

compEpiTools

Mattia Pelizzola, Kamal Kishore

October 17, 2016

Contents

1	Overview	1
2	Counting reads in GRanges	1
3	Annotation of genomic regions	3
4	Functional annotation	4
5	Visualization	5
6	Session Information	9

1 Overview

In this document you can find a brief tutorial on the *compEpiTools* package, a toolkit for the integrative analysis of epigenomics data, which can be complemented with the *methyPipe* package to include support on DNA methylation data. *compEpiTools* offers multiple functionalities covering five main areas: Counting reads in genomic regions, Annotation of genomic regions, Functional annotation, Visualization, and Other (see `?compEpiTools` at the R prompt for a brief overview). Many of these methods and functions concern topics of interest in epigenomics, including: identification of enhancer and matching with putative target regions, identification of long non coding RNAs based on chromatin features, computation of PolIII stalling index, determination of promoter CpG content etc. Finally, functionalities are available to integrate and display heterogenous data-types across multiple genomic regions.

2 Counting reads in GRanges

These are mostly *GRanges* methods facilitating several common operations concerning **counting** in genomic regions aligned reads stored in BAM files. In this example a small BAM file is used to compute the base-level coverage using `GRbaseCoverage`. Next the count of reads is determined for a whole genomic region (`GRcoverage`) and dividing it in 5 equally sized bins (`GRcoverageInbins`). Finally the summit, i.e. the position of maximum coverage is identified and highlighted in the plot of base-level coverage for one of these regions.

```
library(compEpiTools)
require(TxDb.Mmusculus.UCSC.mm9.knownGene)
require(org.Mm.eg.db)
txdb <- TxDb.Mmusculus.UCSC.mm9.knownGene
```

```

bampath <- system.file("extdata", "ex1.bam", package="Rsamtools")
ir <- IRanges(start=c(1000, 100), end=c(2000, 1000))
gr <- GRanges(seqnames=Rle(c('seq1','seq2')), ranges=ir)
res <- GRbaseCoverage(Object=gr, bam=bampath)
GRcoverage(Object=gr, bam=bampath, Nnorm=FALSE, Snorm=FALSE)[1]

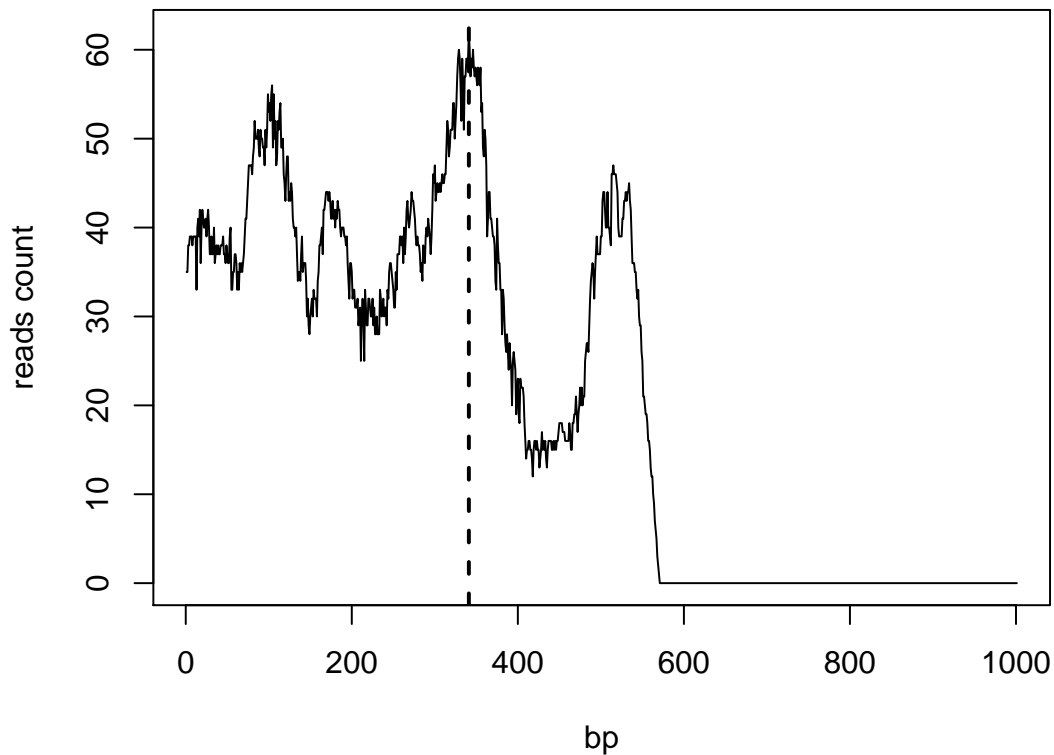
## [1] 20359

GRcoverageInbins(Object=gr, bam=bampath, Nnorm=FALSE, Snorm=FALSE, Nbins=5)[1,]

## [1] 8148 7969 4242 0 0

summit <- GRcoverageSummit(Object=gr, bam=bampath)
plot(res[[1]], type='l', xlab='bp', ylab='reads count')
abline(v=start(summit[1])-start(gr[1])+1, lty=2, lwd=2)

```



Apart from these general functionalities, two specific applications of interest in epigenetic are the determination of ChIP-seq enrichment and of the PolII stalling index.

In (epi)genomics ChIP-seq is an experimental method to identify the genomic regions which are bound by a DNA binding protein (if any), such as transcription factor or an histone mark. Such experiments are typically conducted comparing a sample in which an antibody specific for the target protein was used to enrich for the bound genomic DNA (ChIP), which is then compared to a control sample where the antibody was not used (input). Comparing the signal in the ChIP and input samples allows identifying

regions of significant enrichment, the so called ChIP-seq peaks. The peaks can also be quantitatively scored to give a measure to the binding intensity of the factor of interest. The `GRenrichment` method can be used to determine this **enrichment**. First the number of reads in the ChIP sample within the peak is determined and normalized by the total number of reads aligned in the BAM file (ChIPw), then the same is done for the input sample (inputw). Finally the $\log_2(\text{ChIPw} - \text{inputw})$ is determined which was shown to be approximatively linearly correlated to the log of the qPCR enrichment, typically used for validation.

The **PolII stalling index** is a measure that is commonly determined to estimate the degree of stalling of the Polymeras II in the region around the Transcription Start Sites (TSS), which is somehow related to the elongation rate, the speed at which the PolII is actively elongating and transcribing the open reading frame. The stalling index is determined as the ratio of the PolII ChIP-seq signal in the TSS region compared to the gene body, and can be determined using the `stallingIndex` function and is typically visualized with cumulative plot generated by the `plotStallingIndex` function.

3 Annotation of genomic regions

These methods can be used to assign to genomic ranges annotations concerning gene annotation and user-defined databases. Specifically, the `TSS` and `distanceFromTSS` methods can be used to determine the Transcription Start Sites positions for all transcript in a *TranscriptDb* database and to compute the distance between a set of genomic regions and the most proximal TSS. The `GRangesInPromoters` is a convenience wrapper to subset a *GRanges* object keeping only ranges overlapping (or not) with gene promoter regions. All these methods plus `GRmidpoint` are used by `GRannotateSimple` to partition a set of genomic ranges into those overlapping with promoters, intragenic and intergenic regions.

Apart from genes and promoters, numerous other genomic annotation resources are available for example in the UCSC table browser. These or other user-defined regions of interest, together with genes and promoter annotations can be used with the `GRannotate` method to obtain the complete annotation of a set of genomic ranges. This is illustrated in the following example, where CpG Islands (CGIs) are considered in addition to genes and promoters. In the output each genomic range is put in the context of the nearest TSS ('nearest' columns), the specific location in which it falls ('location' columns) and the overlap with CGIs:

```
TSSpos <- TSS(txdb)
gr <- TSSpos[1:5]
start(gr) <- start(gr) - 1000
end(gr) <- end(gr) - 600
mcols(gr) <- NULL
# retrieving CGI mm9 islands from UCSC annotation tables
cgipath <- system.file("extdata", "CGIgr_mm9.Rdata", package="compEpiTools")
load(cgipath)
res <- GRannotate(Object=GRmidpoint(gr), txdb=txdb, EG2GS=org.Mm.eg.db,
                  upstream=2000, downstream=1000, userAnn=GRangesList(CGI=CGIgr_mm9))
show(res)
```

```
## GRanges object with 5 ranges and 9 metadata columns:
##      seqnames      ranges strand | nearest_tx_name
##      <Rle>        <IRanges> <Rle> | <character>
##  18777 chr1 [4797174, 4797174] + | uc007afg.1
##  18777 chr1 [4797174, 4797174] + | uc007afg.1
##  21399 chr1 [4846975, 4846975] + | uc007afi.2
##  21399 chr1 [4846975, 4846975] + | uc007afi.2
```

```

## 21399 chr1 [4847609, 4847609] + | uc007afi.2
## distance_fromTSS nearest_gene_id nearest_gene_symbol
## <integer> <character> <character>
## 18777 799 18777 Lypla1
## 18777 799 18777 Lypla1
## 21399 799 21399 Tcea1
## 21399 799 21399 Tcea1
## 21399 165 21399 Tcea1
## location location_tx_id
## <character> <character>
## 18777 promoter;promoter uc007afg.1;uc007afh.1
## 18777 promoter;promoter uc007afg.1;uc007afh.1
## 21399 promoter;promoter;promoter uc007afi.2;uc011wht.1;uc011whu.1
## 21399 promoter;promoter;promoter uc007afi.2;uc011wht.1;uc011whu.1
## 21399 promoter;promoter;promoter uc007afi.2;uc011wht.1;uc011whu.1
## location_gene_id location_gene_symbol CGI
## <character> <character> <numeric>
## 18777 18777;18777 Lypla1;Lypla1 0
## 18777 18777;18777 Lypla1;Lypla1 0
## 21399 21399;21399;21399 Tcea1;Tcea1;Tcea1 0
## 21399 21399;21399;21399 Tcea1;Tcea1;Tcea1 0
## 21399 21399;21399;21399 Tcea1;Tcea1;Tcea1 1
## -----
## seqinfo: 35 sequences (1 circular) from mm9 genome

```

The `makeGtfFromDb` method can be used to export a gene or transcript level GTF file to be used with standard aligners, such as TopHat or reads counting tools, such as the one available in HTSeq, in order to be consistent with analyses performed within R and Bioconductor.

4 Functional annotation

A number of functions and methods is available to define epigenetically relevant functional annotations. The `enhancers` allows pointing to putative enhancers based on H3K4me1 (thus pointing to enhancers which could be either active or poised) or H3K27ac (thus pointing to active enhancers) marks. To this purpose, distal peaks of these marks laying outside gene promoters are identified, and required to not overlap with CpG Islands to avoid peaks matching to potentially unannotated promoter regions.

Using the `matchEnhancers` it is possible to match enhancers with putative targets sites. Target sites could either be TSS or transcription factor binding sites localized at the level of promoters. Constraints are given based on a minimum or maximum distance between the enhancer and the target site. At the same time no additional TSS have to be present in between the enhancer and the putative target site (identification of 'direct' enhancers). This does not apply if those TSS belong to isoforms of the same gene. This method returns: (i) a set of reference regions without any interacting direct enhancers, (ii) a set of enhancers sites having putative target regions, and (iii) those of putative target regions under control of enhancer sites. Lists (ii) and (iii) are ordered so that they can be immediately matched. Finally, if TF binding is provided, these two lists will be further divided considering only TF-bound TSS (target regions) which are bound by to enhancer bound from the same TF or not. This could set the foundation for further explorative analyses on nextworks of enhancers and target regions, possibly as a function of the binding of a TF.

`topGOres` and `simplifyGOterms` are convenience functions to deal to GeneOntology enrichment anal-

yses. In particular the latter can be used to keep only the most informative GO terms. This is based on the fact that GeneOntology is composed of three different ontologies (Biological Processes, Molecular Functions and Cellular Components). Within each ontology, a set of GO terms describing those categories are available, together with the relationships linking them. Terms specifying more precisely a biological category are called children (e.g. induction of apoptosis) of more generic parent terms (e.g. apoptosis). Most informative GO terms to keep are defined here as those terms for which an enriched children term mapping to a very similar set of genes has not been also identified. If that happens, the children term is believed to contain most of the information, and typically better specifies the enriched GO category, compared to the more general, less specific, parent term which are thus discarded.

`findLncRNA` is a function to point to putative intergenic long non-coding RNAs (lncRNAs). These are typically identified thanks to their epigenetic signatures, characteristic of transcriptional units unrelated from gene-coding ones, associated to known genes. For simplicity, this function can only point to lncRNAs which are distal from gene transcriptional units, to avoid False Positives due to overlap with coding transcriptional units. While the detection is mostly based on H3K4me3 and H3K4me1 peaks, additional marks and RNA-seq data can be used to have a more robust lncRNAs identification. See the documentation of the `findLncRNA` function for details.

`getPromoterClass` can be used to classify promoters according to their CpG content. In fact, it has been show that promoters with significantly different CpG content can be differently responsive to the presence/absence or even to the level of epigenetic marks such as DNA-methylation (Koga et al, Genome Research 2009).

5 Visualization

`heatmapData` and the associated `heatmapPlot` are very flexible functions that can be used with a range of data types to associate counts for multiple data or annotation tracks to genomic regions of interest (ROIs), and subsequently visualizing similarities in those patterns. Supported data types range from BAM files, *GRanges* objects, *GRanges* metadata, putative methylation sites and their associated absolute and relative methylation level. All these data types are highly relevant for epigenomics integrative analyses and can include but they are not limited to: base-resolution or low-resolution DNA methylation data, histone marks, transcription factor binding, RNA-seq expression, which can be freely combined (data tracks). Importantly, predefined or user-defined genomic annotation(s) (annotation tracks) could be overlaid to further stratify the patterns emerging from the data tracks.

The counts associated to each given data track can be visualized using heatmaps, clustering regions with similar patterns over all (or a subset) of the provided data and annotation tracks (easily dozens of them, covering hundreds or thousands of genomic regions), together with the underlying gene annotation on the forward and reverse strand. ROIs can be potentially divided in a number of bins to provide higher resolution and information about patterns of the data within those regions.

BAM files can be provided to count reads in ROIs; these files could contain any kind of high-throughput sequencing data (including histone marks or transcription factor binding sites) in the case one wants to focus to the density of the reads on the genome, independently from the identification of significantly enriched/scoring regions. On the other hand *GRanges* can be provided to pass for example ChIP-seq peaks, or MeDIP-seq methylated regions, focusing on presence/absence of a given mark in the ROIs. The `mcols` of *GRanges* having the same length of the ROIs could be also pre-populated by the user and used directly to fill in the relevant information of a given data track. DNA methylation relevant data could be provided using *GElist* or *GEcollection* object of the *methyIPipe* package. All these data types can be freely combined in the input list of the `heatmapPlot` function.

Importantly, the colorscale in the heatmap can optionally be set to display the significance of the data associated with any track in each genomic region. In the following example two very simple heatmaps are drawn, reporting putative transcription factor binding sites (actually TSS proximal regions) in the context of gene annotation (introns and exons are reported with dark and light red in the forward and

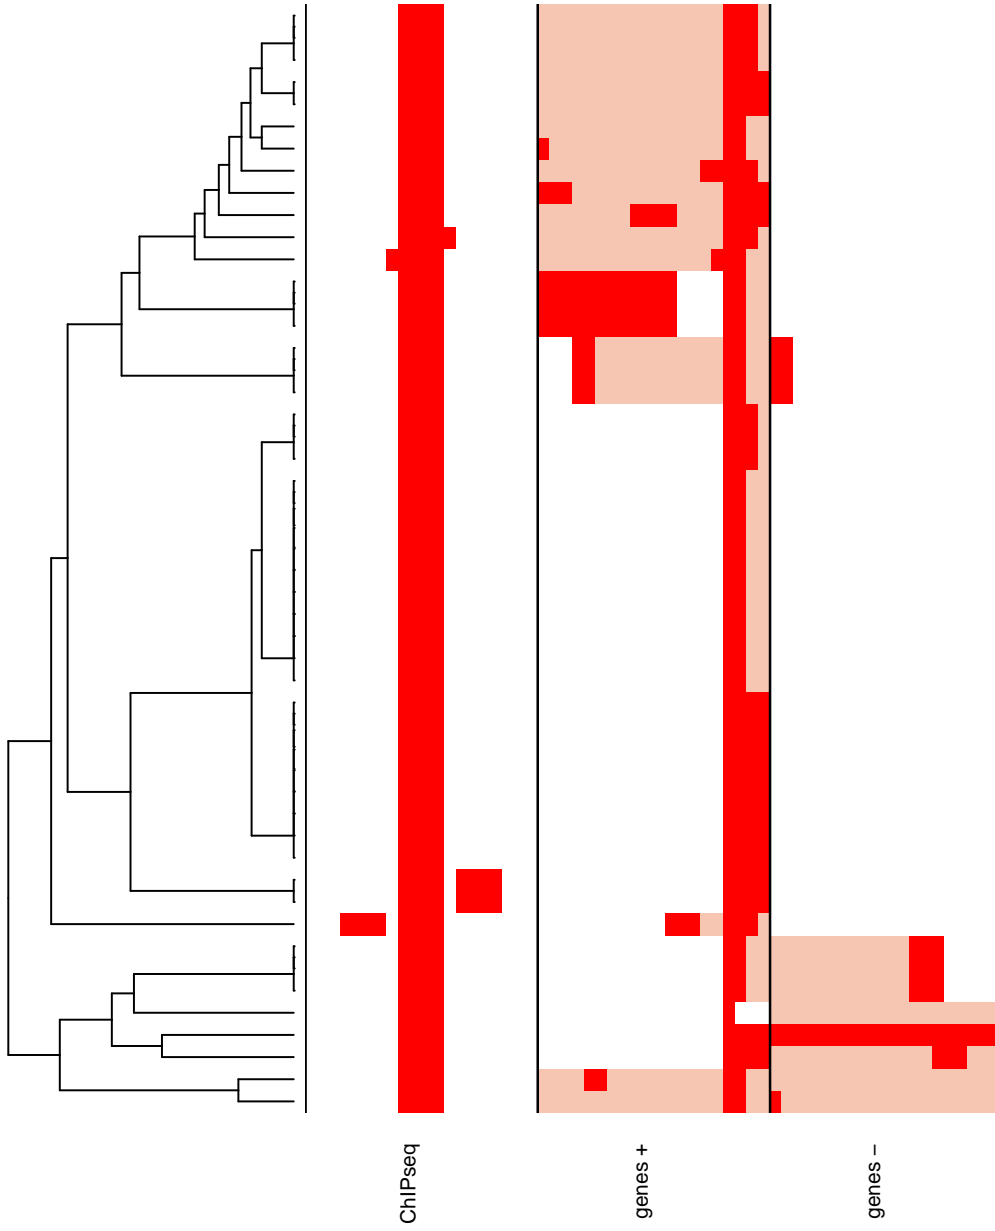
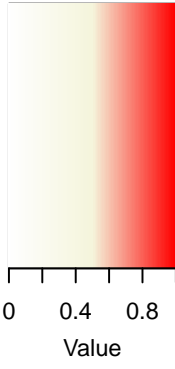
reverse strand, respectively). In the following heatmap the same is displayed according to the p-value for each peak, where regions associated to high p-values are dimmed to a less intense colour.

```
gr <- TSSpos[1:50]
start(gr) <- start(gr) - 1000
end(gr) <- end(gr) - 600
extgr <- GRanges(seqnames(gr), ranges=IRanges(start(gr) - 1000, end(gr) + 1000))
data <- heatmapData(grl=list(ChIPseq=gr), refgr=extgr, type='gr', nbins=20, txdb=txdb)

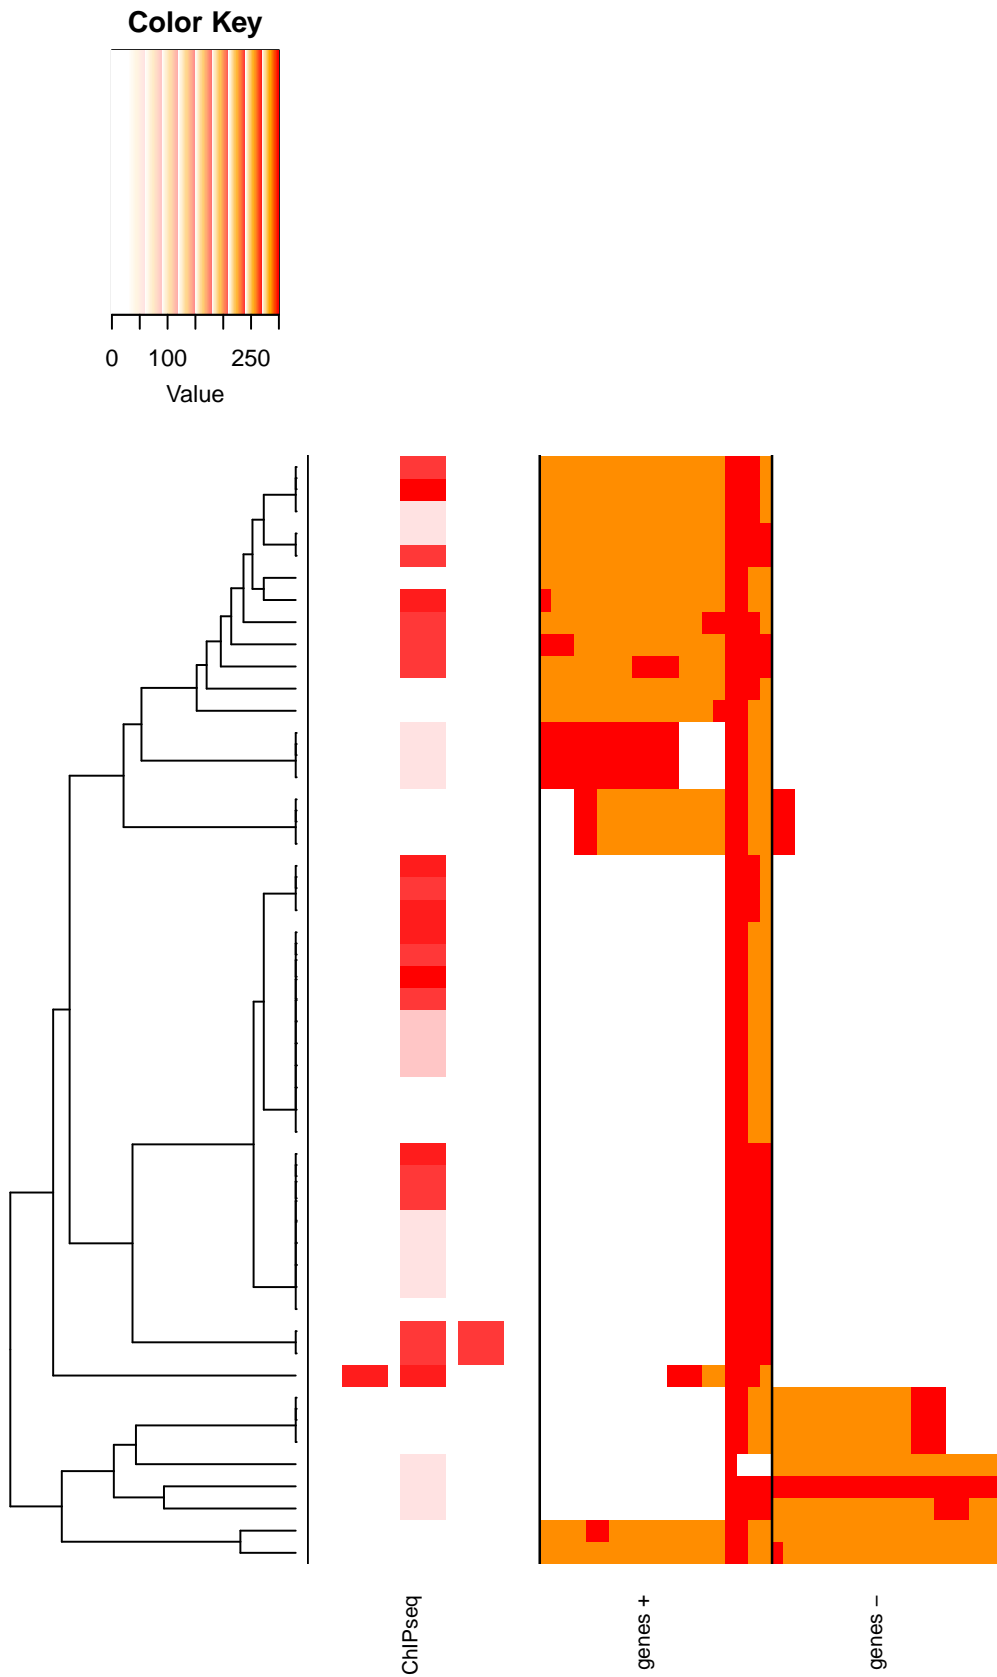
## [1] "ChIPseq"

pvalues <- c(runif(20,1e-20,1e-8), runif(15,1e-4,1e-2), runif(15,0.5,1))
pvalues <- cbind(pvalues, rep(0, 50), rep(0, 50))
rownames(data[[1]][[1]]) <- paste(1:50, signif(pvalues[,1],1), sep=' # ')
heatmapPlot(matList=data[[1]], clusterInds=1:3)
```

Color Key



```
heatmapPlot(matList=data[[1]], sigMat=pvalues, clusterInds=1:3)
```



6 Session Information

```
sessionInfo()

## R version 3.3.1 (2016-06-21)
## Platform: x86_64-pc-linux-gnu (64-bit)
## Running under: Ubuntu 16.04.1 LTS
##
## 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] stats4      parallel  stats      graphics  grDevices  utils
## [7] datasets   methods   base
##
## other attached packages:
## [1] org.Mm.eg.db_3.4.0
## [2] TxDb.Mmusculus.UCSC.mm9.knownGene_3.2.2
## [3] GenomicFeatures_1.26.0
## [4] compEpiTools_1.8.0
## [5] GenomicRanges_1.26.0
## [6] GenomeInfoDb_1.10.0
## [7] topGO_2.26.0
## [8] SparseM_1.72
## [9] GO.db_3.4.0
## [10] AnnotationDbi_1.36.0
## [11] IRanges_2.8.0
## [12] S4Vectors_0.12.0
## [13] Biobase_2.34.0
## [14] graph_1.52.0
## [15] BiocGenerics_0.20.0
##
## loaded via a namespace (and not attached):
## [1] httr_1.2.1              AnnotationHub_2.6.0
## [3] splines_3.3.1           gtools_3.5.0
## [5] Formula_1.2-1           shiny_0.14.1
## [7] interactiveDisplayBase_1.12.0 highr_0.6
## [9] latticeExtra_0.6-28     BSgenome_1.42.0
## [11] Rsamtools_1.26.0        RSQLite_1.0.0
## [13] lattice_0.20-34         biovizBase_1.22.0
## [15] limma_3.30.0            chron_2.3-47
## [17] digest_0.6.10          RColorBrewer_1.1-2
## [19] XVector_0.14.0          colorspace_1.2-7
## [21] htmltools_0.3.5         httpuv_1.3.3
```

```

## [23] Matrix_1.2-7.1          plyr_1.8.4
## [25] XML_3.98-1.4            biomaRt_2.30.0
## [27] zlibbioc_1.20.0        xtable_1.8-2
## [29] scales_0.4.0           gdata_2.17.0
## [31] BiocParallel_1.8.0     ggplot2_2.1.0
## [33] SummarizedExperiment_1.4.0 methylPipe_1.8.0
## [35] nnet_7.3-12            Gviz_1.18.0
## [37] mime_0.5                survival_2.39-5
## [39] magrittr_1.5           evaluate_0.10
## [41] gplots_3.0.1           foreign_0.8-67
## [43] BiocInstaller_1.24.0   tools_3.3.1
## [45] data.table_1.9.6       formatR_1.4
## [47] matrixStats_0.51.0    stringr_1.1.0
## [49] munsell_0.4.3          cluster_2.0.5
## [51] ensemblDb_1.6.0        Biostrings_2.42.0
## [53] caTools_1.17.1         grid_3.3.1
## [55] RCurl_1.95-4.8         dichromat_2.0-0
## [57] marray_1.52.0          VariantAnnotation_1.20.0
## [59] bitops_1.0-6           codetools_0.2-15
## [61] gtable_0.2.0           DBI_0.5-1
## [63] R6_2.2.0               GenomicAlignments_1.10.0
## [65] gridExtra_2.2.1        knitr_1.14
## [67] rtracklayer_1.34.0     Hmisc_3.17-4
## [69] KernSmooth_2.23-15     stringi_1.1.2
## [71] Rcpp_0.12.7            rpart_4.1-10
## [73] acepack_1.3-3.3

```