

# An Introduction to the *REMP* Package

Yinan Zheng

September 25, 2017

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Installation</b>	<b>1</b>
<b>3</b>	<b>REMP: Repetitive Element Methylation Prediction</b>	<b>1</b>
3.1	Groom methylation data . . . . .	2
3.2	Prepare annotation data . . . . .	2
3.3	Run prediction . . . . .	3
3.4	Plot prediction . . . . .	7
<b>4</b>	<b>Session Information</b>	<b>8</b>

## 1 Introduction

*REMP* predicts DNA methylation of locus-specific repetitive elements (RE) by learning surrounding genetic and epigenetic information. *REMP* provides genomewide single-base resolution of DNA methylation on RE that are difficult to measure using array-based or sequencing-based platforms, which enables epigenome-wide association study (EWAS) and differentially methylated region (DMR) analysis on RE.

## 2 Installation

Install *REMP* (release version):

```
> source("http://bioconductor.org/biocLite.R")
> biocLite("REMP")
```

To install devel version:

```
> library(devtools)
> install_github("YinanZheng/REMP")
```

Load *REMP* into the workspace:

```
> library(REMP)
```

## 3 REMP: Repetitive Element Methylation Prediction

Currently *REMP* supports Human (hg19, build 37) Alu and LINE-1 (L1) repetitive element (RE) methylation prediction using Illumina 450k or EPIC array.

### 3.1 Groom methylation data

Appropriate data preprocessing including quality control and normalization of methylation data are recommended before running *REMP*. Many packages are available to carry out these data preprocessing steps, for example, *minfi*, *wateRmelon*, and *methylumi*.

*REMP* is trying to minimize the requirement of the methylation data format. User can maintain the methylation data in *RatioSet* or *GenomicRatioSet* object offered by *minfi*, *data.table*, *data.frame*, *DataFrame*, or *matrix*. User can input either beta value or M-value. There are only two basic requirements of the methylation data:

1. Each row should represent CpG probe and each column should represent sample.
2. The row names should indicate Illumina probe ID (i.e. cg00000029).

However, there are some other common data issues that may prevent *REMP* from running correctly. For example, if the methylation data are in beta value and contain zero methylation values, logit transformation (to create M-value) will create negative infinite value; or the methylation data contain NA, Inf, or NaN data. To tackle these potential issues, *REMP* includes a handy function `groomMethy` which can help detect and fix these issues. We highly recommend to take advantage of this function:

```
> # Get GM12878 methylation data (450k array)
> GM12878_450k <- getGM12878('450k')
> GM12878_450k <- groomMethy(GM12878_450k, verbose = TRUE)
> GM12878_450k
```

```
class: RatioSet
dim: 482421 1
metadata(0):
assays(2): Beta M
rownames(482421): cg00000029 cg00000108 ...
  cg27666046 cg27666123
rowData names(0):
colnames(1): GM12878
colData names(0):
Annotation
  array: IlluminaHumanMethylation450k
  annotation: ilmn12.hg19
Preprocessing
  Method: NA
  minfi version: NA
  Manifest version: NA
```

For zero beta values, `groomMethy` will replace them with smallest non-zero beta value. For NA/NaN/Inf values, `groomMethy` will treat them as missing values and then apply KNN-imputation to complete the dataset. If the imputed value is out of the original range (which is possible when `imputebyrow = FALSE`), mean value will be used instead. Warning: imputed values for multimodal distributed CpGs (across samples) may not be correct. Please check package *ENmix* to identify the CpGs with multimodal distribution.

### 3.2 Prepare annotation data

To run *REMP* for RE methylation prediction, user first needs to prepare some annotation datasets. The function `initREMP` is designed to do the job.

Suppose user will predict Alu methylation using Illumina 450k array data:

```
> data(Alu.demo)
> remparcel <- initREMP(arrayType = "450k", REtype = "Alu",
+                       RE = Alu.demo, ncore = 1)
> remparcel
```

```
REMPParcel object
RE type: Alu
Illumina platform: 450k
Valid (max) RE-CpG flanking window size: 1200
Number of RE: 500
Number of RE-CpG: 5039
```

For demonstration, we only use 500 selected Alu sequence dataset which comes along with the package (`Alu.demo`). We specify `RE = Alu.demo`, so that the annotation dataset will be generated for the 500 selected Alu sequences. In real-world prediction, specifying `RE` is not necessary, as the function will pull up the complete RE sequence dataset from package *AnnotationHub*.

All data are stored in the *REMPParcel* object. It is recommended to specify a working directory so that the data generated can be preserved for later use:

```
> saveParcel(remparcel)
```

Without specifying working directory using option `work.dir`, the annotation dataset will be created under the temporal directory `tempdir()` by default. User can also turn on the `export` parameter in `initREMP` to save the data automatically.

### 3.3 Run prediction

Once the annotation data are ready, user can pass the annotation data parcel to `remp` for prediction:

```
> remp.res <- remp(GM12878_450k, REtype = 'Alu',
+                 parcel = remparcel, ncore = 1, seed = 777)
```

If `parcel` is missing, `remp` will then try to search the *REMPParcel* data file in the directory indicated by `work.dir` (use `tempdir()` if not specified).

By default, `remp` uses Random Forest (`method = 'rf'`) model (package *randomForest*) for prediction. Random Forest model is recommended because it offers more accurate prediction results and it automatically enables Quantile Regression Forest (Nicolai Meinshausen, 2006) for prediction reliability evaluation. `remp` constructs 19 predictors to carry out the prediction. For Random Forest model, the tuning parameter `param = 6` (i.e. `mtry` in *randomForest*) indicates how many predictors will be randomly selected for building the individual trees. The performance of random forest model is often relatively insensitive to the choice of `mtry`. Therefore, auto-tune will be turned off using random forest and `mtry` will be set to one third of the total number of predictors. It is recommended to specify a seed for reproducible prediction results.

`remp` will return a *REMPset* object, which inherits Bioconductor's *RangedSummarizedExperiment* class:

```
> remp.res

class: REMProduct
dim: 4808 1
metadata(8): REannotation RECpG ... GeneStats Seed
assays(3): rempB rempM rempQC
rownames: NULL
rowData names(1): RE.Index
colnames(1): GM12878
colData names(1): mtry

> # Display more detailed information
> details(remp.res)
```

```
RE type: Alu
Methylation profiling platform: 450k
Flanking window size: 1000
```

Prediction model: Random Forest  
QC model: Quantile Regression Forest  
Predicted 4808 CpG sites in 500 Alu

Number of predicted CpGs by chromosome:

chr1	chr2	chr3	chr4	chr5	chr6	chr7	chr8
449	276	293	131	179	397	292	102
chr9	chr10	chr11	chr12	chr13	chr14	chr15	chr16
98	148	254	310	66	127	133	333
chr17	chr18	chr19	chr20	chr21	chr22		
295	81	674	66	37	67		

Coverage information:

There are 500 profiled Alu by Illumina array.  
There are 481 RE-CpGs that have neighboring profiled CpGs are used for model training.  
In total, REMP predicts 500 Alu (4808 RE-CpG).  
Gene coverage by predicted Alu (out of total refSeq Gene):  
492 (1.96%) total genes;  
413 (2.15%) protein-coding genes;  
117 (1.59%) non-coding RNA genes.

Distribution of predicted methylation value (beta value):

Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
0.02877	0.47252	0.66134	0.59484	0.75076	0.91730

Distribution of prediction reliability score:

Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
0.7534	1.4620	1.7652	1.7943	2.0543	5.5620

Prediction results can be obtained by accessors:

```
> # Predicted RE-CpG methylation value (Beta value)
> rempB(remp.res)
```

DataFrame with 4808 rows and 1 column

```
GM12878
<numeric>
1 0.7626404
2 0.7847527
3 0.7938839
4 0.7963254
5 0.7987047
...
4804 0.4579916
4805 0.4535980
4806 0.4639225
4807 0.4860950
4808 0.4947872
```

```
> # Predicted RE-CpG methylation value (M value)
> rempM(remp.res)
```

DataFrame with 4808 rows and 1 column

```
GM12878
```

```

      <numeric>
1      1.683928
2      1.866243
3      1.945471
4      1.967092
5      1.988348
...      ...
4804 -0.24299384
4805 -0.26854843
4806 -0.20855755
4807 -0.08026338
4808 -0.03008277

```

```

> # Genomic location information of the predicted RE-CpG
> # Function inherit from class 'RangedSummarizedExperiment'
> rowRanges(remp.res)

```

GRanges object with 4808 ranges and 1 metadata column:

	seqnames	ranges	strand	RE.Index
	<Rle>	<IRanges>	<Rle>	<Rle>
[1]	chr1	[943927, 943928]	-	Alu_0000527
[2]	chr1	[943935, 943936]	-	Alu_0000527
[3]	chr1	[943968, 943969]	-	Alu_0000527
[4]	chr1	[943974, 943975]	-	Alu_0000527
[5]	chr1	[943991, 943992]	-	Alu_0000527
...	...	...	...	...
[4804]	chr22	[42095154, 42095155]	-	Alu_1170175
[4805]	chr22	[42095161, 42095162]	-	Alu_1170175
[4806]	chr22	[42095170, 42095171]	-	Alu_1170175
[4807]	chr22	[42095198, 42095199]	-	Alu_1170175
[4808]	chr22	[42095214, 42095215]	-	Alu_1170175

-----  
seqinfo: 93 sequences from an unspecified genome; no seqlengths

```

> # Standard error-scaled permutation importance of predictors
> imp(remp.res)

```

DataFrame with 19 rows and 1 column

```

      GM12878
      <numeric>
RE.score      8.127929
RE.Length     7.801685
RE.CpG.density 6.117575
RE.InNM       4.164557
RE.InNR       1.590008
...      ...
distance.min2 12.877820
Methy.min     29.910507
Methy.min2    10.624580
Methy.mean    12.499192
Methy.std     4.422039

```

```

> # Retrieve seed number used for the results
> metadata(remp.res)$Seed

```

```
[1] 777
```

Trim off less reliable predicted results:

```
> # Any predicted CpG values with quality score < threshold (default = 1.7) will be replaced with NA. CpG
> # For mechanism study, more stringent cutoff is recommended.
> remp.res <- trim(remp.res)
> details(remp.res)
```

```
RE type: Alu
Methylation profiling platform: 450k
Flanking window size: 1000
Prediction model: Random Forest - trimmed (1.7)
QC model: Quantile Regression Forest
Predicted 2151 CpG sites in 393 Alu
```

Number of predicted CpGs by chromosome:

```
chr1 chr2 chr3 chr4 chr5 chr6 chr7 chr8
 230  120  140   76   69  188  135   46
```

```
chr9 chr10 chr11 chr12 chr13 chr14 chr15 chr16
  40   84   82  115   23   99   31  186
```

```
chr17 chr18 chr19 chr20 chr21 chr22
 144   27  250   37    5   24
```

Coverage information:

There are 393 profiled Alu by Illumina array.

There are 355 RE-CpGs that have neighboring profiled CpGs are used for model training.

In total, REMP predicts 393 Alu (2151 RE-CpG).

Gene coverage by predicted Alu (out of total refSeq Gene):

```
377 (1.5%) total genes;
312 (1.63%) protein-coding genes;
87 (1.18%) non-coding RNA genes.
```

Distribution of predicted methylation value (beta value):

```
Min. 1st Qu. Median Mean 3rd Qu. Max.
0.06723 0.67452 0.74424 0.70393 0.79153 0.91599
```

Distribution of prediction reliability score:

```
Min. 1st Qu. Median Mean 3rd Qu. Max.
0.7534 1.2923 1.4321 1.4182 1.5657 1.6998
```

To add genomic regions annotation of the predicted REs:

```
> # By default gene symbol annotation will be added
> remp.res <- decodeAnnot(remp.res, ncore = 1)
> annotation(remp.res)
```

Seven genomic region indicators will be added to the annotation data in the input *REMPProduct* object:

- InNM: in protein-coding genes (overlap with refSeq gene's "NM" transcripts + 2000 bp upstream of the transcription start site (TSS))
- InNR: in noncoding RNA genes (overlap with refSeq gene's "NR" transcripts + 2000 bp upstream of the TSS)
- InTSS: in flanking region of 2000 bp upstream of the TSS. Default upstream limit is 2000 bp, which can be modified globally using `remp_options`

- In5UTR: in 5'untranslated regions (UTRs)
- InCDS: in coding DNA sequence regions
- InExon: in exon regions
- In3UTR: in 3'UTRs

Note that intron region and intergenic region information can be derived from the above genomic region indicators: if "InNM" and/or "InNR" is not missing but "InTSS", "In5UTR", "InExon", and "In3UTR" are missing, then the RE is strictly located within intron region; if all indicators are missing, then the RE is strictly located in intergenic region.

### 3.4 Plot prediction

Make a density plot of the predicted methylation (beta values):

```
> plot(remp.res, main = "Alu methylation (GM12878)", col = "blue")
```

## 4 Session Information

R version 3.4.1 (2017-06-30)  
Platform: x86\_64-pc-linux-gnu (64-bit)  
Running under: Ubuntu 16.04.3 LTS

Matrix products: default  
BLAS: /home/biocbuild/bbs-3.5-bioc/R/lib/libRblas.so  
LAPACK: /home/biocbuild/bbs-3.5-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] parallel stats4 stats graphics grDevices  
[6] utils datasets methods base

other attached packages:  
[1] REMP\_1.0.7  
[2] IlluminaHumanMethylationEPICanno.ilm10b2.hg19\_0.6.0  
[3] IlluminaHumanMethylation450kanno.ilmn12.hg19\_0.6.0  
[4] minfi\_1.22.1  
[5] bumpHunter\_1.16.0  
[6] locfit\_1.5-9.1  
[7] iterators\_1.0.8  
[8] foreach\_1.4.3  
[9] Biostrings\_2.44.2  
[10] XVector\_0.16.0  
[11] SummarizedExperiment\_1.6.5  
[12] DelayedArray\_0.2.7  
[13] matrixStats\_0.52.2  
[14] Biobase\_2.36.2  
[15] GenomicRanges\_1.28.5  
[16] GenomeInfoDb\_1.12.2  
[17] IRanges\_2.10.3  
[18] S4Vectors\_0.14.4  
[19] BiocGenerics\_0.22.0  
[20] knitr\_1.17

loaded via a namespace (and not attached):  
[1] colorspace\_1.3-2  
[2] class\_7.3-14  
[3] siggenes\_1.50.0  
[4] mclust\_5.3  
[5] base64\_2.0  
[6] DRR\_0.0.2  
[7] bit64\_0.9-7  
[8] interactiveDisplayBase\_1.14.0  
[9] AnnotationDbi\_1.38.2



[10] proclim\_1.6.1  
[11] lubridate\_1.6.0  
[12] ranger\_0.8.0  
[13] codetools\_0.2-15  
[14] splines\_3.4.1  
[15] doParallel\_1.0.10  
[16] impute\_1.50.1  
[17] robustbase\_0.92-7  
[18] RcppRoll\_0.2.2  
[19] Rsamtools\_1.28.0  
[20] caret\_6.0-77  
[21] annotate\_1.54.0  
[22] ddalpaha\_1.2.1  
[23] kernlab\_0.9-25  
[24] settings\_0.2.4  
[25] shiny\_1.0.5  
[26] compiler\_3.4.1  
[27] httr\_1.3.1  
[28] assertthat\_0.2.0  
[29] Matrix\_1.2-11  
[30] lazyeval\_0.2.0  
[31] limma\_3.32.7  
[32] htmltools\_0.3.6  
[33] tools\_3.4.1  
[34] bindrcpp\_0.2  
[35] gtable\_0.2.0  
[36] glue\_1.1.1  
[37] GenomeInfoDbData\_0.99.0  
[38] reshape2\_1.4.2  
[39] dplyr\_0.7.3  
[40] doRNG\_1.6.6  
[41] Rcpp\_0.12.12  
[42] multtest\_2.32.0  
[43] preprocessCore\_1.38.1  
[44] nlme\_3.1-131  
[45] rtracklayer\_1.36.4  
[46] timeDate\_3012.100  
[47] gower\_0.1.2  
[48] stringr\_1.2.0  
[49] mime\_0.5  
[50] rngtools\_1.2.4  
[51] XML\_3.98-1.9  
[52] beanplot\_1.2  
[53] org.Hs.eg.db\_3.4.1  
[54] AnnotationHub\_2.8.2  
[55] DEoptimR\_1.0-8  
[56] zlibbioc\_1.22.0  
[57] MASS\_7.3-47  
[58] scales\_0.5.0  
[59] ipred\_0.9-6  
[60] BSgenome\_1.44.2  
[61] BiocInstaller\_1.26.1  
[62] GEOquery\_2.42.0  
[63] RColorBrewer\_1.1-2

[64] curl\_2.8.1  
[65] yaml\_2.1.14  
[66] memoise\_1.1.0  
[67] ggplot2\_2.2.1  
[68] pkgmaker\_0.22  
[69] biomaRt\_2.32.1  
[70] rpart\_4.1-11  
[71] reshape\_0.8.7  
[72] stringi\_1.1.5  
[73] RSQLite\_2.0  
[74] genefilter\_1.58.1  
[75] GenomicFeatures\_1.28.5  
[76] BiocParallel\_1.10.1  
[77] lava\_1.5  
[78] rlang\_0.1.2  
[79] pkgconfig\_2.0.1  
[80] bitops\_1.0-6  
[81] nor1mix\_1.2-3  
[82] evaluate\_0.10.1  
[83] lattice\_0.20-35  
[84] purrr\_0.2.3  
[85] bindr\_0.1  
[86] GenomicAlignments\_1.12.2  
[87] recipes\_0.1.0  
[88] CVST\_0.2-1  
[89] bit\_1.1-12  
[90] BSgenome.Hsapiens.UCSC.hg19\_1.4.0  
[91] plyr\_1.8.4  
[92] magrittr\_1.5  
[93] R6\_2.2.2  
[94] dimRed\_0.1.0  
[95] DBI\_0.7  
[96] withr\_2.0.0  
[97] survival\_2.41-3  
[98] RCurl\_1.95-4.8  
[99] nnet\_7.3-12  
[100] tibble\_1.3.4  
[101] grid\_3.4.1  
[102] data.table\_1.10.4  
[103] blob\_1.1.0  
[104] ModelMetrics\_1.1.0  
[105] digest\_0.6.12  
[106] xtable\_1.8-2  
[107] httpuv\_1.3.5  
[108] illuminaio\_0.18.0  
[109] openssl\_0.9.7  
[110] munsell\_0.4.3  
[111] registry\_0.3  
[112] quadprog\_1.5-5