missMethyl: Analysing Illumina HumanMethylation450 BeadChip Data

Belinda Phipson and Jovana Maksimovic

Modified: 10 September 2015. Compiled: May 15, 2016

Contents

1   Introduction

The *missMethyl* package contains functions to analyse methylation data from Illuminaś HumanMethylation450 beadchip. These arrays are a cost-effective alternative to whole genome bisulphite sequencing, and as such are widely used to profile DNA methylation. Specifically, missMethyl contains functions to perform SWAN normalisation [9], perform differential methylation analysis using RUV [14], differential variability analysis [10] and gene set analysis [15]. As our lab's research into specialised analyses of these arrays continues we anticipate that the package will be continuously updated with new functions.

Raw data files are in IDAT format, which can be read into R using the *minfi* package [1]. Statistical analyses are usually performed on M-values, and β values are used for visualisation, both of which can be extracted from *MethylSet* objects, which is a class of object created by *minfi*. For detecting differentially variable CpGs we recommend that the analysis is performed on M-values. All analyses described here are performed at the CpG site level.

2 Reading data into R

We will use the data in the *minfiData* package to demonstrate the functions in *missMethyl*. The example dataset has 6 samples across two slides. The sample information is in the targets file. An essential column in the targets file is the "Basename" column which tells *minfi* where the idat files to be read in are located. The R commands to read in the data are taken from the *minfi* User's Guide. For additional details on how to read the IDAT files into R, as well as information regarding quality control please refer to the *minfi* User's Guide.

```
library(missMethyl)

## Setting options('download.file.method.GEOquery'='auto')

## Setting options('GEOquery.inmemory.gpl'=FALSE)

##

library(limma)
library(minfi)

## Loading required package:  BiocGenerics

## Loading required package:  parallel

##
## Attaching package:  'BiocGenerics'

## The following objects are masked from 'package:parallel':
##
##     clusterApply, clusterApplyLB, clusterCall, clusterEvalQ, clusterExport,
##     clusterMap, parApply, parCapply, parLapply, parLapplyLB, parRapply,
##     parSapply, parSapplyLB

## The following object is masked from 'package:limma':
##
##     plotMA

## The following objects are masked from 'package:stats':
##
##     IQR, mad, xtabs

## The following objects are masked from 'package:base':
##
##     Filter, Find, Map, Position, Reduce, anyDuplicated, append, as.data.frame,
##     cbind, colnames, do.call, duplicated, eval, evalq, get, grep, grepl,
##     intersect, is.unsorted, lapply, lengths, mapply, match, mget, order, paste,
##     pmax, pmax.int, pmin, pmin.int, rank, rbind, rownames, sapply, setdiff,
##     sort, table, tapply, union, unique, unsplit

## Loading required package:  Biobase

## Welcome to Bioconductor
##
##     Vignettes contain introductory material; view with 'browseVignettes()'.  To
##     cite Bioconductor, see 'citation("Biobase")', and for packages
##     'citation("pkgname")'.

## Loading required package:  lattice

## Loading required package:  GenomicRanges

## Loading required package:  S4Vectors
```

```
## Loading required package:  stats4

##
## Attaching package:  'S4Vectors'

## The following objects are masked from 'package:base':
##
##     colMeans, colSums, expand.grid, rowMeans, rowSums

## Loading required package:  IRanges

## Loading required package:  GenomeInfoDb

## Loading required package:  SummarizedExperiment

## Loading required package:  Biostrings

## Loading required package:  XVector

## Loading required package:  bumphunter

## Loading required package:  foreach

## Loading required package:  iterators

## Loading required package:  locfit

## locfit 1.5-9.1   2013-03-22
```

```r
library(minfiData)
```

```
## Loading required package:  IlluminaHumanMethylation450kmanifest

## Loading required package:  IlluminaHumanMethylation450kanno.ilmn12.hg19
```

```r
baseDir <- system.file("extdata", package = "minfiData")
targets <- read.450k.sheet(baseDir)
```

```
## Warning:  'read.450k.sheet' is deprecated.
## Use 'read.metharray.sheet' instead.
## See help("Deprecated")
```

```
## [read.450k.sheet] Found the following CSV files:
## [1] "/home/biocbuild/bbs-3.3-bioc/R/library/minfiData/extdata/SampleSheet.csv"
```

```r
targets[,1:9]
```

```
##   Sample_Name Sample_Well Sample_Plate Sample_Group Pool_ID person age sex status
## 1    GroupA_3          H5           NA       GroupA      NA     id3  83   M normal
## 2    GroupA_2          D5           NA       GroupA      NA     id2  58   F normal
## 3    GroupB_3          C6           NA       GroupB      NA     id3  83   M cancer
## 4    GroupB_1          F7           NA       GroupB      NA     id1  75   F cancer
## 5    GroupA_1          G7           NA       GroupA      NA     id1  75   F normal
## 6    GroupB_2          H7           NA       GroupB      NA     id2  58   F cancer
```

```r
targets[,10:12]
```

```
##    Array       Slide
## 1 R02C02 5723646052
## 2 R04C01 5723646052
## 3 R05C02 5723646052
## 4 R04C02 5723646053
## 5 R05C02 5723646053
## 6 R06C02 5723646053
```

```
##                                                                      Basename
## 1 /home/biocbuild/bbs-3.3-bioc/R/library/minfiData/extdata/5723646052/5723646052_R02C02
## 2 /home/biocbuild/bbs-3.3-bioc/R/library/minfiData/extdata/5723646052/5723646052_R04C01
## 3 /home/biocbuild/bbs-3.3-bioc/R/library/minfiData/extdata/5723646052/5723646052_R05C02
## 4 /home/biocbuild/bbs-3.3-bioc/R/library/minfiData/extdata/5723646053/5723646053_R04C02
## 5 /home/biocbuild/bbs-3.3-bioc/R/library/minfiData/extdata/5723646053/5723646053_R05C02
## 6 /home/biocbuild/bbs-3.3-bioc/R/library/minfiData/extdata/5723646053/5723646053_R06C02

rgSet <- read.450k.exp(targets = targets)

## Warning:  'read.450k.exp' is deprecated.
## Use 'read.metharray.exp' instead.
## See help("Deprecated")

## Warning:  'read.450k' is deprecated.
## Use 'read.metharray' instead.
## See help("Deprecated")
```

The data is now an *RGChannelSet* object and needs to be normalised and converted to a *MethylSet* object.

3   Subset-quantile within array normalization (SWAN)

SWAN (subset-quantile within array normalization) is a within-array normalization method for Illumina 450k BeadChips. Technical differencs have been demonstrated to exist between the Infinium I and Infinium II assays on a single 450K array [2, 3]. Using the SWAN method substantially reduces the technical variability between the assay designs whilst maintaining important biological differences. The SWAN method makes the assumption that the number of CpGs within the 50bp probe sequence reflects the underlying biology of the region being interrogated. Hence, the overall distribution of intensities of probes with the same number of CpGs in the probe body should be the same regardless of assay type. The method then uses a subset quantile normalization approach to adjust the intensities of each array [9].

SWAN can take a *MethylSet*, *RGChannelSet* or *MethyLumiSet* as input. It should be noted that, in order to create the normalization subset, SWAN randomly selects Infinium I and II probes that have one, two and three underlying CpGs; as such, we recommend using set.seed before SWAN to ensure that the normalized intensities will be identical, if the normalization is repeated.

The technical differences between Infinium I and II assay designs can result in aberrant beta value distributions (Figure 1, panel "Raw"). Using SWAN corrects for the technical differences between the Infinium I and II assay designs and produces a smoother overall β value distribution (Figure 1, panel "SWAN").

```
mSet <- preprocessRaw(rgSet)
mSetSw <- SWAN(mSet,verbose=TRUE)

## [SWAN] Preparing normalization subset
## [SWAN] Normalizing methylated channel
## [SWAN] Normalizing array 1 of 6
## [SWAN] Normalizing array 2 of 6
## [SWAN] Normalizing array 3 of 6
## [SWAN] Normalizing array 4 of 6
## [SWAN] Normalizing array 5 of 6
## [SWAN] Normalizing array 6 of 6
## [SWAN] Normalizing unmethylated channel
## [SWAN] Normalizing array 1 of 6
## [SWAN] Normalizing array 2 of 6
## [SWAN] Normalizing array 3 of 6
## [SWAN] Normalizing array 4 of 6
```

```
## [SWAN] Normalizing array 5 of 6
## [SWAN] Normalizing array 6 of 6
```

```
par(mfrow=c(1,2), cex=1.25)
densityByProbeType(mSet[,1], main = "Raw")
densityByProbeType(mSetSw[,1], main = "SWAN")
```
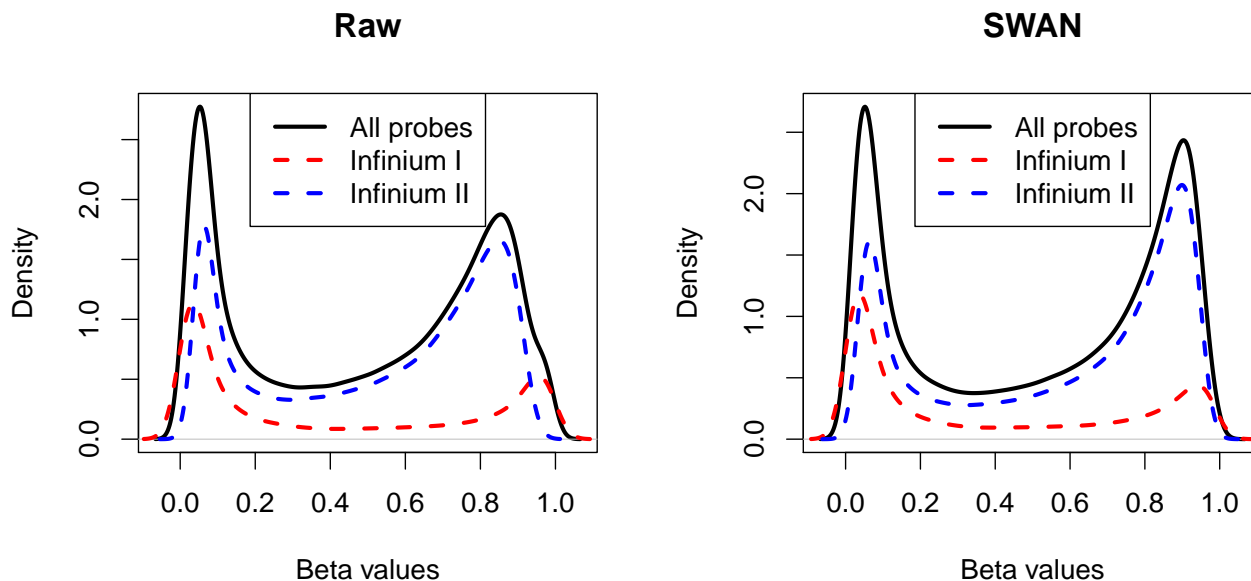

Figure 1: Density distributions of β values before and after using SWAN.

4 Filter out poor quality probes

Poor quality probes can be filtered out based on the detection p-value. For this example, to retain a CpG for further analysis, we require that the detection p-value is less than 0.01 in all samples.

```
detP <- detectionP(rgSet)
keep <- rowSums(detP < 0.01) == ncol(rgSet)
mSetSw <- mSetSw[keep,]
```

5 Extracting Beta and M-values

Now that the data has been SWAN normalised we can extract β and M-values from the *MethylSet* object. We prefer to add an offset to the methylated and unmethylated intensities when calculating M-values, hence we extract the methylated and unmethylated channels separately and perform our own calculation. For all subsequent analysis we use a random selection of 20 000 CpGs to reduce computation time.

```
mset_reduced <- mSetSw[sample(1:nrow(mSetSw), 20000),]
meth <- getMeth(mset_reduced)
unmeth <- getUnmeth(mset_reduced)
Mval <- log2((meth + 100)/(unmeth + 100))
```

```
beta <- getBeta(mset_reduced)
dim(Mval)
```

```
## [1] 20000     6
```

```
par(mfrow=c(1,1))
plotMDS(Mval, labels=targets$Sample_Name, col=as.integer(factor(targets$status)))
legend("topleft",legend=c("Cancer","Normal"),pch=16,cex=1.2,col=1:2)
```


Figure 2: A multi-dimensional scaling (MDS) plot of cancer and normal samples.

An MDS plot (Figure 2) is a good sanity check to make sure samples cluster together according to the main factor of interest, in this case, cancer and normal.

6 Testing for differential methylation using limma

To test for differential methylation we use the *limma* package [12], which employs an empirical Bayes framework based on Guassian model theory. First we need to set up the design matrix. There are a number of ways to do this, the most straightforward is directly from the targets file. There are a number of variables, with the status column indicating cancer/normal samples. From the person column of the targets file, we see that the cancer/normal samples are matched, with 3 individuals each contributing both a cancer and normal sample. Since the limma model framework can handle

any experimental design which can be summarised by a design matrix, we can take into account the paired nature of the data in the analysis. For more complicated experimental designs, please refer to the *limma* User's Guide.

```
group <- factor(targets$status,levels=c("normal","cancer"))
id <- factor(targets$person)
design <- model.matrix(~id + group)
design

##   (Intercept) idid2 idid3 groupcancer
## 1           1     0     1           0
## 2           1     1     0           0
## 3           1     0     1           1
## 4           1     0     0           1
## 5           1     0     0           0
## 6           1     1     0           1
## attr(,"assign")
## [1] 0 1 1 2
## attr(,"contrasts")
## attr(,"contrasts")$id
## [1] "contr.treatment"
##
## attr(,"contrasts")$group
## [1] "contr.treatment"
```

Now we can test for differential methylation using the `lmFit` and `eBayes` functions from *limma*. As input data we use the matrix of M-values.

```
fit.reduced <- lmFit(Mval,design)
fit.reduced <- eBayes(fit.reduced)
```

The numbers of hyper-methylated (1) and hypo-methylated (-1) can be displayed using the `decideTests` function in *limma* and the top 10 differentially methylated CpGs for cancer versus normal outputted using `topTable`.

```
summary(decideTests(fit.reduced))

##    (Intercept) idid2 idid3 groupcancer
## -1        6954     0   146         485
## 0         3352 20000 19847       19086
## 1         9694     0     7         429

top<-topTable(fit.reduced,coef=4)
top

##                logFC     AveExpr        t      P.Value  adj.P.Val        B
## cg21101720  3.956688 -1.32541753 15.71947 1.208689e-05 0.03212398 3.846384
## cg11946336  4.058335 -1.34573254 15.54663 1.280138e-05 0.03212398 3.805601
## cg03735888  5.150500 -1.23542597 14.78159 1.663267e-05 0.03212398 3.615267
## cg24862510  4.182276 -0.74663716 14.57860 1.786821e-05 0.03212398 3.561930
## cg02855633  4.738628 -0.74231987 14.46792 1.858758e-05 0.03212398 3.532323
## cg09043127  3.980513 -1.40197284 14.19464 2.051606e-05 0.03212398 3.457580
## cg16306898  5.400654 -2.24243473 14.01361 2.192448e-05 0.03212398 3.406751
## cg21900495  3.601777  0.09010265 13.83240 2.345052e-05 0.03212398 3.354786
## cg01978558  3.635464 -0.82462802 13.79553 2.377623e-05 0.03212398 3.344078
## cg21294812  3.623604 -1.08712229 13.58340 2.575780e-05 0.03212398 3.281561
```

Note that since we performed our analysis on M-values, the logFC and AveExpr columns are computed on the M-value scale. For interpretability and visualisation we can look at the β values. The beta values for the top 4 differentially methylated CpGs shown in Figure 3.

```
cpgs <- rownames(top)
par(mfrow=c(2,2))
for(i in 1:4){
stripchart(beta[rownames(beta)==cpgs[i],]~design[,4],method="jitter",
group.names=c("Normal","Cancer"),pch=16,cex=1.5,col=c(4,2),ylab="Beta values",
vertical=TRUE,cex.axis=1.5,cex.lab=1.5)
title(cpgs[i],cex.main=1.5)
}
```
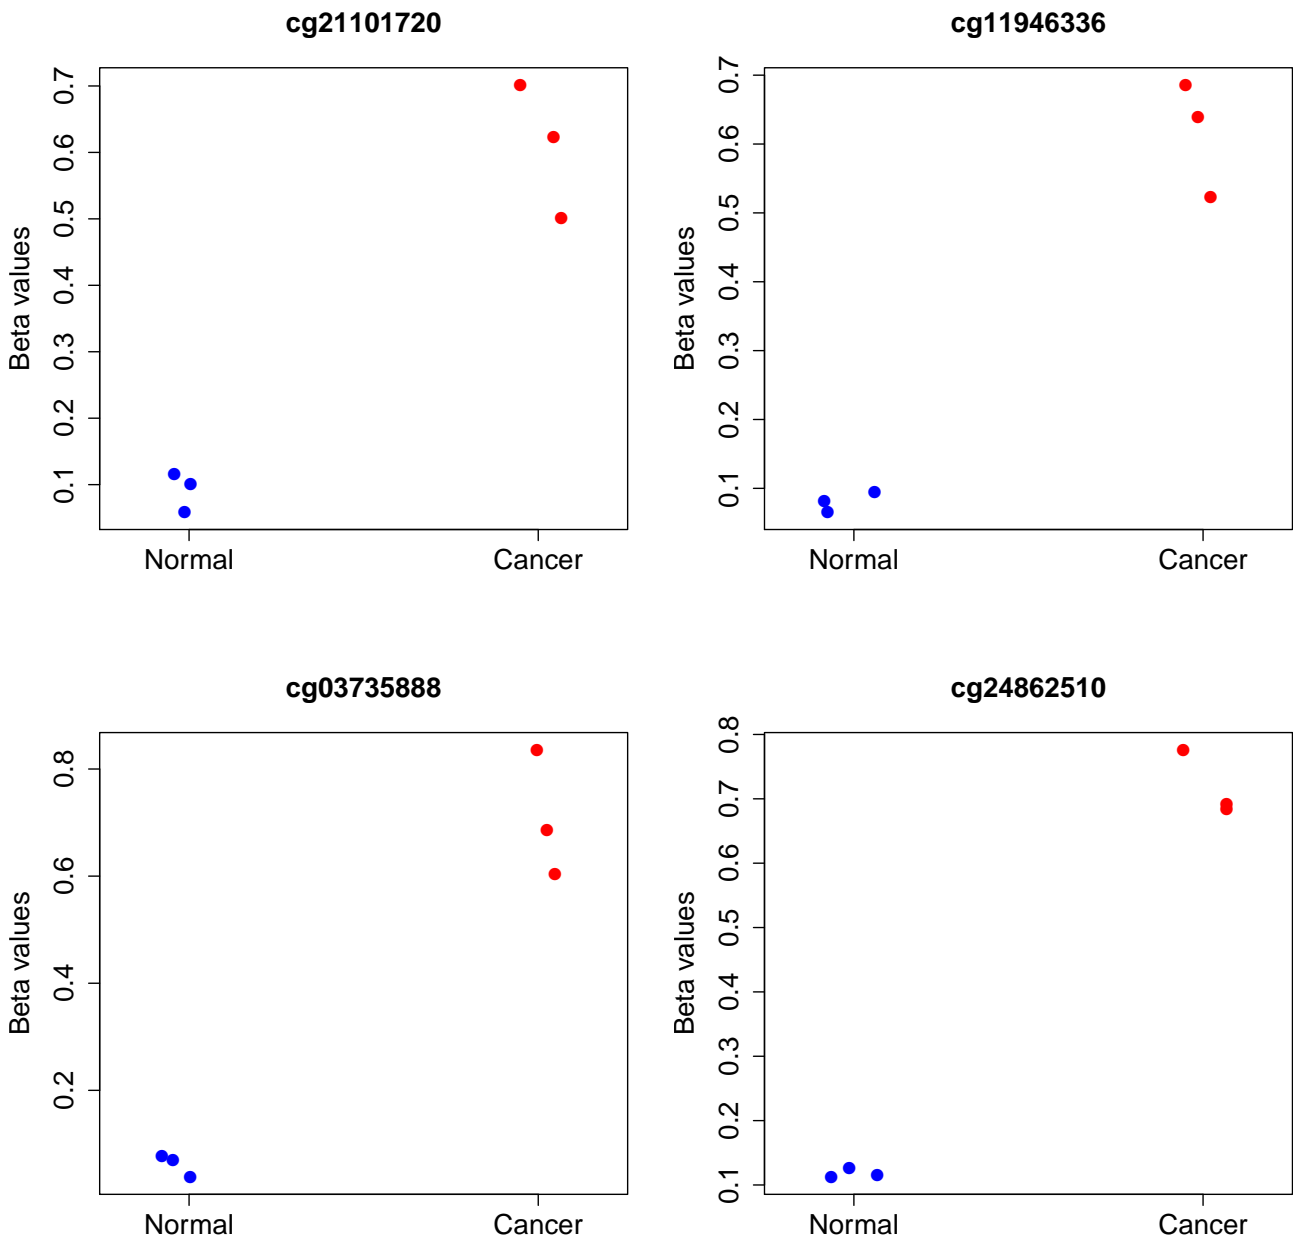

Figure 3: The β values for the top 4 differentially methylated CpGs.

7   Removing unwanted variation when testing for differential methylation

Like other platforms, 450k array studies are subject to unwanted technical variation such as batch effects and other, often unknown, sources of variation. The adverse effects of unwanted variation have been extensively documented in gene expression array studies and have been shown to be able to both reduce power to detect true differences and to increase the number of false discoveries. As such, when it is apparent that data is significantly affected by unwanted variation, it is advisable to perform an adjustment to mitigate its effects.

missMethyl provides a *limma*-like interface to functions from the CRAN *ruv* package that enables the removal of unwanted variation when performing a differential analysis [14]. All of the *ruv* methods rely on negative control features to accurately estimate the components of unwanted variation. Negative control features are probes/genes/etc. that are known *a priori* to not truly be associated with the biological factor of interest, but are affected by unwanted variation. For example, in a microarray gene expression study, these could be house-keeping genes or a set of spike-in controls. Negative control features are extensively discussed in Gagnon-Bartsch and Speed (2012) [4] and Gagnon-Bartsch et al. (2013) [5]. Once the unwanted factors are accurately estimated from the data, they are adjusted for in the linear model that describes the differential analysis.

If the negative control features are not known *a priori*, they can be identified empirically. This can be achieved using a 2-stage approach, *RUVm*, based on **RUV-inverse**. Stage 1 involves performing a differential methylation analysis using **RUV-inverse** and the 613 Illumina negative controls (INCs) as negative control features. This will produce a list of CpGs ranked by p-value according to their level of association with the factor of interest. This list can then be used to identify a set of empirical control probes (ECPs), which will capture more of the unwanted variation than using the INCs alone. ECPs are selected by designating a proportion of the CpGs least associated with the factor of interest as negative control features; this can be done based on either an FDR cut-off or by taking a fixed percentage of probes from the bottom of the ranked list. Stage 2 involves performing a second differential methylation analysis on the original data using **RUV-inverse** and the ECPs. For simplicity, we are ignoring the paired nature of the cancer and normal samples in this example.

```
# get M-values for ALL probes
meth <- getMeth(mSet)
unmeth <- getUnmeth(mSet)
M <- log2((meth + 100)/(unmeth + 100))

grp <- factor(targets$status,levels=c("normal","cancer"))
des <- model.matrix(~grp)
des

##   (Intercept) grpcancer
## 1           1         0
## 2           1         0
## 3           1         1
## 4           1         1
## 5           1         0
## 6           1         1
## attr(,"assign")
## [1] 0 1
## attr(,"contrasts")
## attr(,"contrasts")$grp
## [1] "contr.treatment"

INCs <- getINCs(rgSet)
head(INCs)

##          5723646052_R02C02 5723646052_R04C01 5723646052_R05C02 5723646053_R04C02
## 13792480        -0.3299654        -1.0955482        -0.5266103        -0.6374299
## 69649505        -1.0354488        -1.4943396        -1.0067050        -0.8854881
```

```
## 34772371        -1.1286422        -0.2995603        -0.8192636        -0.6895514
## 28715352        -0.5553373        -0.7599489        -0.7186973        -1.7903619
## 74737439        -1.1169178        -0.8656399        -0.6429681        -0.8872082
## 33730459        -0.7714684        -0.5622424        -0.7724825        -1.5623138
##           5723646053_R05C02 5723646053_R06C02
## 13792480        -1.116598        -0.4332793
## 69649505        -1.586679        -0.9217329
## 34772371        -1.161155        -0.6186795
## 28715352        -1.348105        -1.0067259
## 74737439        -1.064986        -0.9841833
## 33730459        -2.079184        -1.0445246
```

```r
Mc <- rbind(M,INCs)
ctl <- rownames(Mc) %in% rownames(INCs)
table(ctl)
```

```
## ctl
##  FALSE    TRUE
## 485512     613
```

```r
rfit1 <- RUVfit(data=Mc, design=des, coef=2, ctl=ctl) # Stage 1 analysis
rfit2 <- RUVadj(rfit1)
```

Now that we have performed an initial differential methylation analysis to rank the CpGs with respect to their association with the factor of interest, we can designate the CpGs that are least associated with the factor of interest based on FDR-adjusted p-value as ECPs.

```r
top1 <- topRUV(rfit2, num=Inf)
head(top1)
```

```
##            coefficients        t            p      p.BH      p.ebayes p.ebayes.BH
## cg04743961     4.838190 26.74467 3.812882e-05 0.1401969 3.516091e-07  0.01017357
## cg07155336     5.887409 17.62103 1.608653e-04 0.1401969 3.583107e-07  0.01017357
## cg20925841     4.790211 26.69524 3.837354e-05 0.1401969 3.730375e-07  0.01017357
## cg03607359     4.394397 34.74068 1.542013e-05 0.1401969 4.721205e-07  0.01017357
## cg10566121     4.787914 21.80693 7.717708e-05 0.1401969 5.238865e-07  0.01017357
## cg07655636     4.571758 22.99708 6.424216e-05 0.1401969 6.080091e-07  0.01017357
```

```r
ctl <- rownames(M) %in% rownames(top1[top1$p.ebayes.BH > 0.5,])
table(ctl)
```

```
## ctl
##  FALSE    TRUE
## 172540 312972
```

We can then use the ECPs to perform a second differential methylation with **RUV-inverse**, which is adjusted for the unwanted variation estimated from the data.

```r
# Perform RUV adjustment and fit
rfit1 <- RUVfit(data=M, design=des, coef=2, ctl=ctl) # Stage 2 analysis
rfit2 <- RUVadj(rfit1)

# Look at table of top results
topRUV(rfit2)
```

```
##            coefficients        t           p      p.BH      p.ebayes p.ebayes.BH
## cg07155336     5.769286 15.345069 0.002005546 0.3431163 1.434834e-55 6.966293e-50
## cg06463958     5.733093 15.434797 0.001978272 0.3431163 6.749298e-55 1.638433e-49
## cg00024472     5.662959 15.946200 0.001832444 0.3431163 1.319390e-53 2.135266e-48
```

```
## cg02040433    5.651399 10.054445 0.005389436 0.3431163 2.146210e-53 2.605027e-48
## cg13355248    5.595396  9.963702 0.005504213 0.3431163 2.234891e-52 2.022589e-47
## cg02467990    5.592707  6.859614 0.013008521 0.3431163 2.499534e-52 2.022589e-47
## cg00817367    5.527501 13.070583 0.002921656 0.3431163 3.710480e-51 2.573547e-46
## cg11396157    5.487992 10.931263 0.004436178 0.3431163 1.873636e-50 1.137091e-45
## cg16306898    5.466780  5.573935 0.020790127 0.3431163 4.448085e-50 2.399554e-45
## cg03735888    5.396242 15.482605 0.001963955 0.3431163 7.700032e-49 3.738458e-44
```

Note, at present `RUVfit` does not support contrasts, so only one factor of interest can be interrogated at a time using a
design matrix with an intercept term.

8   Testing for differential variability (DiffVar)

8.1   Methylation data

Rather than testing for differences in mean methylation, we may be interested in testing for differences between group
variances. For example, it has been hypothesised that highly variable CpGs in cancer are important for tumour progression
[7]. Hence we may be interested in CpG sites that are consistently methylated in the normal samples, but variably
methylated in the cancer samples.

In general we recommend at least 10 samples in each group for accurate variance estimation, however for the purpose
of this vignette we perform the analysis on 3 vs 3. In this example, we are interested in testing for differential variability
in the cancer versus normal group. Note that when we specify the *coef* parameter, which corresponds to the columns
of the design matrix to be used for testing differential variability, we need to specify both the intercept and the fourth
column. The ID variable is a nuisance parameter and not used when obtaining the absolute deviations, however it can be
included in the linear modelling step. For methylation data, the `varFit` function will take either a matrix of M-values,
β values or a *MethylSet* object as input. If β values are supplied, a logit transformation is performed. Note that as a
default, `varFit` uses the robust setting in the *limma* framework, which requires the use of the *statmod* package.

```
fitvar <- varFit(Mval, design = design, coef = c(1,4))
```

The numbers of hyper-variable (1) and hypo-variable (-1) genes in cancer vs normal can be obtained using `decideTests`.

```
summary(decideTests(fitvar))
```

```
##    (Intercept) idid2 idid3 groupcancer
## -1           0     2     2           0
## 0        19679 19991 19947       19987
## 1          321     7    51          13
```

```
topDV <- topVar(fitvar, coef=4)
topDV
```

```
##            SampleVar LogVarRatio DiffLevene        t      P.Value  Adj.P.Value
## cg19089337  7.367676    2.459879   2.571048 6.509713 7.618375e-11 1.523675e-06
## cg05196820  7.468081    4.637874   2.959574 5.659684 1.526873e-08 1.526873e-04
## cg23209255  7.909645    4.061495   2.769041 5.234016 1.667017e-07 1.111344e-03
## cg10192047  7.857573    2.162623   2.162711 5.130033 2.910411e-07 1.455206e-03
## cg07210840  4.463025    2.774897   2.354299 4.628899 3.687687e-06 1.475075e-02
## cg12902357  6.719135    1.356153   1.686641 4.574926 4.778177e-06 1.592726e-02
## cg19557340  7.827875    2.166208   2.031353 4.468747 7.889478e-06 2.254136e-02
## cg22587602  5.522689    2.100709   1.982476 4.426572 9.599549e-06 2.399887e-02
## cg04089426  5.180724    5.724316   2.728900 4.331737 1.482978e-05 3.103121e-02
## cg07509935  4.282834    2.672635   2.171044 4.318218 1.576720e-05 3.103121e-02
```

An alternate parameterisation of the design matrix that does not include an intercept term can also be used, and specific contrasts tested with `contrasts.varFit`. Here we specify the design matrix such that the first two columns correspond to the normal and cancer groups, respectively.

```
design2 <- model.matrix(~0+group+id)
fitvar.contr <- varFit(Mval, design=design2, coef=c(1,2))
contr <- makeContrasts(groupcancer-groupnormal,levels=colnames(design2))
fitvar.contr <- contrasts.varFit(fitvar.contr,contrasts=contr)
```

The results are identical to before.

```
summary(decideTests(fitvar.contr))

##     groupcancer - groupnormal
## -1                          0
## 0                       19987
## 1                          13

topVar(fitvar.contr,coef=1)

##            SampleVar LogVarRatio DiffLevene        t      P.Value  Adj.P.Value
## cg19089337  7.367676    2.459879   2.571048 6.509713 7.618375e-11 1.523675e-06
## cg05196820  7.468081    4.637874   2.959574 5.659684 1.526873e-08 1.526873e-04
## cg23209255  7.909645    4.061495   2.769041 5.234016 1.667017e-07 1.111344e-03
## cg10192047  7.857573    2.162623   2.162711 5.130033 2.910411e-07 1.455206e-03
## cg07210840  4.463025    2.774897   2.354299 4.628899 3.687687e-06 1.475075e-02
## cg12902357  6.719135    1.356153   1.686641 4.574926 4.778177e-06 1.592726e-02
## cg19557340  7.827875    2.166208   2.031353 4.468747 7.889478e-06 2.254136e-02
## cg22587602  5.522689    2.100709   1.982476 4.426572 9.599549e-06 2.399887e-02
## cg04089426  5.180724    5.724316   2.728900 4.331737 1.482978e-05 3.103121e-02
## cg07509935  4.282834    2.672635   2.171044 4.318218 1.576720e-05 3.103121e-02
```

The β values for the top 4 differentially variable CpGs can be seen in Figure 4.

8.2 RNA-Seq expression data

Testing for differential variability in expression data is straightforward if the technology is gene expression microarrays. The matrix of expression values can be supplied directly to the `varFit` function. For RNA-Seq data, the mean-variance relationship that occurs in count data needs to be taken into account. In order to deal with this issue, we apply a `voom` transformation [8] to obtain observation weights, which are then used in the linear modelling step. For RNA-Seq data, the `varFit` function will take a *DGEList* object as input.

To demonstrate this, we use data from the *tweeDEseqCountData* package. This data is part of the International HapMap project, consisting of RNA-Seq profiles from 69 unrelated Nigerian individuals [11]. The only covariate is gender, so we can look at differentially variable expression between males and females. We follow the code from the *limma* vignette to read in and process the data before testing for differential variability.

First we load up the data and extract the relevant information.

```
library(tweeDEseqCountData)
data(pickrell1)
counts<-exprs(pickrell1.eset)
dim(counts)

## [1] 38415    69

gender <- pickrell1.eset$gender
table(gender)
```

```
cpgsDV <- rownames(topDV)
par(mfrow=c(2,2))
for(i in 1:4){
stripchart(beta[rownames(beta)==cpgsDV[i],]~design[,4],method="jitter",
group.names=c("Normal","Cancer"),pch=16,cex=1.5,col=c(4,2),ylab="Beta values",
vertical=TRUE,cex.axis=1.5,cex.lab=1.5)
title(cpgsDV[i],cex.main=1.5)
}
```
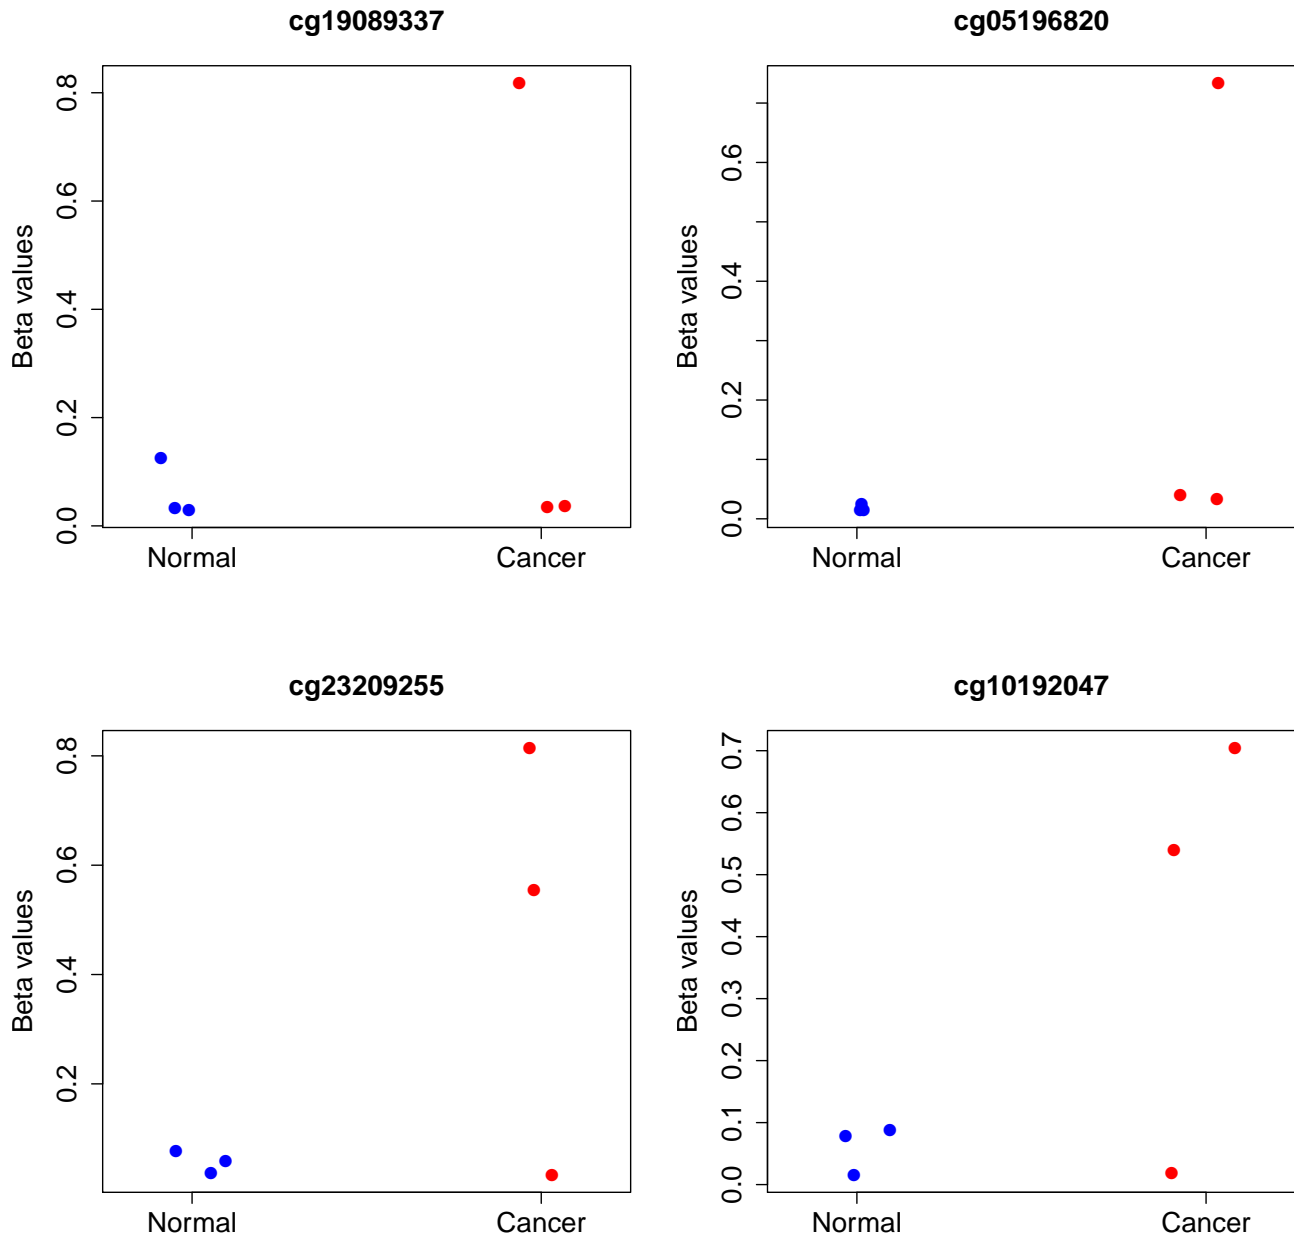

Figure 4: The β values for the top 4 differentially variable CpGs.

```
## gender
## female    male
##     40      29
```

```
rm(pickrell1.eset)
data(genderGenes)
data(annotEnsembl63)
annot <- annotEnsembl63[,c("Symbol","Chr")]
rm(annotEnsembl63)
```

We now have the counts, gender of each sample and annotation (gene symbol and chromosome) for each Ensemble gene.
We can form a *DGEList* object using the *edgeR* package.

```
library(edgeR)
y <- DGEList(counts=counts, genes=annot[rownames(counts),])
```

We filter out lowly expressed genes by keeping genes with at least 1 count per million reads in at least 20 samples, as
well as genes that have defined annotation. Finally we perform scaling normalisation.

```
isexpr <- rowSums(cpm(y)>1) >= 20
hasannot <- rowSums(is.na(y$genes))==0
y <- y[isexpr & hasannot,,keep.lib.sizes=FALSE]
dim(y)
```

```
## [1] 17310    69
```

```
y <- calcNormFactors(y)
```

We set up the design matrix and test for differential variability. In this case there are no nuisance parameters, so *coef*
does not need to be explicitly specified.

```
design.hapmap <- model.matrix(~gender)
fitvar.hapmap <- varFit(y, design = design.hapmap)
```

```
## Converting counts to log counts-per-million using voom.
```

```
fitvar.hapmap$genes <- y$genes
```

We can display the results of the test:

```
summary(decideTests(fitvar.hapmap))
```

```
##    (Intercept) gendermale
## -1           0          2
## 0            0      17308
## 1        17310          0
```

```
topDV.hapmap <- topVar(fitvar.hapmap,coef=ncol(design.hapmap))
topDV.hapmap
```

```
##                       Symbol Chr  SampleVar LogVarRatio DiffLevene        t      P.Value
## ENSG00000213318 RP11-331F4.1  16 5.69839463   -2.562939 -0.9859943 -8.031075 7.238901e-12
## ENSG00000129824        RPS4Y1   Y 2.32497726   -2.087025 -0.4585620 -4.957078 3.959119e-06
## ENSG00000233864        TTTY15   Y 6.79004140   -2.245369 -0.6085233 -4.612613 1.497880e-05
## ENSG00000176171         BNIP3  10 0.41317384    1.199292  0.3632133  4.219459 6.439806e-05
## ENSG00000197358        BNIP3P1 14 0.39969125    1.149754  0.3353288  4.058182 1.147653e-04
## ENSG00000025039          RRAGD   6 0.91837213    1.091229  0.4926839  3.977116 1.527080e-04
## ENSG00000103671          TRIP4  15 0.07456448   -1.457139 -0.1520583 -3.911199 1.921645e-04
## ENSG00000171100           MTM1   X 0.44049558   -1.133295 -0.3334619 -3.896210 2.024119e-04
## ENSG00000149476            DAK  11 0.13289523   -1.470460 -0.1919880 -3.785886 2.956246e-04
## ENSG00000064886         CHI3L2   1 2.70234584    1.468059  0.8449434  3.782149 2.994084e-04
```

```
##                  Adj.P.Value
## ENSG00000213318 1.253054e-07
## ENSG00000129824 3.426618e-02
## ENSG00000233864 8.642769e-02
## ENSG00000176171 2.786826e-01
## ENSG00000197358 3.973173e-01
## ENSG00000025039 4.379688e-01
## ENSG00000103671 4.379688e-01
## ENSG00000171100 4.379688e-01
## ENSG00000149476 4.428665e-01
## ENSG00000064886 4.428665e-01
```

The log counts per million for the top 4 differentially variable genes can be seen in Figure 5.

9   Gene ontology analysis

Once a differential methylation or differential variability analysis has been performed, it may be of interest to know which gene pathways are targeted by the significant CpG sites. It is not entirely clear from the literature how best to perform such an analysis, however Geeleher et al. [6] showed there is a severe bias when performing gene ontology analysis with methylation data. This is due to the fact that there are differing numbers of probes per gene on several different array technologies. For the Illumina Infinium HumanMethylation450 array the number of probes per gene ranges from 1 to 1299, with a median of 15 probes per gene. This means that when mapping CpG sites to genes, a gene is more likely to be "selected" if there are many CpG sites associated with the gene.

One way to take into account this selection bias is to model the relationship between the number of probes per gene and the probability of being selected. This can be performed by adapting the *goseq* method of Young et al. [13]. Each gene then has a prior probability associated with it, and a modified version of a hypergeometric test can be performed, testing for over-representation of the selected genes in each gene set.

The function `gometh` performs gene set testing on GO categories or KEGG pathways [15]. The function `gsameth` is a more generalised gene set testing function which can take as input a list of user specified gene sets. Note that for gsameth, the format for the gene ids for each gene in the gene set needs to be Entrez Gene IDs. For example, the entire curated gene set list (C2) from the Broad's Molecular Signatures Database can be specified as input. The R version of these lists can be downloaded from http://bioinf.wehi.edu.au/software/MSigDB/index.html. Both functions take a vector of significant CpG probe names as input.

To illustrate how to use `gometh`, consider the results from the differential methylation analysis with ruv.

```
topRUV(rfit2)

##            coefficients         t           p       p.BH     p.ebayes  p.ebayes.BH
## cg07155336     5.769286 15.345069 0.002005546 0.3431163 1.434834e-55 6.966293e-50
## cg06463958     5.733093 15.434797 0.001978272 0.3431163 6.749298e-55 1.638433e-49
## cg00024472     5.662959 15.946200 0.001832444 0.3431163 1.319390e-53 2.135266e-48
## cg02040433     5.651399 10.054445 0.005389436 0.3431163 2.146210e-53 2.605027e-48
## cg13355248     5.595396  9.963702 0.005504213 0.3431163 2.234891e-52 2.022589e-47
## cg02467990     5.592707  6.859614 0.013008521 0.3431163 2.499534e-52 2.022589e-47
## cg00817367     5.527501 13.070583 0.002921656 0.3431163 3.710480e-51 2.573547e-46
## cg11396157     5.487992 10.931263 0.004436178 0.3431163 1.873636e-50 1.137091e-45
## cg16306898     5.466780  5.573935 0.020790127 0.3431163 4.448085e-50 2.399554e-45
## cg03735888     5.396242 15.482605 0.001963955 0.3431163 7.700032e-49 3.738458e-44

table(rfit2$p.ebayes.BH < 0.01)

##
##  FALSE   TRUE
```

```
genesDV <- rownames(topDV.hapmap)
par(mfrow=c(2,2))
for(i in 1:4){
stripchart(cpm(y,log=TRUE)[rownames(y)==genesDV[i],]~design.hapmap[,ncol(design.hapmap)],method="jitter",
group.names=c("Female","Male"),pch=16,cex=1.5,col=c(4,2),ylab="Log counts per million",
vertical=TRUE,cex.axis=1.5,cex.lab=1.5)
title(genesDV[i],cex.main=1.5)
}
```
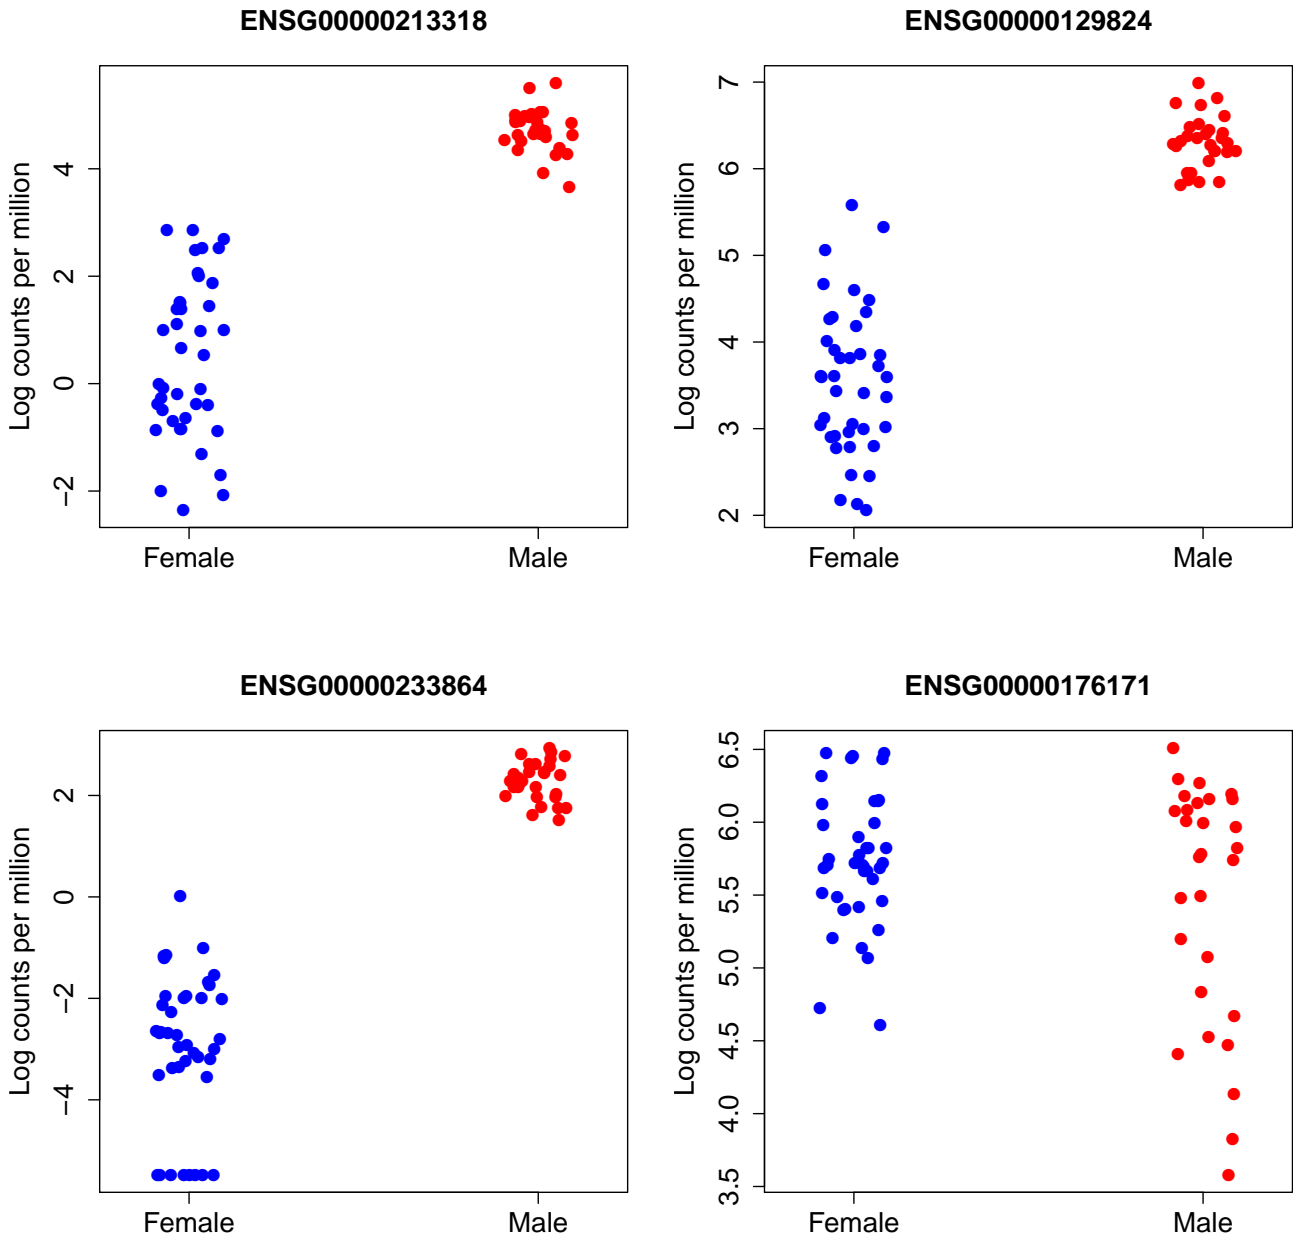

Figure 5: The log counts per million for the top 4 differentially variably expressed genes.

```
## 424168  61344
```

At a 1% false discovery rate cut-off, there are still tens of thousands of CpG sites differentially methylated. These will undoubtably map to almost all the genes in the genome, making a gene ontology analysis irrelevant. One option for selecting CpGs in this context is to apply not only a false discovery rate cut-off, but also a $\delta\beta$ cut-off. However, for this dataset, taking a relatively large $\delta\beta$ cut-off of 0.25 still leaves more than 30 000 CpGs differentially methylated.

```
beta <- getBeta(mSet)
beta_norm <- rowMeans(beta[,des[,2]==0])
beta_can <- rowMeans(beta[,des[,2]==1])
delta_beta <- beta_can - beta_norm
sigDM <- rfit2$p.ebayes.BH < 0.01 & abs(delta_beta) > 0.25
table(sigDM)

## sigDM
##  FALSE    TRUE
## 451760   33748
```

Instead, we take the top 10 000 CpG sites as input to gometh.

```
topCpGs<-topRUV(rfit2,number=10000)
sigCpGs <- rownames(topCpGs)
sigCpGs[1:10]

##  [1] "cg07155336" "cg06463958" "cg00024472" "cg02040433" "cg13355248" "cg02467990"
##  [7] "cg00817367" "cg11396157" "cg16306898" "cg03735888"
```

The gometh takes as input a character vector of CpG names, and optionally, a character vector of all CpG sites tested. If the *all.cpg* argument is omitted, all the CpGs on the array are used as background. If the *plot.bias* argument is TRUE, a figure showing the relationship between the probability of being selected and the number of probes per gene will be displayed. For testing of GO terms, the *collection* argument takes the value "GO", which is the default setting. For KEGG pathway analysis, set *collection* to "KEGG". The function topGO shows the top enriched GO categories. For KEGG testing, use the topKEGG function. The *limma* functions goana and kegga are called by gometh for GO and KEGG pathway analysis respectively.

For GO testing on our example dataset:

```
library(IlluminaHumanMethylation450kanno.ilmn12.hg19)
gst <- gometh(sig.cpg=sigCpGs, all.cpg=rownames(rfit2), collection="GO")

## Warning in alias2SymbolTable(flat$symbol):  Multiple symbols ignored for one or more aliases

topGO(gst)

##                                             Term Ont    N   DE          P.DE          FDR
## GO:0044707 single-multicellular organism process  BP 6002 1145 6.469346e-36 1.346853e-31
## GO:0048731                      system development  BP 4343  913 1.402302e-35 1.459726e-31
## GO:0048856      anatomical structure development  BP 5117 1017 1.673360e-33 1.161256e-29
## GO:0032501        multicellular organismal process  BP 6916 1221 3.020933e-33 1.572320e-29
## GO:0007275    multicellular organism development  BP 4893  982 4.216205e-33 1.755544e-29
## GO:0007399          nervous system development  BP 2358  589 4.553971e-32 1.580152e-28
## GO:0005886                        plasma membrane  CC 4754  877 3.528908e-30 1.049548e-26
## GO:0044767 single-organism developmental process  BP 5671 1075 5.373794e-30 1.398463e-26
## GO:0032502                developmental process  BP 5755 1082 5.039218e-29 1.165683e-25
## GO:0071944                        cell periphery  CC 4859  887 1.522801e-28 3.170318e-25
## GO:0048513            animal organ development  BP 3106  665 7.256169e-27 1.373329e-23
## GO:0030154                    cell differentiation  BP 3731  772 9.051315e-27 1.570328e-23
## GO:0048869        cellular developmental process  BP 3980  800 4.258878e-25 6.820429e-22
## GO:0007417    central nervous system development  BP  927  277 5.133624e-25 7.634066e-22
```

```
## GO:0009653    anatomical structure morphogenesis  BP 2750  618 9.491879e-25 1.317409e-21
## GO:0048699                  generation of neurons  BP 1600  422 1.300641e-24 1.692378e-21
## GO:0030182                neuron differentiation  BP 1490  399 1.739839e-24 2.130689e-21
## GO:0007267                    cell-cell signaling  BP 1205  313 3.371772e-24 3.899829e-21
## GO:0022008                           neurogenesis  BP 1682  435 4.351450e-24 4.768044e-21
## GO:0031224       intrinsic component of membrane  CC 5207  858 8.151750e-23 8.485564e-20
```

For KEGG testing on our example dataset:

```
gst.kegg <- gometh(sig.cpg=sigCpGs, all.cpg=rownames(rfit2), collection="KEGG")

## Warning in alias2SymbolTable(flat$symbol):  Multiple symbols ignored for one or more aliases

topGO(gst.kegg)

##                                                          Pathway    N DE
## path:hsa04080          Neuroactive ligand-receptor interaction 273 66
## path:hsa05200                                  Pathways in cancer 396 83
## path:hsa04020                            Calcium signaling pathway 178 55
## path:hsa04151                            PI3K-Akt signaling pathway 328 71
## path:hsa04024                             cAMP signaling pathway 199 52
## path:hsa04514                    Cell adhesion molecules (CAMs) 138 43
## path:hsa04510                                      Focal adhesion 202 50
## path:hsa04022                        cGMP-PKG signaling pathway 166 44
## path:hsa04015                            Rap1 signaling pathway 211 48
## path:hsa05205                            Proteoglycans in cancer 203 47
## path:hsa04360                                        Axon guidance 174 46
## path:hsa05206                                  MicroRNAs in cancer 284 48
## path:hsa04072                  Phospholipase D signaling pathway 143 40
## path:hsa04014                            Ras signaling pathway 226 47
## path:hsa04810                    Regulation of actin cytoskeleton 214 45
## path:hsa04060             Cytokine-cytokine receptor interaction 249 37
## path:hsa04550 Signaling pathways regulating pluripotency of stem cells 141 36
## path:hsa04970                                  Salivary secretion  89 27
## path:hsa04724                              Glutamatergic synapse 114 33
## path:hsa04725                                Cholinergic synapse 111 32
##                     P.DE          FDR
## path:hsa04080 2.041302e-41 6.164731e-39
## path:hsa05200 2.095870e-37 3.164763e-35
## path:hsa04020 7.600910e-35 7.651583e-33
## path:hsa04151 1.424062e-34 1.075167e-32
## path:hsa04024 5.684418e-30 3.433388e-28
## path:hsa04514 5.781373e-28 2.909958e-26
## path:hsa04510 2.277061e-25 9.823891e-24
## path:hsa04022 4.216484e-25 1.591723e-23
## path:hsa04015 2.163487e-23 7.259702e-22
## path:hsa05205 4.183137e-23 1.263307e-21
## path:hsa04360 4.733606e-23 1.299590e-21
## path:hsa05206 1.076287e-22 2.708656e-21
## path:hsa04072 1.749168e-22 4.063453e-21
## path:hsa04014 6.897267e-22 1.487839e-20
## path:hsa04810 3.118187e-21 6.277950e-20
## path:hsa04060 1.033028e-19 1.949841e-18
## path:hsa04550 1.486761e-19 2.641188e-18
## path:hsa04970 1.798783e-19 3.017957e-18
## path:hsa04724 2.432492e-19 3.866382e-18
```

```
## path:hsa04725 1.671168e-18 2.523464e-17
```

For a more generalised version of gene set testing in methylation data where the user can specify the gene set to be tested, the gsameth function can be used. To display the top 20 pathways, topGSA can be called. gsameth can take a single gene set, or a list of gene sets. The gene identifiers in the gene set must be Entrez Gene IDs. To demonstrate gsameth a toy example is shown below, with gene sets made up of randomly selected genes from the *org.Hs.eg.db* package.

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

## Loading required package:  AnnotationDbi

genes <- toTable(org.Hs.egSYMBOL2EG)
set1 <- sample(genes$gene_id,size=80)
set2 <- sample(genes$gene_id,size=100)
set3 <- sample(genes$gene_id,size=30)
genesets <- list(set1,set2,set3)
gsa <- gsameth(sig.cpg=sigCpGs, all.cpg=rownames(rfit2), collection=genesets)

## Warning in alias2SymbolTable(flat$symbol):  Multiple symbols ignored for one or more aliases

topGSA(gsa)

##         N DE       P.DE        FDR
## [1,]   30  2 0.06919509 0.2075853
## [2,]   80  2 0.40975858 0.6146379
## [3,] 100  2 0.81034291 0.8103429
```

Note that if it is of interest to obtain the Entrez Gene IDs that the significant CpGs are mapped to, the getMappedEntrezIDs can be called.

10 Session information

```
toLatex(sessionInfo())
```

- R version 3.3.0 (2016-05-03), x86_64-pc-linux-gnu
- Locale: LC_CTYPE=en_US.UTF-8, LC_NUMERIC=C, LC_TIME=en_US.UTF-8, LC_COLLATE=C,
  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, grDevices, graphics, methods, parallel, stats, stats4, utils
- Other packages: AnnotationDbi 1.34.2, Biobase 2.32.0, BiocGenerics 0.18.0, Biostrings 2.40.0,
  GenomeInfoDb 1.8.2, GenomicRanges 1.24.0, IRanges 2.6.0,
  IlluminaHumanMethylation450kanno.ilmn12.hg19 0.2.1, IlluminaHumanMethylation450kmanifest 0.4.0,
  S4Vectors 0.10.0, SummarizedExperiment 1.2.2, XVector 0.12.0, bumphunter 1.12.0, edgeR 3.14.0, foreach 1.4.3,
  iterators 1.0.8, lattice 0.20-33, limma 3.28.4, locfit 1.5-9.1, minfi 1.18.2, minfiData 0.14.0, missMethyl 1.6.2,
  org.Hs.eg.db 3.3.0, tweeDEseqCountData 1.10.0
- Loaded via a namespace (and not attached): BiasedUrn 1.07, BiocParallel 1.6.2, BiocStyle 2.0.2, DBI 0.4-1,
  GEOquery 2.38.4, GO.db 3.3.0, GenomicAlignments 1.8.0, GenomicFeatures 1.24.2, MASS 7.3-45, Matrix 1.2-6,
  R6 2.1.2, RColorBrewer 1.1-2, RCurl 1.95-4.8, RSQLite 1.0.0, Rcpp 0.12.5, Rsamtools 1.24.0, XML 3.98-1.4,
  annotate 1.50.0, base64 2.0, beanplot 1.2, biomaRt 2.28.0, bitops 1.0-6, chron 2.3-47, codetools 0.2-14,
  colorspace 1.2-6, data.table 1.9.6, digest 0.6.9, doRNG 1.6, evaluate 0.9, formatR 1.4, genefilter 1.54.2, grid 3.3.0,
  highr 0.6, httr 1.1.0, illuminaio 0.14.0, knitr 1.13, magrittr 1.5, matrixStats 0.50.2, mclust 5.2, methylumi 2.18.2,
  multtest 2.28.0, nlme 3.1-128, nor1mix 1.2-1, openssl 0.9.3, pkgmaker 0.22, plyr 1.8.3, preprocessCore 1.34.0,
  quadprog 1.5-5, registry 0.3, reshape 0.8.5, rngtools 1.2.4, rtracklayer 1.32.0, ruv 0.9.6, siggenes 1.46.0,
  splines 3.3.0, statmod 1.4.24, stringi 1.0-1, stringr 1.0.0, survival 2.39-4, tools 3.3.0, xtable 1.8-2, zlibbioc 1.18.0

References

[1] Aryee, M. J., Jaffe, A. E., Corrada-Bravo, H., Ladd-Acosta, C., Feinberg, A. P., Hansen, K. D., and Irizarry, R. A. (2014). Minfi: a flexible and comprehensive Bioconductor package for the analysis of Infinium DNA methylation microarrays. *Bioinformatics*, **30**(10):1363-1369.

[2] Bibikova, M., Barnes, B., Tsan, C., Ho, V., Klotzle, B, Le, J.M., Delano, D., Zhang, L., Schroth, G.P., Gunderson, K.L., Fan, J., and Shen, R. (2011). High density DNA methylation array with single CpG site resolution. *Genomics*, **98**(4):288-295.

[3] Dedeurwaerder, S., Defrance, M., Calonne, E., Denis, H., Sotiriou, C., and Fuks, F. (2011). Evaluation of the Infinium Methylation 450K technology. *Epigenomics*, **3**(6):771-784.

[4] Gagnon-Bartsch JA, Speed TP. (2012). Using control genes to correct for unwanted variation in microarray data. *Biostatistics*. **13**(3), 539-52. Available at: http://biostatistics.oxfordjournals.org/content/13/3/539.full.

[5] Gagnon-Bartsch, Jacob, and Speed. (2013). Removing Unwanted Variation from High Dimensional Data with Negative Controls. Available at: http://statistics.berkeley.edu/tech-reports/820.

[6] Geeleher, P., Hartnett, L., Egan, L. J., Golden, A., Ali, R. A. R., and Seoighe, C. (2013). Gene-set analysis is severely biased when applied to genome-wide methylation data. *Bioinformatics*, **29**(15), 1851-1857.

[7] Hansen, K.D., Timp, W., Bravo, H.C., Sabunciyan, S., Langmead, B., McDonald, O.G., Wen, B., Wu, H., Liu, Y., Diep, D., Briem, E., Zhang, K., Irizarry, R., and Feinberg, A.P. (2011). Increased methylation variation in epigenetic domains across cancer types. *Nature Genetics*, **43**:768-75.

[8] Law, C.W., Chen, Y., Shi, W., Smyth, G.K. (2014). Voom: precision weights unlock linear model analysis tools for RNA-seq read counts. *Genome Biology*, 15, R29.

[9] Maksimovic, J., Gordon, L., and Oshlack, A. (2012). SWAN: Subset-quantile within array normalization for illumina infinium HumanMethylation450 BeadChips. *Genome Biology*, **13**(6):R44.

[10] Phipson, B., and Oshlack, A. DiffVar: A new method for detecting differential variability with application to methylation in cancer and ageing. *Genome Biology, Accepted 10 September 2014.*

[11] Pickrell, J.K., Marioni, J.C., Pai, A.A., Degner, J.F., Engelhardt, B.E., Nkadori, E., Veyrieras, J.B., Stephens, M., Gilad, Y., and Pritchard, J.K. (2010). Understanding mechanisms underlying human gene expression variation with RNA sequencing. *Nature*, 464, 768-72.

[12] Smyth, G.K. (2005). Limma: linear models for microarray data. In: 'Bioinformatics and Computational Biology Solutions using R and Bioconductor'. R. Gentleman, V. Carey, S. Dudoit, R. Irizarry, W. Huber (eds), Springer, New York, pages 397-420.

[13] Young, M. D., Wakefield, M. J., Smyth, G. K., and Oshlack, A. (2010). Gene ontology analysis for RNA-seq: accounting for selection bias. *Genome Biology*, 11, R14.

[14] Maksimovic, J., Gagnon-Bartsch, J.A., Speed, T.P., and Oshlack, A. (2015). Removing unwanted variation in a differential methylation analysis of Illumina HumanMethylation450 array data. *Nucleic acids research*, gkv526.

[15] Phipson, B., Maksimovic, J., and Oshlack, A. (2015). missMethyl: an R package for analysing methylation data from Illuminas HumanMethylation450 platform. *Bioinformatics*, btv560.