

An Introduction to the **methylumi** package

Sean Davis and Sven Bilke

October 28, 2009

1 Introduction

Gene expression patterns are very important in understanding any biologic system. The regulation of gene expression is a complex, poorly understood process. However, DNA methylation is one biological process that is known to be important in gene regulation. In short, DNA methylation is a chemical modification of DNA CpG sites in which a methyl group is added number 5 carbon of the cytosine pyrimidine ring. In humans, the DNA methyltransferases (DNMT1, DNMT3a, and DNMT3b) are the enzymes responsible for carrying out the methylation.

The Illumina GoldenGate methylation profiling technology specifically targets more than 1500 CpG sites throughout the genome, specifically targeting approximately 700 “cancer genes”. Samples are run in 96-well format, making the technology very high-throughput. After a two-color hybridization, a laser captures the intensities associated with the methylated state and the accompanying unmethylated state. The Illumina BeadStudio software is then used for quality control and basic visualization tasks. A newer platform, the Illumina Infinium Methylation platform provides a more “whole-genome” view of DNA methylation. Utilizing the Infinium profiling technology on bisulfite-treated DNA, the methylation status of more than 25,000 individual CpG sites is assayed simultaneously. This package can handle both types of data.

The **methylumi** package provides convenient mechanisms for loading the results of the Illumina methylation platform into R/Bioconductor. Classes based on common Bioconductor classes for encapsulating the data and facilitate data manipulation are at the core of the package, with methods for quality control, normalization, and plotting.

2 Loading Data

After exporting the data from BeadStudio, **methylumi** can read them in with a single command. To include rich sample annotation, it is possible to supply a `data.frame` including that sample annotation. This can be read from disk or constructed on-the-fly for flexibility. If used, a *SampleID* column must be present and match the sample IDs used in the BeadStudio export file. Also, if a column called *SampleLabel* is present in the data frame and it

includes unique names, the values from that column will be used for the sampleNames of the resulting *MethyLumiSet*.

Two different formats can be read by `methylumi`. The “Final Report” format includes all the data in a single file. The package will look “[Header]” in the file to determine when that file format is to be used. The data block “[Sample Methylation Profile]” needs to be present in the “Final Report” format. If the data block “[Control Probe Profile]” is present, these data will be included in the *QCdata* of the resulting *MethyLumiSet* object. The second format can be a simple tab-delimited text file with headers from BeadStudio. If this format is used, the sample data and the QC data can be in separate files. The QC data need not be present for either format, but it can be helpful for quality control of data. For the examples in this vignette, a small sample set run on the Illumina GoldenGate platform will be used and the file format is the tab-delimited format.

```
> library(methylumi)
> samps <- read.table(system.file("extdata/samples.txt",
+   package = "methylumi"), sep = "\t",
+   header = TRUE)
> mldat <- methylumiR(system.file("extdata/extendedata.samples.txt",
+   package = "methylumi"), qcfile = system.file("extdata/extendedata.controls.txt",
+   package = "methylumi"), sampleDescriptions = samps)
```

Only a subset of an entire plate is included here for illustration purposes. The `mldat` object now contains the data (in an *eSet*-like object) and quality control information (available as `QCdata(mldat)`, also an *eSet*-like object) from a set of experiments. The details of what was loaded can be viewed:

```
> mldat
```

Object Information:

MethyLumiSet (storageMode: environment)

assayData: 1536 features, 10 samples

element names: Avg_NBEADS, BEAD_STDERR, betas, methylated, pvals, unmethylated
phenoData

sampleNames: M_1, M_2, ..., F_10 (10 total)

varLabels and varMetadata description:

sampleID: sampleID

SampleLabel: SampleLabel

Sample: Sample

Gender: Gender

featureData

featureNames: AATK_E63_R, AATK_P519_R, ..., ZP3_P220_F (1536 total)

fvarLabels and fvarMetadata description:

TargetID: NA

```

    ProbeID: NA
    ...: ...
    PRODUCT: NA
    (17 total)
experimentData: use 'experimentData(object)'
Annotation:
Major Operation History:
      submitted          finished
1 2009-10-28 02:36:48 2009-10-28 02:36:49

1 methylumiR(filename = system.file("extdata/exampledata.samples.txt", package = "me

```

Accessors for various slots of the resulting object are outlined in the online help. It is worth noting, though, that the *QCdata* will return another *eSet*-like object of class *MethylumiQC*; the data contained here can be useful if an array has failed.

Note that the assayData names have been changed from the original column identifiers in the data file from Illumina. The mappings are available via the function `getAssayDataNameSubstitutions`.

```
> getAssayDataNameSubstitutions()
```


	regex	result
1	AVG_Beta	betas
2	Detection Pval	pvals
3	Signal CY3	unmethylated
4	Signal CY5	methylated
5	Signal_Red	unmethylated
6	Signal_Grn	methylated
7	Signal_A	unmethylated
8	Signal_B	methylated

3 Quality Control

The data that are included with the methylumi package are all normal samples from the same tissue. The samples are labeled with the presumed gender. The data are meant to be illustrative of some typical quality-control and sample-labeling problems. In order to get a quick overview of the samples, it is useful to look at an MDS plot of the samples, using only probes on the X chromosome. Since females undergo X-inactivation, they should show something approximating hemi-methylation on that chromosome while males should show very little methylation on the X chromosome.


```
> md <- cmdscale(dist(t(exprs(mldat)[fData(mldat)$CHROMOSOME ==
+ "X", ])), 2)
```

```
> plot(md, pch = c("F", "M")[pData(mldat)$Gender],
+       col = c("red", "blue")[pData(mldat)$Gender])
```



The MDS plot shows that the males and females are quite distinct except for a single male that groups with the females. Upon consultation with the laboratory investigator, the sample was found to be mislabeled. Also, it is worth noting that the males do not cluster nearly as tightly as the females. A quick evaluation of the p-values for detection for the samples will show what the problem is:

```
> avgPval <- colMeans(pvals(mldat))
> par(las = 2)
> barplot(avgPval, ylab = "Average P-Value")
```



So, it is quite obvious that there are two arrays that fail QC with a large percentage of the reporters showing lack of measurement.


It is also possible to use the *qcplot* method in combination with *controlTypes* to examine the QC data in more detail. The control types for the GoldenGate platform are:

```
> controlTypes(mldat)

[1] "ALLELE SPECIFIC EXTENSION"
[2] "EXTENSION GAP"
[3] "FIRST HYBRIDIZATION"
[4] "GENDER"
[5] "NEGATIVE"
[6] "PCR CONTAMINATION"
[7] "SECOND HYBRIDIZATION"
```

Looking more closely at the hybridization controls (“FIRST HYBRIDIZATION”) is telling here:

```
> qcplot(mldat, "FIRST HYBRIDIZATION")
```



So, it appears that the hybridization controls (at least) failed for samples “M_1” and “M_4”, which might help explain why the samples failed.

4 Normalization

The Illumina platform shows a significant dye bias in the two channels which will lead to bias in the estimates of Beta. Therefore, some normalization is required. The function, `normalizeMethyLumiSet` does this normalization. Basically, it looks at the median intensities in the methylated and unmethylated channels at very low and very high beta values and sets these medians equal. Using the transformed unmethylated and methylated values, new beta values are calculated using one of two “map” functions. The `ratio` function is the default and is the same as used by Illumina in the BeadStudio software, but values using the `atan` selection should be similar.

First, a bit of cleanup is needed. The two samples with significantly poorer quality are removed. The gender of the mis-labeled sample is also corrected.

```
> toKeep <- (avgPval < 0.05)
> pData(mldat)$Gender[9] <- "F"
> mldat.norm <- normalizeMethyLumiSet(mldat[,
+   toKeep])
```

5 Example Analysis

As a simple example of an analysis, we can look for the differences between males and females. We already know there is a strong difference based on simple unsupervised methods (the MDS plot). However, methylation is particularly informative for sex differences because females undergo X-inactivation and are, therefore, expected to have one copy of the X chromosome largely methylated.

```
> library(limma)
> dm <- model.matrix(~1 + Gender, data = pData(mldat.norm))
> colnames(dm)

[1] "(Intercept)" "GenderM"

> fit1 <- lmFit(exprs(mldat.norm), dm)
> fit2 <- eBayes(fit1)
> tt <- topTable(fit2, coef = 2, genelist = fData(mldat.norm)[,
+   c("SYMBOL", "CHROMOSOME")], number = 1536)
> x <- aggregate(tt$adj.P.Val, by = list(tt$CHROMOSOME),
+   median)
> colnames(x) <- c("Chromosome", "Median adjusted P-value")

> library(xtable)
> xt <- xtable(x, label = "tab:chromosomepvals",
+   caption = "The median adjusted p-value for each chromosome showing that the X-ch.",
> digits(xt) <- 6
> print(xt, include.rownames = FALSE, align = "cr")
```

Looking at the median adjusted p-values for the resulting differences (calculated using limma), one can quickly see that the X chromosome is, indeed, quite significantly different, on the whole, between males and females. The actual p-values are plotted to show the distribution.

```
> pdf("gender_probes_by_chrom.pdf")
> print(xyplot(-log10(adj.P.Val) ~ CHROMOSOME,
+   tt, ylab = "-log10(Adjusted P-value)",
+   main = "P-values for probes\ndistinguishing males from females"))
> invisible(dev.off())
```

Chromosome	Median adjusted P-value
1	0.997002
10	0.997002
11	0.997002
12	0.997002
13	0.997002
14	0.997002
15	0.997002
16	0.997002
17	0.997002
18	0.997002
19	0.997002
2	0.997002
20	0.997002
21	0.997002
22	0.997002
3	0.997002
4	0.997002
5	0.997002
6	0.997002
7	0.997002
8	0.997002
9	0.997002
X	0.004904


Table 1: The median adjusted p-value for each chromosome showing that the X-chromosome is highly significantly different between males and females

6 sessionInfo

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

- R version 2.10.0 (2009-10-26), x86_64-unknown-linux-gnu
- Locale: LC_CTYPE=en_US.UTF-8, LC_NUMERIC=C, LC_TIME=en_US.UTF-8, LC_COLLATE=en_US.UTF-8, LC_MONETARY=C, LC_MESSAGES=en_US.UTF-8, LC_PAPER=en_US.UTF-8, LC_NAME=C, LC_ADDRESS=C, LC_TELEPHONE=C, LC_MEASUREMENT=en_US.UTF-8, LC_IDENTIFICATION=C
- Base packages: base, datasets, graphics, grDevices, methods, stats, tools, utils
- Other packages: Biobase 2.6.0, lattice 0.17-26, limma 3.2.1, methylumi 1.0.0, xtable 1.5-5

**P-values for probes
distinguishing males from females**



- Loaded via a namespace (and not attached): grid 2.10.0