

msgbsR: an R package to analyse methylation sensitive genotyping by sequencing (MS-GBS) data

Benjamin Mayne

November 4, 2025

Contents

1	Introduction	2
2	Reading data into R	2
3	Confirmation of correct cut sites	3
4	Visualization of read counts	4
5	Differential methylation analysis	6
6	Visualization of cut site locations	6
7	Session Information	8
8	References	9

1 Introduction

Current data analysis tools do not fulfil all experimental designs. For example, GBS experiments using methylation sensitive restriction enzymes (REs), which is also known as methylation sensitive genotyping by sequencing (MS-GBS), is an effective method to identify differentially methylated sites that may not be accessible in other technologies such as microarrays and methyl capture sequencing. However, current data analysis tools do not satisfy the requirements for these types of experimental designs.

Here we present msgbsR, an R package for data analysis of MS-GBS experiments. Read counts and cut sites from a MS-GBS experiment can be read directly into the R environment from a sorted and indexed BAM file(s).

2 Reading data into R

The analysis with the msgbsR pipeline begins with a directory which contains sorted and indexed BAM file(s). msgbsR contains an example data set containing 6 samples from a MS-GBS experiment using the restriction enzyme MspI. In this example the 6 samples are from the prostate of a rat and have been truncated for chromosome 20. 3 of the samples were fed a control diet and the other 3 were fed an experimental high fat diet.

To read in the data directly into the R environment can be done using the `rawCounts()` function, which requires the directory path to where the sorted and indexed files are located and the desired number of threads to be run (Default = 1).

```
> library(msgbsR)
> library(GenomicRanges)
> library(SummarizedExperiment)
> my_path <- system.file("extdata", package = "msgbsR")
> se <- rawCounts(bamFilepath = my_path)
> dim(assay(se))
```

```
[1] 16047      6
```

The result is an `RangedSummarizedExperiment` object containing the read counts. The columns are samples and the rows contain the location of each unique cut sites. Each cut site has been given a unique ID (chr:position:position:strand). The cut site IDs can be turned into a `GRanges` object. Information regarding the samples such as treatment or other groups can be added into the return object as shown below

```
> colData(se) <- DataFrame(Group = c(rep("Control", 3), rep("Experimental", 3)),
+                             row.names = colnames(assay(se)))
```

3 Confirmation of correct cut sites

After the data has been generated into the R environment, the next step is to confirm that the cut sites were the correctly generated sites. In this example, the methylated sensitive restriction enzyme that has been used is MspI which recognizes a 4bp sequence (C/CGG). MspI cuts between the two cytosines when the outside cytosine is methylated.

The first step is to extract the location of the cut sites from `se` and adjust the cut sites such that the region will cover the recognition sequence of MspI. It is important to note that in this example the user must adjust the region over the cut sites specifically for each strand. In other words although the enzyme cuts at C/CGG on the minus strand this would appear as CCG/G. The code below shows how to adjust the postioining of the cut sites to cover the recgnition site on each strand.

```
> cutSites <- rowRanges(se)
> # # Adjust the cut sites to overlap recognition site on each strand
> start(cutSites) <- ifelse(test = strand(cutSites) == '+',
+                           yes = start(cutSites) - 1, no = start(cutSites) - 2)
> end(cutSites) <- ifelse(test = strand(cutSites) == '+',
+                          yes = end(cutSites) + 2, no = end(cutSites) + 1)
```

The object `cutSites` is a `GRanges` object that contains the start and end position of the MspI sequence length around the cut sites. These cut sites can now be checked if the sequence matches the MspI sequence.

`msgbsR` offer two approaches to checking the cut sites. The first approach is to use a `BSSgenome` which can be obtained from Bioconductor. In this example, `BSSgenome.Rnornvegicus.UCSC.rn6` will be used.

```
> library(BSSgenome.Rnornvegicus.UCSC.rn6)
> correctCuts <- checkCuts(cutSites = cutSites, genome = "rn6", seq = "CCGG")
```

If a `BSSgenome` is unavailable for a species of interest, another option to checking the cut sites is to use a fasta file which can be used through the `checkCuts()` function.

The `correctCuts` data object is in the format of a `GRanges` object and contains the correct sites that contained the recognition sequence. These sites can be kept within `se` by using the `subsetByOverlaps` function.

The incorrect MspI cut sites can be filtered out of `datCounts`:

```
> se <- subsetByOverlaps(se, correctCuts)
> dim(assay(se))
```

```
[1] 13983      6
```

`se` now contains the correct cut sites and can now be used in downstream analyses.

4 Visualization of read counts

Before any further downstream analyses with the data, the user may want to filter out samples that did not generate a sufficient number of read counts or cut sites. The `msgbsR` package contains a function which plots the total number of read counts against the total number of cut sites produced per sample. The user can also use the function to visualise if different categories or groups produced varying amount of cut sites or total amount of reads.

To visualize the total number of read counts against the total number of cut sites produced per sample:

```
> plotCounts(se = se, category = "Group")
```

This function generates a plot (Figure 1) where the x axis and y axis represents the total number of reads and the total number of cut sites produced for each sample respectively.



Figure 1: The distribution of the total number of reads and cut sites produced by each sample.

5 Differential methylation analysis

msgbsR utilizes edgeR in order to determine which cut sites are differentially methylated between groups. Since MS-GBS experiments can have multiple groups or conditions msgbsR offers a wrapper function of edgeR (Zhou et al., 2014) tools to automate differential methylation analyses.

To determine which cut sites are differentially methylated between groups:

```
> top <- diffMeth(se = se, category = "Group",  
+               condition1 = "Control", condition2 = "Experimental",  
+               cpmThreshold = 1, thresholdSamples = 1)
```

The top object now contains a data frame of the cut sites that had a CPM > 1 in at least 1 sample and which cut sites are differentially methylated between the two groups.

6 Visualization of cut site locations

The msgbsR package contains a function to allow visualization of the location of the cut sites. Given the lengths of the chromosomes the cut sites can be visualized in a circos plot (Figure 2).

Firstly, define the length of the chromosome.

```
> ratChr <- seqlengths(BSgenome.Rnorvegicus.UCSC.rn6)["chr20"]
```

Extract the differentially methylated cut sites.

```
> my_cuts <- GRanges(top$site[which(top$FDR < 0.05)])
```

To generate a circos plot:

```
> plotCircos(cutSites = my_cuts, seqlengths = ratChr,  
+           cutSite.colour = "red", seqlengths.colour = "blue")
```




Figure 2: A circos plot of chromosome 20 representing cut sites defined by the user.

7 Session Information

This analysis was conducted on:

```
> sessionInfo()
```

R Under development (unstable) (2025-10-20 r88955)

Platform: x86_64-pc-linux-gnu

Running under: Ubuntu 24.04.3 LTS

Matrix products: default

BLAS: /home/biocbuild/bbs-3.23-bioc/R/lib/libRblas.so

LAPACK: /usr/lib/x86_64-linux-gnu/lapack/liblapack.so.3.12.0 LAPACK version 3.12.0

locale:

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

time zone: America/New_York

tzcode source: system (glibc)

attached base packages:

```
[1] stats4      stats      graphics  grDevices  utils      datasets  methods
[8] base
```

other attached packages:

```
[1] BSgenome.Rnorvegicus.UCSC.rn6_1.4.1 BSgenome_1.79.1
[3] rtracklayer_1.71.0                    BiocIO_1.21.0
[5] Biostrings_2.79.1                      XVector_0.51.0
[7] SummarizedExperiment_1.41.0            Biobase_2.71.0
[9] MatrixGenerics_1.23.0                  matrixStats_1.5.0
[11] msgbsR_1.35.0                          GenomicRanges_1.63.0
[13] Seqinfo_1.1.0                          IRanges_2.45.0
[15] S4Vectors_0.49.0                       BiocGenerics_0.57.0
[17] generics_0.1.4
```

loaded via a namespace (and not attached):

```
[1] RColorBrewer_1.1-3      rstudioapi_0.17.1      jsonlite_2.0.0
[4] magrittr_2.0.4          GenomicFeatures_1.63.1 farver_2.1.2
[7] rmarkdown_2.30          vctrs_0.6.5            memoise_2.0.1
[10] Rsamtools_2.27.0        RCurl_1.98-1.17        base64enc_0.1-3
[13] htmltools_0.5.8.1       S4Arrays_1.11.0        progress_1.2.3
[16] curl_7.0.0              SparseArray_1.11.1     Formula_1.2-5
```


[19] easyRNASeq_2.47.0	htmlwidgets_1.6.4	plyr_1.8.9
[22] httr2_1.2.1	cachem_1.1.0	GenomicAlignments_1.47.0
[25] lifecycle_1.0.4	pkgconfig_2.0.3	Matrix_1.7-4
[28] R6_2.6.1	fastmap_1.2.0	digest_0.6.37
[31] colorspace_2.1-2	ShortRead_1.69.1	OrganismDbi_1.53.2
[34] AnnotationDbi_1.73.0	Hmisc_5.2-4	RSQLite_2.4.3
[37] hwriter_1.3.2.1	labeling_0.4.3	filelock_1.0.3
[40] httr_1.4.7	abind_1.4-8	compiler_4.6.0
[43] intervals_0.15.5	withr_3.0.2	bit64_4.6.0-1
[46] htmlTable_2.4.3	S7_0.2.0	backports_1.5.0
[49] BiocParallel_1.45.0	DBI_1.2.3	R.utils_2.13.0
[52] biomaRt_2.67.0	rappdirs_0.3.3	DelayedArray_0.37.0
[55] rjson_0.2.23	tools_4.6.0	foreign_0.8-90
[58] nnet_7.3-20	R.oo_1.27.1	glue_1.8.0
[61] restfulr_0.0.16	grid_4.6.0	checkmate_2.3.3
[64] cluster_2.1.8.1	reshape2_1.4.4	gtable_0.3.6
[67] R.methodsS3_1.8.2	ensembldb_2.35.0	data.table_1.17.8
[70] hms_1.1.4	pillar_1.11.1	stringr_1.6.0
[73] limma_3.67.0	dplyr_1.1.4	LSD_4.1-0
[76] BiocFileCache_3.1.0	lattice_0.22-7	bit_4.6.0
[79] deldir_2.0-4	biovizBase_1.59.0	RBGL_1.87.0
[82] tidyselect_1.2.1	locfit_1.5-9.12	knitr_1.50
[85] gridExtra_2.3	ProtGenerics_1.43.0	ggbio_1.59.0
[88] edgeR_4.9.0	xfun_0.54	genomeIntervals_1.67.0
[91] statmod_1.5.1	stringi_1.8.7	UCSC.utils_1.7.0
[94] lazyeval_0.2.2	yaml_2.3.10	evaluate_1.0.5
[97] codetools_0.2-20	cigarillo_1.1.0	interp_1.1-6
[100] tibble_3.3.0	graph_1.89.0	BiocManager_1.30.26
[103] cli_3.6.5	rpart_4.1.24	dichromat_2.0-0.1
[106] Rcpp_1.1.0	GenomeInfoDb_1.47.0	dbplyr_2.5.1
[109] png_0.1-8	XML_3.99-0.19	parallel_4.6.0
[112] ggplot2_4.0.0	blob_1.2.4	prettyunits_1.2.0
[115] latticeExtra_0.6-31	jpeg_0.1-11	AnnotationFilter_1.35.0
[118] bitops_1.0-9	pwalgn_1.7.0	VariantAnnotation_1.57.0
[121] scales_1.4.0	crayon_1.5.3	rlang_1.1.6
[124] KEGGREST_1.51.0		

8 References

Zhou X, Lindsay H, Robinson MD (2014). Robustly detecting differential expression in RNA sequencing data using observation weights. *Nucleic Acids Research*, 42(11), e91.