samples
  element names: exprs
protocolData: none
phenoData
  sampleNames: GSM36793 GSM36796 ... GSM36924 (100 total)
  varLabels: PID GEOaccession ... BrainRelapse (7 total)
  varMetadata: labelDescription
featureData: none
experimentData: use 'experimentData(object)'
Annotation: hgu133a
```

For illustration purposes we selected the first 1000 genes and the first
100 samples, created a continuous variable (named *Tumor.size*) and added a
new category to the categorical variable *lymph.node.status* to illustrate the
functionality of the package.

```
> Tumor.size <- rnorm(ncol(eset), 50, 2)
> pData(eset) <- cbind(pData(eset), Tumor.size)
> pData(eset)[1:20, "lymph.node.status"] <- "positive"
```

2 Individual gene(s) association with phenotype(s)

2.1 Creating an `epheno`

The `epheno` object will contain the univariate association between a list of
phenotype variables and the gene expression from the given `Expression-
Set`. We will use the `ExpressionPhenoTest` function to create the `epheno`
object. We will have to tell this function which phenotype variables we want
to test and the type of these variables (if they are *ordinal*, *continuous*, *cat-
egorical* or *survival* variables). For this purpose we will create a variable
called (for instance) `vars2test`. This variable has to be of class `list` with
components *continuous*, *categorical*, *ordinal* and *survival* indicating which
phenotype variables should be tested. *continuous*, *categorical* and *ordinal*
must be character vectors, *survival* a matrix with columns named *time* and
event. The names must match names in `names(pData(eset))` (being `eset`
the `ExpressionSet` of the cohort we are interested in).

```
> head(pData(eset))
```

```
          PID GEOaccession lymph.node.status Months2Relapse Relapse ER.Status BrainRelapse
GSM36793   3     GSM36793          positive            101       0         0            0
GSM36796   5     GSM36796          positive            118       0         1            0
GSM36797   6     GSM36797          positive              9       1         0            0
GSM36798   7     GSM36798          positive            106       0         0            0
GSM36800   8     GSM36800          positive             37       1         0            0
GSM36801   9     GSM36801          positive            125       0         1            0
          Tumor.size
GSM36793   50.18250
GSM36796   47.91716
GSM36797   50.81062
GSM36798   47.30843
GSM36800   50.45073
GSM36801   51.36586

> survival <- matrix(c("Relapse", "Months2Relapse"), ncol = 2, byrow = TRUE)
> colnames(survival) <- c("event", "time")
> vars2test <- list(survival = survival, categorical = "lymph.node.status", continuous = "Tumor.size")
> vars2test


$survival
     event      time
[1,] "Relapse" "Months2Relapse"

$categorical
[1] "lymph.node.status"

$continuous
[1] "Tumor.size"
```

Now we have everything we need to create the `epheno` object:

```
> epheno <- ExpressionPhenoTest(eset, vars2test, p.adjust.method = "none")

Performing analysis for continuous variable Tumor.size
Performing analysis for categorical variable lymph.node.status
Performing analysis for survival variable Relapse

> epheno

Object of class 'epheno'
featureNames: 1007_s_at, 1053_at, 117_at ... (1000) feature(s)
phenoNames: Tumor.size, lymph.node.status, Relapse. (3) phenotype(s)
P-value adjustment method:  none
Annotation: hgu133a

Type "showMethods(classes='epheno')" for a list of ALL methods
```

P values can also be adjusted afterwards:

```
> p.adjust.method(epheno)

[1] "none"

> epheno <- pAdjust(epheno, method = "BH")
> p.adjust.method(epheno)

[1] "BH"
```

The `epheno` object extends the `ExpressionSet` object and therefore methods that are available for `ExpressionSet`s are also available for `epheno`s.

The effect of both *continuous*, *categorical* and *ordinal* phenotype variables on gene expression levels are tested via `lmFit` from package `limma` (Smyth [2005]). For *ordinal* variables a single coefficient is used to test its effect on gene expression (trend test), which is then used to obtain a P-value. Gene expression effects on *survival* are tested via Cox proportional hazards model (Cox [1972]), as implemented in function `coxph` from package `survival`.

`ExpressionPhenoTest` implements parallel computing via the function `mclapply` from the package `multicore`. Currently `multicore` only operates on Unix systems. If you are a windows user you should set `mc.cores`=1 (the default).

2.2   Useful methods for the `epheno` object

Some of the methods for the epheno objects are shown here.

The object can be subseted by phenotype names:

```
> phenoNames(epheno)

[1] "Tumor.size"       "lymph.node.status" "Relapse"

> epheno[, "Tumor.size"]


Object of class 'epheno'
featureNames: 1007_s_at, 1053_at, 117_at ... (1000) feature(s)
phenoNames: Tumor.size. (1) phenotype(s)
P-value adjustment method:  BH
Annotation: hgu133a

Type "showMethods(classes='epheno')" for a list of ALL methods


> epheno[, 2]


Object of class 'epheno'
featureNames: 1007_s_at, 1053_at, 117_at ... (1000) feature(s)
phenoNames: lymph.node.status. (1) phenotype(s)
P-value adjustment method:  BH
Annotation: hgu133a

Type "showMethods(classes='epheno')" for a list of ALL methods


> epheno[, c(FALSE, FALSE, TRUE)]


Object of class 'epheno'
featureNames: 1007_s_at, 1053_at, 117_at ... (1000) feature(s)
phenoNames: Relapse. (1) phenotype(s)
P-value adjustment method:  BH
Annotation: hgu133a

Type "showMethods(classes='epheno')" for a list of ALL methods
```

or by class (class can be ordinal, continuous, categorical or survival):

```
> phenoClass(epheno)

      Tumor.size lymph.node.status            Relapse
    "continuous"      "categorical"        "survival"

> epheno[, phenoClass(epheno) == "survival"]

Object of class 'epheno'
featureNames: 1007_s_at, 1053_at, 117_at ... (1000) feature(s)
phenoNames: Relapse. (1) phenotype(s)
P-value adjustment method:  BH
Annotation: hgu133a

Type "showMethods(classes='epheno')" for a list of ALL methods
```

epheno objects contain information summarizing the association between genes and phenotypes. getMeans can be used to obtain the average expression for each group in categorical and ordinal variables, as well as for categorized version of the continuous variables.

```
> head(getMeans(epheno))

         Tumor.size.[45.5,49.1) Tumor.size.[49.1,50.6) Tumor.size.[50.6,55.6]
1007_s_at             11.964886             11.753518             11.668944
1053_at                7.716940              7.812752              7.692608
117_at                 7.876629              7.918348              7.902192
121_at                10.300963             10.230651             10.166576
1255_g_at              6.215478              5.924199              5.551650
1294_at                9.394689              9.519671              9.423615
         lymph.node.status.negative lymph.node.status.positive
1007_s_at                 11.779693                  11.868594
1053_at                    7.700153                   7.902028
117_at                     7.882203                   7.965347
121_at                    10.217469                  10.297187
1255_g_at                  5.856978                   6.073553
1294_at                    9.444534                   9.449257
```

Here we see that tumor size has been categorized into 3 groups. The number of categories can be changed with the argument continuousCategories in the call to ExpressionPhenoTest.

epheno objects also contain fold changes and hazard ratios (for survival variables). These can be accessed with getSummaryDif, getFc and getHr.

```
> head(getSummaryDif(epheno))

         Tumor.size.fc.[49.1,50.6) Tumor.size.fc.[50.6,55.6] lymph.node.status.positive.fc
1007_s_at                -1.157785                 -1.227687                      1.063559
1053_at                   1.068667                 -1.017009                      1.150192
117_at                    1.029339                  1.017876                      1.059324
121_at                   -1.049944                 -1.097626                      1.056812
1255_g_at                -1.223725                 -1.584280                      1.161972
1294_at                   1.090494                  1.020252                      1.003280
         Relapse.HR
```

```
1007_s_at  -1.299526
1053_at    -1.138198
117_at      1.175962
121_at      1.100616
1255_g_at   1.093210
1294_at     1.006512

> head(getFc(epheno))

        Tumor.size.fc.[49.1,50.6) Tumor.size.fc.[50.6,55.6] lymph.node.status.positive.fc
1007_s_at              -1.157785                 -1.227687                      1.063559
1053_at                 1.068667                 -1.017009                      1.150192
117_at                  1.029339                  1.017876                      1.059324
121_at                 -1.049944                 -1.097626                      1.056812
1255_g_at              -1.223725                 -1.584280                      1.161972
1294_at                 1.090494                  1.020252                      1.003280

> head(getHr(epheno))

        Relapse.HR
1007_s_at  -1.299526
1053_at    -1.138198
117_at      1.175962
121_at      1.100616
1255_g_at   1.093210
1294_at     1.006512
```

ExpressionPhenoTest also computes P-values. eBayes from limma package is used for continuous, categorical and ordinal phenotypes. A Cox proportional hazards likelihood-ratio test is used for survival phenotypes. P-values can be accessed with getPvals. Notice that a single P-value is reported for each phenotype variable. For categorical variables these corresponds to the overall null hypothesis that there are no differences between groups.

```
> head(getPvals(epheno))

        Tumor.size lymph.node.status   Relapse
1007_s_at  0.8297632         0.9576191 0.9197093
1053_at    0.9589053         0.9197093 0.9388141
117_at     0.9794646         0.9670237 0.9220701
121_at     0.9197093         0.9463586 0.9650314
1255_g_at  0.8965376         0.9310620 0.9670237
1294_at    0.9833547         0.9970126 0.9970126
```

We can also ask for the variables we sent to the ExpressionPhenoTest function:

```
> getVars2test(epheno)

$continuous
[1] "Tumor.size"

$categorical
[1] "lymph.node.status"

$survival
     event       time
[1,] "Relapse" "Months2Relapse"
```

2.3 Export an `epheno`

Functions `export2csv` and `epheno2html` can be used to export to a comma separated value (csv) or an html file. The html file will have useful links to online databases that will provide information about each known gene. For more information about how to use these functions and examples read their help manuals.

3 Gene set(s) association with phenotype(s)

Gene sets can be stored in a list object. Each element of the list will contain one gene set. The names of the list will be the names of the gene sets. Here we select genes at random to build our gene sets:

```
> sign1 <- featureNames(eset)[100:160]
> sign2 <- featureNames(eset)[300:320]
> mySignature <- list(sign1, sign2)
> names(mySignature) <- c("My first signature", "Another signature")
> mySignature

$`My first signature`
 [1] "200079_s_at" "200080_s_at" "200081_s_at" "200082_s_at" "200083_at"   "200084_at"
 [7] "200085_s_at" "200086_s_at" "200087_s_at" "200088_x_at" "200089_s_at" "200090_at"
[13] "200091_s_at" "200092_s_at" "200093_s_at" "200094_s_at" "200095_x_at" "200096_s_at"
[19] "200097_s_at" "200098_s_at" "200099_s_at" "200593_s_at" "200594_x_at" "200595_s_at"
[25] "200596_s_at" "200597_at"   "200598_s_at" "200599_s_at" "200600_at"   "200601_at"
[31] "200602_at"   "200603_at"   "200604_s_at" "200605_s_at" "200606_at"   "200607_s_at"
[37] "200608_s_at" "200609_s_at" "200610_s_at" "200611_s_at" "200612_s_at" "200613_at"
[43] "200614_at"   "200615_s_at" "200616_s_at" "200617_at"   "200618_at"   "200619_at"
[49] "200620_at"   "200621_at"   "200622_x_at" "200623_s_at" "200624_s_at" "200625_s_at"
[55] "200626_s_at" "200627_at"   "200628_s_at" "200629_at"   "200630_x_at" "200631_s_at"
[61] "200632_s_at"

$`Another signature`
 [1] "200772_x_at" "200773_x_at" "200774_at"   "200775_s_at" "200776_s_at" "200777_s_at"
 [7] "200778_s_at" "200779_at"   "200780_x_at" "200781_s_at" "200782_at"   "200783_s_at"
[13] "200784_s_at" "200785_s_at" "200786_at"   "200787_s_at" "200788_s_at" "200789_at"
[19] "200790_at"   "200791_s_at" "200792_at"
```

Gene sets can also be stored in gene set collection objects. From here on all functions have methods for gene sets stored as `list`s, `GeneSet`s or `GeneSetCollection`s. You can use the one you feel more confortable with. We will work with `GeneSetCollection`:

```
> library(GSEABase)
> myGeneSetA <- GeneSet(geneIds = sign1, setName = "My first signature")
> myGeneSetB <- GeneSet(geneIds = sign2, setName = "Another signature")
> mySignature <- GeneSetCollection(myGeneSetA, myGeneSetB)
> mySignature
```
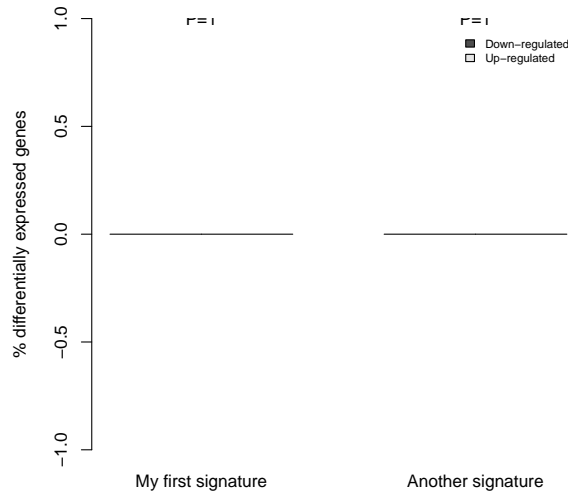
Figure 1: `barplotSignifSignatures`: Number of diferentially expressed genes in each gene set that are statistically significant. P-values test for differences in each signature between the number of up and down regulated genes.

```
GeneSetCollection
  names: My first signature, Another signature (2 total)
  unique identifiers: 200079_s_at, 200080_s_at, ..., 200792_at (82 total)
  types in collection:
    geneIdType: NullIdentifier (1 total)
    collectionType: NullCollection (1 total)
```

3.1 Plots that use epheno as input

`barplotSignifSignatures` will plot the percentage of up regulated and down regulated genes that are statistically significant in each signature. In our random selection of genes we did not find any statistically significant genes. Therefore, and just to show the plot we set the alpha value 0.85. The plot can be seen in Figure 1.

```
> barplotSignifSignatures(epheno[, "lymph.node.status"], mySignature, alpha = 0.85)
```

By default `barplotSignifSignatures` performs a binomial test (`binom.test` from package `stats`) for each signature to test if the proportions of up regulated and down regulated genes are different. For example, Figure 1 indicates
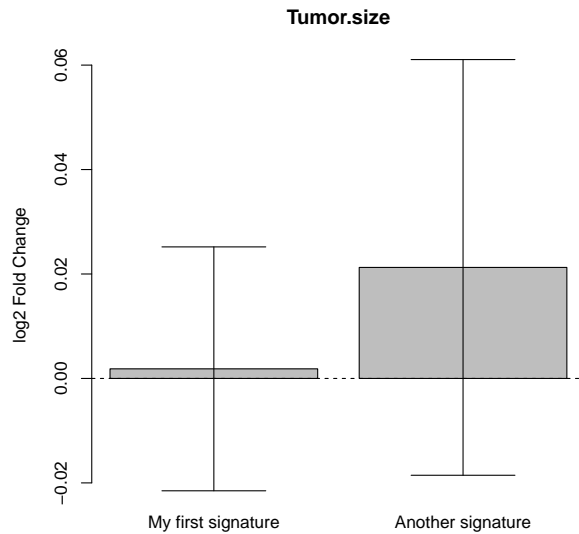
Figure 2: `barplotSignatures`: Averge fold change or hazard ratio.

that in the first signature the proportion of up regulated genes is higher than the proportion of down regulated genes. The second signature shows no significant statistical differences.

Sometimes we want to compare the proportions of up and down regulated genes in our signature with the proportions of up and down regulated of all genes in the genome. In this case we may provide a reference signature via the argument `referenceSignature`. When providing the `referenceSignature` argument a chi-square test comparing the proportion of up and down regulated genes in each signature with the proportion in the reference set will be computed.

When a reference gene set is provided and parameter `testUpDown` is `TRUE` (by default it is `FALSE`) the proportion of up regulated genes is compared with those of the reference gene set. The same is done for down regulated genes.

`barplotSignatures` plots the average log2 fold change or hazard ratio of each phenotype for each gene set. Figure 2 shows an example of it.

```
> barplotSignatures(epheno[, "Tumor.size"], mySignature)
```

We can also cluster our samples in two clusters based on the expression levels of one gene set of genes and then test the effect of cluster on phenotypes. For *ordinal* and *continuous* variables a Kruskal-Wallis Rank Sum
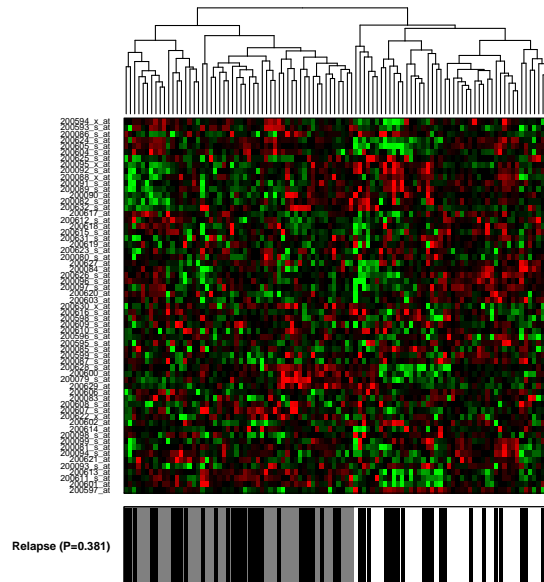
Figure 3: Heatmap produced with `heatmapPhenoTest` function. All variables in `vars2test` that are of class *logical* will be plotted under the heatmap.

test is used, for *categorical* variables a chi-square test is used and for *survival* variables a Cox proportional hazards likelihood-ratio test is used. The `heatmapPhenoTest` function can be used to this end. Its results can be seen in Figure 3 and 4.

```
> pvals <- heatmapPhenoTest(eset, mySignature[[1]], vars2test = vars2test, heat.kaplan = "heat")

> pvals

       Tumor.size lymph.node.status     Months2Relapse
     "(P=0.6883)"      "(P=0.4508)"       "(P=0.381)"


> pvals <- heatmapPhenoTest(eset, mySignature[[1]], vars2test = vars2test, heat.kaplan = "kaplan")

> pvals

       Tumor.size lymph.node.status     Months2Relapse
     "(P=0.6883)"      "(P=0.4531)"       "(P=0.381)"
```
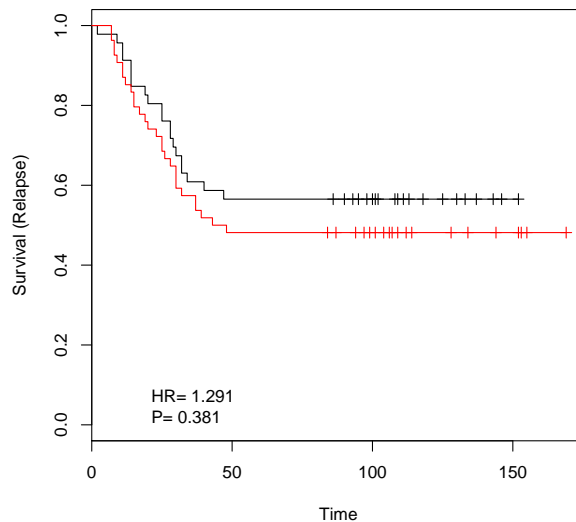
Figure 4: Kaplan-Meier produced with `heatmapPhenoTest` function.

3.2 GSEA (Gene Set Enrichment Analysis)

A popular way to test association between gene sets' gene expression and phenotype is GSEA (Subramanian [2005]). The main idea is to test the association between the gene set *as a whole* and a phenotype.

Although GSEA and several extensions are already available in other *Biconductor* packages, here we implement a slightly different extension. Most GSEA-like approaches assess statistical significance by permuting the values of the phenotype of interest. From a statisticall point of view this tests the null hypothesis that no genes are associated with phenotype. However in many applications one is actually interested in testing if the proportion of genes associated with phenotype in the gene set is greater than that outside of the gene set. As a simple example, imagine a cancer study where 25% of the genes are differentially expressed. In this setup a randomly chosen gene set will have around 25% of differentially expressed genes, and classical GSEA-like approaches will tend to flag the gene set as statistically significant. In contrast, our implementation will tend to select only gene sets with more than 25% of differentially expressed genes.

On a first step we will use the `gseaSignatures` method to compute *enrichment scores* (see Subramanian [2005] for details about the enrichment scores) and *simulated enrichment scores* (by permuting the selection of

11

genes). The *simulated enrichment scores* will be used to compute P-values and FDR by the `gseaSignificance` method. We can summarize the results obtained using the `summary` method. The following chunk of code is an illustrative example of it:

```
> set.seed("777")
> my.gseaSignatures <- gseaSignatures(x = epheno, signatures = mySignature, B = 1000)
> my.gseaSignatures


Object of class 'gseaSignaturesVar'
You can use the getEs, getEsSim and getFcHr methods to easily acces its data
-The tested variables are: Tumor.size.fc.[49.1,50.6) Tumor.size.fc.[50.6,55.6] lymph.node.status.positive.fc Relapse.
-The tested signatures (for each variable) are: My first signature Another signature

> my.gseaSignificance <- gseaSignificance(my.gseaSignatures, p.adjust = "BH")
> my.gseaSignificance


Object of class 'gseaSignificanceVar'
You can use the summary method to easily acces its data
-The tested variables are: Tumor.size.fc.[49.1,50.6) Tumor.size.fc.[50.6,55.6] lymph.node.status.positive.fc Relapse.
-The tested signatures (for each variable) are: My first signature Another signature
-P adjust method: BH

> my.summary <- summary(my.gseaSignificance)
> my.summary


                          variable           signature         es        nes   pval.es  pval.nes
1       Tumor.size.fc.[49.1,50.6) My first signature -0.2171991 -0.8097059 0.8099268 0.7822591
2       Tumor.size.fc.[49.1,50.6)  Another signature  0.3281258  0.9930753 0.8099268 0.7822591
3       Tumor.size.fc.[50.6,55.6] My first signature  0.2541469  1.0283811 0.8699020 0.8721015
4       Tumor.size.fc.[50.6,55.6]  Another signature  0.2105614  0.6755951 0.8922045 0.9229477
5 lymph.node.status.positive.fc My first signature  0.3640407  1.3338321 0.1525067 0.2191314
6 lymph.node.status.positive.fc  Another signature -0.1928397 -0.5826575 0.9657338 0.9789755
7                     Relapse.HR My first signature -0.3266072 -1.3023822 0.2012930 0.2430182
8                     Relapse.HR  Another signature -0.2502964 -0.7698751 0.7811233 0.8281786
        fdr
1 0.7822591
2 0.4756465
3 0.8721015
4 0.9229477
5 0.1095657
6 0.9789755
7 0.2430182
8 0.8281786
```

We can produce plots as follows:

```
> plot(my.gseaSignatures, my.gseaSignificance)
```

This will produce two plots (one for *enrichment score* and another for *normalised enrichment score*) for every phenotype and gene set (in our case 12 plots). Following code shows an example on plotting only *enrichment score* for variable *Relapse* on the first gene set of genes. Plot can be seen in Figure 5.
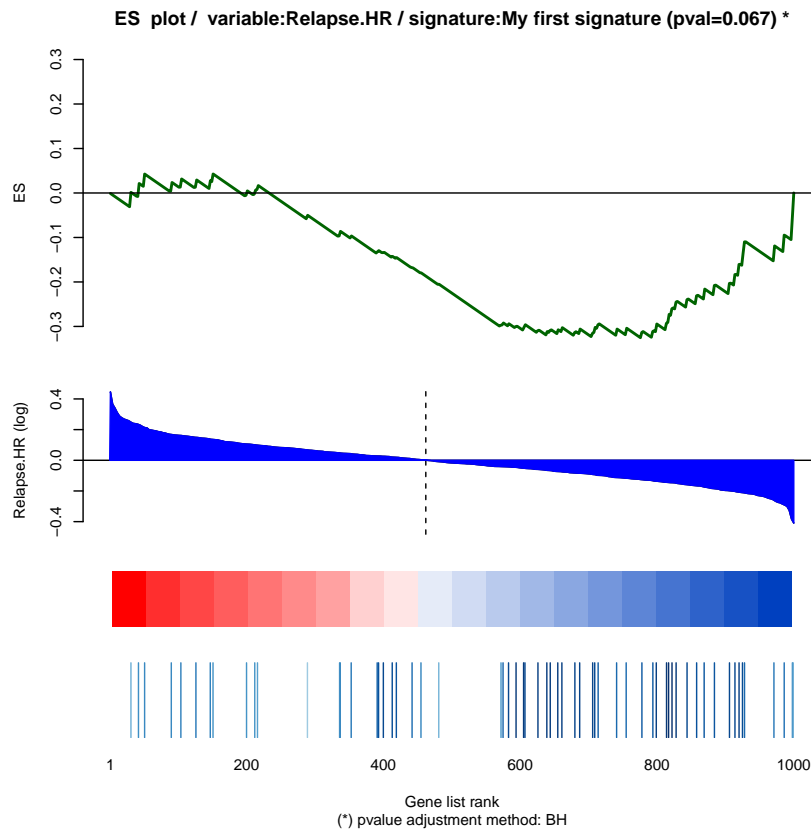
Figure 5: GSEA plot.

```
> my.gseaSignatures <- gseaSignatures(x = epheno[, "Relapse"], signatures = mySignature[1],
+     B = 100)
> my.gseaSignificance <- gseaSignificance(my.gseaSignatures, p.adjust = "BH")
> summary(my.gseaSignificance)

    variable      signature       es       nes      pval.es    pval.nes        fdr
1 Relapse.HR My first signature -0.3266072 -1.377891 0.06745943 0.06745616 0.06745616

> plot(my.gseaSignatures, my.gseaSignificance, es.nes = "es")
```

gseaSignatures can be used not only with epheno objects but also with
objects of class numeric or matrix. For more information read the gseaSig-
natures function help.

Following similar ideas to Virtaneva [2001] we also implemented a Wilcoxon
test. This can be used instead of the permutation test which can be slow if
we use a lot of permutations and we can not use the multicore package. The
plot we will obtain will also be different. Instead of plotting the *enrichment
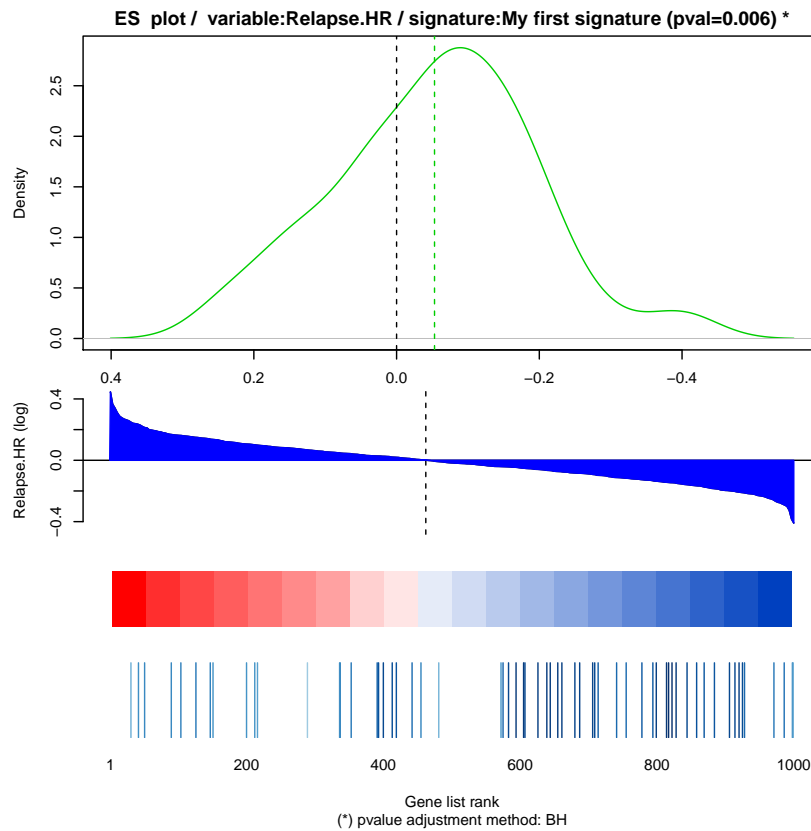scores* we will plot the density function and the mean log2 fold change or

Figure 6: GSEA plot using Wilcoxon test.

hazard ratio of the genes that belong to our gene set. This will allow us to compare how similar/different from 0 the mean of our gene set is. The plot using Wilcoxon test can be seen in Figure 6.

```
> my.gseaSignatures <- gseaSignatures(x = epheno[, "Relapse"], signatures = mySignature[1],
+    B = 100, test = "wilcox")
> my.gseaSignificance <- gseaSignificance(my.gseaSignatures, p.adjust = "BH")
> summary(my.gseaSignificance)

    variable        signature           es         pval
1 Relapse.HR My first signature -0.04643351 0.005812314

> plot(my.gseaSignatures, my.gseaSignificance)
```

Notice that using a Wilcoxon test is conceptually very similar to the average gene set fold change presented in figure 2.

A current limitation of gseaSignatures is that it does not consider the existance of dependence between genes in the gene set. This will be addressed

in future versions. Nevertheless we believe `gseaSignatures` is usefull in that it targets the correct null hypothesis that gene set is as enriched as a randomly selected gene set, opposed to testing that there are no enriched genes in the set as is done in GSEA.

References

D.R. Cox. Regression models and life tables. *Journal of the Royal Statistical Society Series B*, 34:187–220, 1972.

G.K. Smyth. Limma: linear models for microarray data. In R. Gentleman, V. Carey, S. Dudoit, R. Irizarry, and W. Huber, editors, *Bioinformatics and Computational Biology Solutions using R and Bioconductor*, pages 397–420. Springer, New York, 2005.

Aravind Subramanian. Gene set enrichment analysis: A knowledge-based approach for interpreting genome-wide expression profiles. *PNAS*, 102, 2005.

K. Virtaneva. Expression profiling reveals fundamental biological differences in acute myeloid leukemia with isolated trisomy 8 and normal cytogenetics. *Proc Natl Acad Sci U S A*, 98(98):1124–1129, January 2001. doi: http://dx.doi.org/10.1073/pnas.98.3.1124. URL `http://dx.doi.org/10.1073/pnas.98.3.1124`.