

The `DMRcate` package user's guide

Peters TJ, Buckley MJ, Statham A, Pidsley R, Clark SJ, Molloy PL

October 30, 2018

Summary

`DMRcate` extracts the most differentially methylated regions (DMRs) and variably methylated regions (VMRs) from both Whole Genome Bisulphite Sequencing (WGBS) and Illumina® Infinium BeadChip Array samples via kernel smoothing.

```
if (!require("BiocManager"))
  install.packages("BiocManager")
BiocManager::install("DMRcate")
```

Load `DMRcate` into the workspace:

```
library(DMRcate)
```

Illumina® Array Workflow


We now can load in the test data set of beta values. We assume at this point that normalisation and filtering out bad-quality probes via their detection p -values have already been done. Many packages are available for these purposes, including `minfi`, `watermelon` and `methylyumi`. M-values (logit-transform of beta) are preferable to beta values for significance testing via `limma` because of increased sensitivity, but we will retain the beta matrix for visualisation purposes later on.

The TCGA (Cancer Genome Atlas - colorectal cancer) 450K data in `myBetas` only comes from chromosome 20, but `DMRcate` will have no problem taking in the approximately half million probes as input for this pipeline either.

```
data(dmrcatedata)
myMs <- logit2(myBetas)
```

Some of the methylation measurements on the array may be confounded by proximity to SNPs, and cross-hybridisation to other areas of the genome[1, 2]. In particular, probes that are 0, 1, or 2 nucleotides from the methylcytosine of

Figure 1: Beta distribution of 450K probes from publically available data from blood samples of healthy individuals [3] by their proximity to a SNP. “All SNP probes” refers to the 153 113 probes listed by Illumina® whose values may potentially be confounded by a SNP.



interest show a markedly different distribution to those farther away, in healthy tissue (Figure 1).

It is with this in mind that we filter out probes 2 nucleotides or closer to a SNP that have a minor allele frequency greater than 0.05, and the approximately 63,000 [1, 2] cross-reactive probes on either 450K and/or EPIC, so as to reduce confounding. Here we use a combination of *in silico* analysis from [1, 2] and Illumina®'s database of approximately 150,000 potentially SNP-confounded probes, to filter these probes out. About 600 are removed from our M-matrix of approximately 10,000:

```
nrow(snpsall)
## [1] 208568

nrow(myMs)
## [1] 10042

myMs.noSNPs <- rmSNPandCH(myMs, dist=2, mafcut=0.05)
nrow(myMs.noSNPs)
## [1] 9382
```

Next we want to annotate our matrix of M-values with relevant information. We also use the backbone of the `limma` pipeline for differential array analysis to get *t*-statistics changes and, optionally, filter probes by their *fdr*-corrected *p*-value. Here we have 38 patients with 2 tissue samples each taken from them. We want to compare within patients across tissue samples, so we set up our variables for a standard `limma` pipeline, and set `coef=39` in `cpg.annotate` since this corresponds to the phenotype comparison in `design`.

`cpg.annotate()` takes either a data matrix with Illumina probe IDs, or an already prepared `GenomicRatioSet` from `minfi`.

```
patient <- factor(sub("-.*", "", colnames(myMs)))
type <- factor(sub(".*-", "", colnames(myMs)))
design <- model.matrix(~patient + type)
myannotation <- cpg.annotate("array", myMs.noSNPs, what="M", arraytype = "450K",
                             analysis.type="differential", design=design, coef=39)

## Loading required package: IlluminaHumanMethylation450kanno.ilmn12.hg19
## Your contrast returned 6091 individually significant probes. We
## recommend the default setting of pcutoff in dmrcate().

#Or, alternatively
grset <- makeGenomicRatioSetFromMatrix(myMs.noSNPs, array = "IlluminaHumanMethylation450k",
                                       annotation = "ilmn12.hg19", mergeManifest = TRUE,
                                       what = "M")
```

```
myannotation <- cpg.annotate("array", grset, analysis.type="differential", design=design,
                             coef=39)

## Your contrast returned 6091 individually significant probes. We
## recommend the default setting of pcutoff in dmrcate().
```

Now we can find our most differentially methylated regions with `dmrcate()`. For each chromosome, two smoothed estimates are computed: one weighted with `myannotation$stat` and one not, for a null comparison. The two estimates are compared via a Satterthwaite approximation[4], and a significance test is calculated at all hg19 coordinates that an input probe maps to. After *fdr*-correction, regions are then agglomerated from groups of significant probes where the distance to the next consecutive probe is less than `lambda` nucleotides.

```
dmrcoutput <- dmr cate(myannotation, lambda=1000, C=2)

## Fitting chr20...
## Demarcating regions...
## Done!
```

We can convert our DMR list to a `GRanges` object, which uses the `genome` argument to annotate overlapping promoter regions (+/- 2000 bp from TSS). and pass it to `DMR.plot`, which uses the `Gviz` package as a backend for contextualising each DMR. We'll choose one associated with the `GATA5` locus.

```
results.ranges <- extractRanges(dmrcoutput, genome = "hg19")
results.ranges

## GRanges object with 743 ranges and 6 metadata columns:
##      seqnames      ranges strand |   no.cpgs      minfdr
##      <Rle>        <IRanges> <Rle> | <integer> <numeric>
##    946 chr20 61049813-61051915 * |      27          0
##    846 chr20 57424521-57431303 * |      77          0
##    281 chr20 24448859-24452131 * |      21          0
##    252 chr20 21491781-21498921 * |      26 3.34227117448411e-208
##   1005 chr20 61806628-61810795 * |      23          0
##     ...   ...           ...   ... .   ...           ...
##     93 chr20  3026614-3027467 * |       5 6.14728786414342e-17
##    115 chr20  3451292-3451627 * |       8 1.41332210315143e-12
##    108 chr20  3214756-3214926 * |       3 2.34834021398005e-13
##    599 chr20 43729808-43730241 * |       9 1.82254248848938e-16
##    654 chr20 44541804-44542136 * |       2 4.14071181754309e-12
##           Stouffer           maxbetafc           meanbetafc
##           <numeric>           <numeric>           <numeric>
##    946           0 0.477068010284825 0.35455081132964
##    846 2.72563652014737e-268 -0.208426806412668 0.0772863134287036
```

```

##      281 1.43446488803657e-236  0.426352239962428  0.282913167917303
##      252 4.59434687170942e-231  0.438500227266554  0.256988028662146
##     1005 8.95352503528324e-215  0.418203372124153  0.241141203405213
##      ...
##      93  0.00229092578505795  0.0269616535027719  0.0107706371669906
##     115  0.00315251351546967  0.123047303794311  0.0295704660451269
##     108  0.00549122544888556  0.051584188775636  0.0163448981725268
##     599  0.0193040395983799 -0.117089975663635  0.0159797583684682
##     654  0.0194702652563527  0.043268319382771  0.0164531983802135
##
##
##     946
##     846 GNAS-050, GNAS-037, GNAS-058, GNAS-001, GNAS-009, GNAS-049, GNAS-AS1-001, GNAS-030
##     281
##     252
##     1005
##     ...
##     93
##     115
##     108
##     599
##     654
##     -----
##     seqinfo: 1 sequence from an unspecified genome; no seqlengths

```


PLTP-003, PLTP

Now we can plot a significant DMR. We use functionality from the `Gviz` package as a backend for this purpose. We will plot a DMR associated with the `GATA5` locus for the first 6 tumour/normal matched pairs.

```

groups <- c(Tumour="magenta", Normal="forestgreen")
cols <- groups[as.character(type)]
samps <- c(1:6, 38+(1:6))
DMR.plot(ranges=results.ranges, dmr=1, CpGs=myBetas, what="Beta", arraytype = "450K",
         phen.col=cols, genome="hg19", samps=samps)

```



WGBS Workflow

WGBS is a little different. Because the data is represented binomially (that is, by the number of methylated reads followed by the total coverage for that particular CpG site) rather than the continuous distribution afforded by array intensities, we must model the differential methylation signal in a way that respects this. A popular way of doing this is via the beta-binomial distribution. We currently recommend using the method implemented in the DSS package[5], because it uses dispersion shrinkage via a Bayesian framework - similar to `edgeR` for RNA-Seq count data.

The `CpGs` GRanges object contains simulated data for 3 Treatment vs. 3 Control samples for 10^5 CpG sites, generated by WGBSSuite[6].

```
CpGs <- CpGs[1:5000]
CpGs
```

```

## GRanges object with 5000 ranges and 12 metadata columns:
##      seqnames      ranges strand | Treatment1.C Treatment1.cov
##      <Rle> <IRanges> <Rle> |   <integer>   <integer>
## [1] chr1          1      * |          11         13
## [2] chr1          54      * |           9         15
## [3] chr1          58      * |          14         20
## [4] chr1         320      * |          12         15
## [5] chr1         325      * |          10         19
## ...      ...      ...      ... .      ...      ...
## [4996] chr1    1006512      * |          17         20
## [4997] chr1    1006514      * |          12         14
## [4998] chr1    1006527      * |          13         18
## [4999] chr1    1006530      * |          13         16
## [5000] chr1    1006537      * |          18         19
##      Treatment2.C Treatment2.cov Treatment3.C Treatment3.cov
##      <integer>      <integer>      <integer>      <integer>
## [1]          9          14          16          19
## [2]         16          26          18          20
## [3]         19          20          19          27
## [4]         14          14          17          20
## [5]         13          18          14          22
## ...      ...      ...      ...      ...
## [4996]         18          19           9          17
## [4997]         22          25          17          25
## [4998]         16          17          13          16
## [4999]         15          19           9          21
## [5000]         14          15          20          24
##      Control1.C Control1.cov Control2.C Control2.cov Control3.C
##      <integer>      <integer>      <integer>      <integer>      <integer>
## [1]          11          15          16          23          11
## [2]          17          18          10          17          19
## [3]          16          16          12          14          15
## [4]          13          25          15          21          18
## [5]           5          14          16          23          20
## ...      ...      ...      ...      ...
## [4996]         12          16          10          14          12
## [4997]         11          13          16          17          16
## [4998]         12          14          21          22          17
## [4999]         14          17          16          26          12
## [5000]         17          17          17          23          16
##      Control3.cov
##      <integer>
## [1]          14
## [2]          21
## [3]          19

```

```
##      [4]          22
##      [5]          26
##      ...          ...
## [4996]          16
## [4997]          17
## [4998]          17
## [4999]          21
## [5000]          20
## -----
## seqinfo: 1 sequence from an unspecified genome; no seqlengths
```

Note the structure of the metadata columns for this object: samples come in column pairs, with the number of methylated reads followed by the total coverage for that CpG site. Naturally, $\langle \text{sample} \rangle.\text{cov}$ must always be $\geq \langle \text{sample} \rangle.\text{C}$. This structure must be in place in order for downstream tasks such as `DMR.plot()` to be run. Using this structure, we can now extract the methylation and coverage counts, and prepare a `bsseq` object as we would for DSS, and call differentially methylated CpG sites.

```
meth <- as.data.frame(CpGs)[,c(1:2, grep(".C$", colnames(as.data.frame(CpGs))))]
coverage <- as.data.frame(CpGs)[,c(1:2, grep(".cov$", colnames(as.data.frame(CpGs))))]

treat1 <- data.frame(chr=coverage$seqnames, pos=coverage$start,
                    N=coverage$Treatment1.cov, X=meth$Treatment1.C)

treat2 <- data.frame(chr=coverage$seqnames, pos=coverage$start,
                    N=coverage$Treatment2.cov, X=meth$Treatment2.C)

treat3 <- data.frame(chr=coverage$seqnames, pos=coverage$start,
                    N=coverage$Treatment3.cov, X=meth$Treatment3.C)

ctrl1 <- data.frame(chr=coverage$seqnames, pos=coverage$start,
                   N=coverage$Control1.cov, X=meth$Control1.C)

ctrl2 <- data.frame(chr=coverage$seqnames, pos=coverage$start,
                   N=coverage$Control2.cov, X=meth$Control2.C)

ctrl3 <- data.frame(chr=coverage$seqnames, pos=coverage$start,
                   N=coverage$Control3.cov, X=meth$Control3.C)

samples <- list(treat1, treat2, treat3, ctrl1, ctrl2, ctrl3)
samplnames <- sub("\\.*", "", colnames(meth))[-c(1:2)]
obj_bsseq <- makeBSseqData(samples, samplnames)
DSSres <- DMLtest(obj_bsseq, group1=samplnames[1:3], group2=samplnames[4:6], smoothing=FALSE)

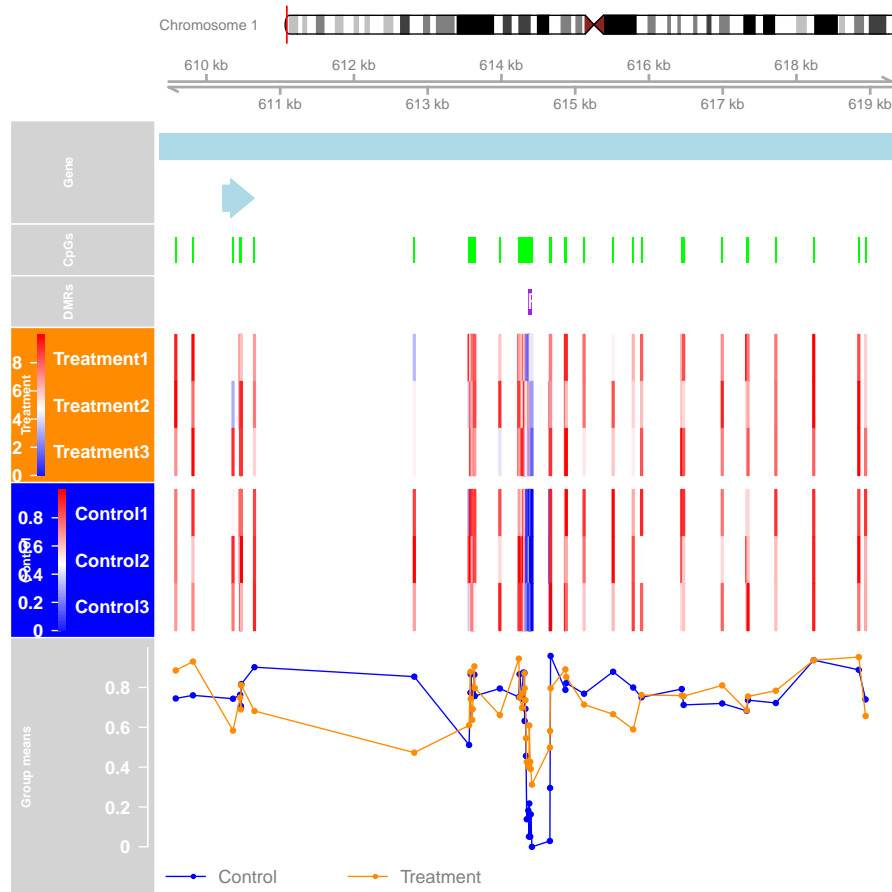
## Estimating dispersion for each CpG site, this will take a while ...
```


We can now enter `DSSres` into the `DMRcate` workflow. Because CpGs are much closer together than they are when represented by Illumina arrays, we will shrink the kernel size by increasing `C`. We will also run this in serial (`mc.cores=1`). If you want to run `dmrcate()` in parallel (1 chromosome per core), please check your processor specifications by running `detectCores()`.

```
wgbsannot <- cpg.annotate("sequencing", DSSres)
wgbs.DMRs <- dmrcate(wgbsannot, lambda = 1000, C = 50, pcutoff = 0.05, mc.cores = 1)

## Fitting chr1...
## Demarcating regions...
## Done!

wgbs.ranges <- extractRanges(wgbs.DMRs, genome = "hg19")
groups <- c(Treatment="darkorange", Control="blue")
cols <- groups[sub("[0-9]", "", sampnames)]
DMR.plot(ranges=wgbs.ranges, dmr=3, CpGs=CpGs, phen.col=cols, genome="hg19")
```



```

sessionInfo()

## R version 3.5.1 Patched (2018-07-12 r74967)
## Platform: x86_64-pc-linux-gnu (64-bit)
## Running under: Ubuntu 16.04.5 LTS
##
## Matrix products: default
## BLAS: /home/biocbuild/bbs-3.8-bioc/R/lib/libRblas.so
## LAPACK: /home/biocbuild/bbs-3.8-bioc/R/lib/libRlapack.so
##
## locale:
## [1] LC_CTYPE=en_US.UTF-8      LC_NUMERIC=C
## [3] LC_TIME=en_US.UTF-8       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
##
## attached base packages:
## [1] splines      stats4      parallel    stats      graphics  grDevices  utils
## [8] datasets    methods    base
##
## other attached packages:
## [1] IlluminaHumanMethylation450kanno.ilmn12.hg19_0.6.0
## [2] DMRcate_1.18.0
## [3] DMRcatedata_1.17.0
## [4] DSS_2.30.0
## [5] bsseq_1.18.0
## [6] minfi_1.28.0
## [7] bumphunter_1.24.0
## [8] locfit_1.5-9.1
## [9] iterators_1.0.10
## [10] foreach_1.4.4
## [11] Biostrings_2.50.0
## [12] XVector_0.22.0
## [13] SummarizedExperiment_1.12.0
## [14] DelayedArray_0.8.0
## [15] BiocParallel_1.16.0
## [16] matrixStats_0.54.0
## [17] Biobase_2.42.0
## [18] GenomicRanges_1.34.0
## [19] GenomeInfoDb_1.18.0
## [20] IRanges_2.16.0
## [21] S4Vectors_0.20.0
## [22] BiocGenerics_0.28.0
##

```

```
## loaded via a namespace (and not attached):
## [1] backports_1.1.2
## [2] Hmisc_4.1-1
## [3] plyr_1.8.4
## [4] lazyeval_0.2.1
## [5] ggplot2_3.1.0
## [6] digest_0.6.18
## [7] ensemblDb_2.6.0
## [8] htmltools_0.3.6
## [9] GO.db_3.7.0
## [10] magrittr_1.5
## [11] checkmate_1.8.5
## [12] memoise_1.1.0
## [13] BSgenome_1.50.0
## [14] cluster_2.0.7-1
## [15] limma_3.38.0
## [16] readr_1.1.1
## [17] annotate_1.60.0
## [18] R.utils_2.7.0
## [19] siggenes_1.56.0
## [20] prettyunits_1.0.2
## [21] colorspace_1.3-2
## [22] blob_1.1.1
## [23] BiasedUrn_1.07
## [24] dplyr_0.7.7
## [25] crayon_1.3.4
## [26] RCurl_1.95-4.11
## [27] genefilter_1.64.0
## [28] bindr_0.1.1
## [29] GEOquery_2.50.0
## [30] IlluminaHumanMethylationEPICmanifest_0.3.0
## [31] VariantAnnotation_1.28.0
## [32] survival_2.43-1
## [33] glue_1.3.0
## [34] ruv_0.9.7
## [35] registry_0.5
## [36] gtable_0.2.0
## [37] zlibbioc_1.28.0
## [38] IlluminaHumanMethylationEPICanno.ilm10b4.hg19_0.6.0
## [39] Rhdf5lib_1.4.0
## [40] HDF5Array_1.10.0
## [41] scales_1.0.0
## [42] DBI_1.0.0
## [43] rngtools_1.3.1
## [44] bibtex_0.4.2
```

```
## [45] Rcpp_0.12.19
## [46] xtable_1.8-3
## [47] progress_1.2.0
## [48] htmlTable_1.12
## [49] foreign_0.8-71
## [50] bit_1.1-14
## [51] mclust_5.4.1
## [52] preprocessCore_1.44.0
## [53] Formula_1.2-3
## [54] missMethyl_1.16.0
## [55] htmlwidgets_1.3
## [56] httr_1.3.1
## [57] RColorBrewer_1.1-2
## [58] acepack_1.4.1
## [59] pkgconfig_2.0.2
## [60] reshape_0.8.8
## [61] XML_3.98-1.16
## [62] R.methodsS3_1.7.1
## [63] Gviz_1.26.0
## [64] nnet_7.3-12
## [65] tidyselect_0.2.5
## [66] rlang_0.3.0.1
## [67] AnnotationDbi_1.44.0
## [68] munsell_0.5.0
## [69] tools_3.5.1
## [70] RSQLite_2.1.1
## [71] evaluate_0.12
## [72] stringr_1.3.1
## [73] org.Hs.eg.db_3.7.0
## [74] knitr_1.20
## [75] bit64_0.9-7
## [76] beanplot_1.2
## [77] methylumi_2.28.0
## [78] purrr_0.2.5
## [79] AnnotationFilter_1.6.0
## [80] bindrcpp_0.2.2
## [81] nlme_3.1-137
## [82] doRNG_1.7.1
## [83] nor1mix_1.2-3
## [84] R.oo_1.22.0
## [85] xml2_1.2.0
## [86] biomaRt_2.38.0
## [87] compiler_3.5.1
## [88] rstudioapi_0.8
## [89] curl_3.2
```

```
## [90] statmod_1.4.30
## [91] tibble_1.4.2
## [92] stringi_1.2.4
## [93] highr_0.7
## [94] GenomicFeatures_1.34.0
## [95] lattice_0.20-35
## [96] ProtGenerics_1.14.0
## [97] Matrix_1.2-14
## [98] IlluminaHumanMethylation450kmanifest_0.4.0
## [99] permute_0.9-4
## [100] multtest_2.38.0
## [101] pillar_1.3.0
## [102] data.table_1.11.8
## [103] bitops_1.0-6
## [104] rtracklayer_1.42.0
## [105] R6_2.3.0
## [106] latticeExtra_0.6-28
## [107] gridExtra_2.3
## [108] codetools_0.2-15
## [109] dichromat_2.0-0
## [110] MASS_7.3-51
## [111] gtools_3.8.1
## [112] assertthat_0.2.0
## [113] rhdf5_2.26.0
## [114] openssl_1.0.2
## [115] pkgmaker_0.27
## [116] withr_2.1.2
## [117] GenomicAlignments_1.18.0
## [118] Rsamtools_1.34.0
## [119] GenomeInfoDbData_1.2.0
## [120] hms_0.4.2
## [121] quadprog_1.5-5
## [122] grid_3.5.1
## [123] rpart_4.1-13
## [124] tidyr_0.8.2
## [125] base64_2.0
## [126] DelayedMatrixStats_1.4.0
## [127] illuminaio_0.24.0
## [128] biovizBase_1.30.0
## [129] base64enc_0.1-3
```

References

- [1] Pidsley R, Zotenko E, Peters TJ, Lawrence MG, Risbridger GP, Molloy P, Van Dijk S, Muhlhausler B, Stirzaker C, Clark SJ. Critical evaluation of the Illumina MethylationEPIC BeadChip microarray for whole-genome DNA methylation profiling. *Genome Biology*. 2016 17(1), 208.
- [2] Chen YA, Lemire M, Choufani S, Butcher DT, Grafodatskaya D, Zanke BW, Gallinger S, Hudson TJ, Weksberg R. Discovery of cross-reactive probes and polymorphic CpGs in the Illumina Infinium HumanMethylation450 microarray. *Epigenetics*. 2013 Jan 11;8(2).
- [3] Heyn H, Li N, Ferreira HJ, Moran S, Pisano DG, Gomez A, Esteller M. Distinct DNA methylomes of newborns and centenarians. *Proceedings of the National Academy of Sciences*. 2012 **109**(26), 10522-7.
- [4] Satterthwaite FE. An Approximate Distribution of Estimates of Variance Components., *Biometrics Bulletin*. 1946 **2**: 110-114
- [5] Feng H, Conneely KN, Wu H. A Bayesian hierarchical model to detect differentially methylated loci from single nucleotide resolution sequencing data. *Nucleic Acids Research*. 2014 **42**(8), e69.
- [6] Rackham, OJL, Dellaportas P, Petretto E, Bottolo, L. WGBSSuite: Simulating Whole Genome Bisulphite Sequencing data and benchmarking differential DNA methylation analysis tools. *Bioinformatics* 2015. (Oxford, England), (March).