The chromstaR user's guide

Aaron Taudt *

*aaron.taudt@gmail.com

November 7, 2021

Contents

1 Introduction

ChIP-seq has become the standard technique for assessing the genome-wide chromatin state of DNA. *chromstaR* provides functions for the joint analysis of multiple ChIP-seq samples. It allows peak calling for transcription factor binding and histone modifications with a narrow (e.g. H3K4me3, H3K27ac, ...) or broad (e.g. H3K36me3, H3K27me3, ...) profile. All analysis can be performed on each sample individually (=univariate), or in a joint analysis considering all samples simultaneously (=multivariate). The joint analysis is generally recommended because it is more powerful for detecting differences.

2 Citation

If you use *chromstaR* for chromatin state analysis, please cite [1]:

Taudt, A., Nguyen, M. A., Heinig, M., Johannes, F. and Colome-Tatche, M. **chromstaR: Tracking combinatorial chromatin state dynamics in space and time.** bioRxiv (Cold Spring Harbor Labs Journals, 2016). doi:10.1101/038612

If you use *chromstaR* for differential ChIP-seq analysis, please cite [1] and [2]:

Hanna, C. W., Taudt, A., Huang, J., Gahurova, L., Kranz, A., Andrews, S., Dean, W., Stewart, A. F., Colome-Tatche, M. and Kelsey, G. **MLL2 conveys transcription-independent H3K4 trimethylation in oocytes.** Nat. Struct. Mol. Biol. 1 (2018). doi:10.1038/s41594-017-0013-5

3 Outline of workflow

Every analysis with the *chromstaR* package starts from aligned reads in either BAM or BED format. In the first step, the genome is partitioned into non-overlapping, equally sized bins and the reads that fall into each bin are counted. These read counts serve as the basis for both the univariate and the multivariate peak- and broad-region calling. Univariate peak calling is done by fitting a three-state Hidden Markov Model to the binned read counts. Multivariate peak calling for $\mathcal{S}$ samples is done by fitting a $2^{\mathcal{S}}$-state Hidden Markov Model to all binned read counts.

4 Univariate analysis

4.1 Task 1: Peak calling for a narrow histone modification

Examples of histone modifications with a narrow profile are H3K4me3, H3K9ac and H3K27ac in most human tissues. For such peak-like modifications, the bin size should be set to a value between 200bp and 1000bp.

```
library(chromstaR)

## === Step 1: Binning ===
# Get an example BAM file
file <- system.file("extdata","euratrans","lv-H3K4me3-BN-male-bio2-tech1.bam",
                    package="chromstaRData")
# Get the chromosome lengths (see ?GenomeInfoDb::fetchExtendedChromInfoFromUCSC)
# This is only necessary for BED files. BAM files are handled automatically.
data(rn4_chrominfo)
head(rn4_chrominfo)
```

```
##   chromosome   length
## 1       chrM    16300
## 2      chr12 46782294
## 3      chr20 55268282
## 4      chr19 59218465
## 5      chr18 87265094
## 6      chr11 87759784

# We use bin size 1000bp and chromosome 12 to keep the example quick
binned.data <- binReads(file, assembly=rn4_chrominfo, binsizes=1000,
                        stepsizes=500, chromosomes='chr12')
print(binned.data)

## GRanges object with 46782 ranges and 1 metadata column:
##          seqnames               ranges strand |   counts
##             <Rle>            <IRanges>  <Rle> | <matrix>
##      [1]    chr12              1-1000      * |      0:0
##      [2]    chr12           1001-2000      * |      0:0
##      [3]    chr12           2001-3000      * |      0:0
##      [4]    chr12           3001-4000      * |      0:0
##      [5]    chr12           4001-5000      * |      0:0
##      ...      ...                  ...    ... .      ...
##  [46778]    chr12 46777001-46778000      * |      2:3
##  [46779]    chr12 46778001-46779000      * |      1:0
##  [46780]    chr12 46779001-46780000      * |      0:0
##  [46781]    chr12 46780001-46781000      * |      2:3
##  [46782]    chr12 46781001-46782000      * |      1:0
##  -------
##  seqinfo: 1 sequence from an unspecified genome
```
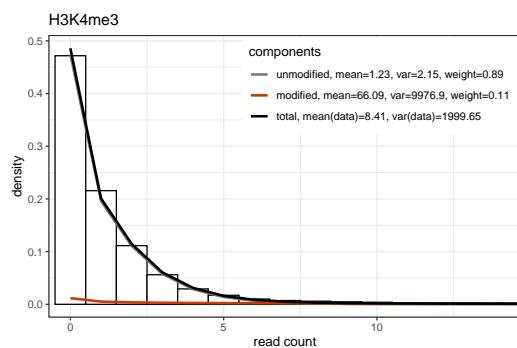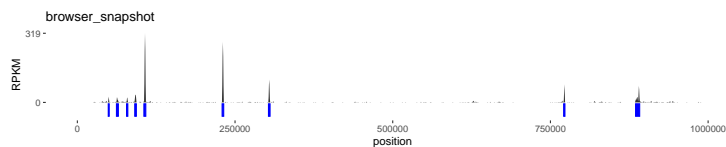
```
## === Step 2: Peak calling ===
model <- callPeaksUnivariate(binned.data, verbosity=0)
```

```
## === Step 3: Checking the fit ===
# For a narrow modification, the fit should look something like this,
# with the 'modified'-component near the bottom of the figure
plotHistogram(model) + ggtitle('H3K4me3')
```



```
# We can also check a browser snapshot of the data
plotGenomeBrowser(model, chr='chr12', start=1, end=1e6)[[1]]
```



```
## === Step 4: Working with peaks ===
# Get the number and average size of peaks
length(model$peaks); mean(width(model$peaks))

## [1] 1245
## [1] 4008.434

# Adjust the sensitivity and get number of peaks
model <- changeMaxPostCutoff(model, maxPost.cutoff=0.9999)
length(model$peaks); mean(width(model$peaks))

## [1] 913
## [1] 4861.993
```

```
## === Step 5: Export to genome browser ===
# We can export peak calls and binned read counts with
exportPeaks(model, filename=tempfile())
exportCounts(model, filename=tempfile())
```

!! It is important that the distributions are fitted correctly !! Please check section 7.4 for examples of how this plot should *not* look like and what can be done to get a correct fit.

## 4.2 Task 2: Peak calling for a broad histone modification

Examples of histone modifications with a broad profile are H3K9me3, H3K27me3, H3K36me3, H4K20me1 in most human tissues. These modifications usually cover broad domains of the genome, and the enrichment is best captured with a bin size between 500bp and 2000bp.

```
library(chromstaR)
```

```
## === Step 1: Binning ===
# Get an example BAM file
file <- system.file("extdata","euratrans","lv-H3K27me3-BN-male-bio2-tech1.bam",
                     package="chromstaRData")
# Get the chromosome lengths (see ?GenomeInfoDb::fetchExtendedChromInfoFromUCSC)
# This is only necessary for BED files. BAM files are handled automatically.
data(rn4_chrominfo)
head(rn4_chrominfo)
# We use bin size 1000bp and chromosome 12 to keep the example quick
binned.data <- binReads(file, assembly=rn4_chrominfo, binsizes=1000,
                        stepsizes=500, chromosomes='chr12')
```

```
## === Step 2: Peak calling ===
model <- callPeaksUnivariate(binned.data, verbosity=0)
```

```
## === Step 3: Checking the fit ===
# For a broad modification, the fit should look something like this,
# with a 'modified'-component that fits the thick tail of the distribution.
plotHistogram(model) + ggtitle('H3K27me3')
```



```
plotGenomeBrowser(model, chr='chr12', start=1, end=1e6)[[1]]
```
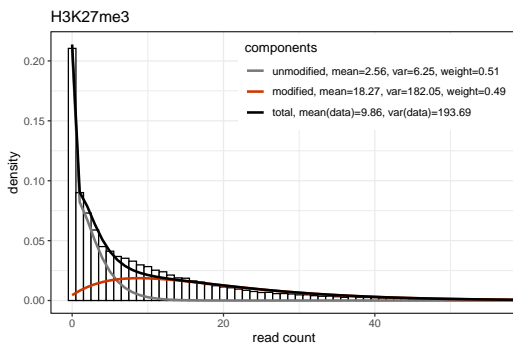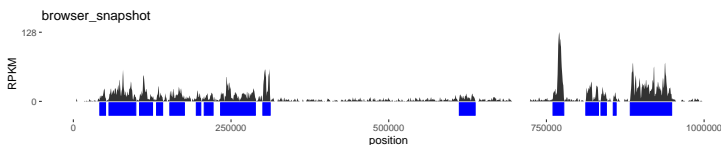


```
## === Step 4: Working with peaks ===
peaks <- model$peaks
length(peaks); mean(width(peaks))
```

```
## [1] 523
## [1] 43522.94
```

```
# Adjust the sensitivity and get number of peaks
model <- changeMaxPostCutoff(model, maxPost.cutoff=0.9999)
peaks <- model$peaks
length(peaks); mean(width(peaks))
```

```
## [1] 416
## [1] 52582.93
```

```
## === Step 5: Export to genome browser ===
# We can export peak calls and binned read counts with
exportPeaks(model, filename=tempfile())
exportCounts(model, filename=tempfile())
```

```
## === Step 1-3: Another example for mark H4K20me1 ===
file <- system.file("extdata","euratrans","lv-H4K20me1-BN-male-bio1-tech1.bam",
                    package="chromstaRData")
data(rn4_chrominfo)
binned.data <- binReads(file, assembly=rn4_chrominfo, binsizes=1000,
                    stepsizes=500, chromosomes='chr12')
model <- callPeaksUnivariate(binned.data, max.time=60, verbosity=0)
plotHistogram(model) + ggtitle('H4K20me1')
```
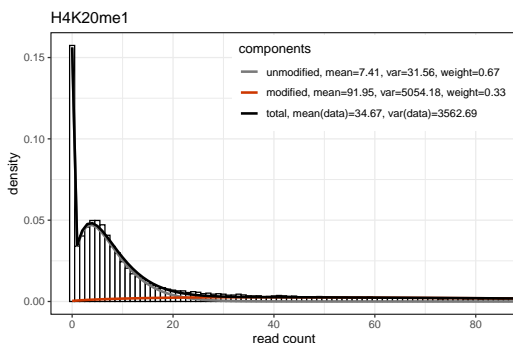


```
# We can also check a browser snapshot of the data
plotGenomeBrowser(model, chr='chr12', start=1, end=1e6)[[1]]
```


!! It is important that the distributions are fitted correctly !! Please check section 7.4 for examples of how this plot should *not* look like and what can be done to get a correct fit.

4.3 Task 3: Peak calling for ATAC-seq, DNase-seq, FAIRE-seq, ...

Peak calling for ATAC-seq and DNase-seq is similar to the peak calling of a narrow histone modification (section 4.1). FAIRE-seq experiments seem to exhibit a broad profile with our model, so the procedure is similar to the domain calling of a broad histone modification (section 4.2).

5 Multivariate analysis

5.1 Task 1: Integrating multiple replicates

chromstaR can be used to call peaks with multiple replicates, without the need of prior merging. The underlying statistical model integrates information from all replicates to identify common peaks. It is, however, important to note that replicates with poor quality can affect

the joint peak calling negatively. It is therefore recommended to first check the replicate quality and discard poor-quality replicates. The necessary steps for peak calling for an example ChIP-seq experiment with 4 replicates are detailed below.

Please note that also the other tasks in this section (Task 5.2, 5.3 and 5.4) can handle multiple replicates via specification of the `experiment.table` parameter. The following example demonstrates how to explicitly use multiple replicates for peak calling and their correlation as a basic quality control.

```r
library(chromstaR)
```

```r
#=== Step 1: Preparation ===
# Let's get some example data with 3 replicates in spontaneous hypertensive rat (SHR)
file.path <- system.file("extdata","euratrans", package='chromstaRData')
files.good <- list.files(file.path, pattern="H3K27me3.*SHR.*bam$", full.names=TRUE)[1:3]
# We fake a replicate with poor quality by taking a different mark entirely
files.poor <- list.files(file.path, pattern="H4K20me1.*SHR.*bam$", full.names=TRUE)[1]
files <- c(files.good, files.poor)
# Obtain chromosome lengths. This is only necessary for BED files. BAM files are
# handled automatically.
data(rn4_chrominfo)
head(rn4_chrominfo)

##    chromosome   length
## 1        chrM    16300
## 2       chr12 46782294
## 3       chr20 55268282
## 4       chr19 59218465
## 5       chr18 87265094
## 6       chr11 87759784

# Define experiment structure
exp <- data.frame(file=files, mark='H3K27me3', condition='SHR', replicate=1:4,
                  pairedEndReads=FALSE, controlFiles=NA)

# Peaks could now be called with
# Chromstar(inputfolder=file.path, experiment.table=exp, outputfolder=tempdir(),
#           mode = 'separate')
# However, to get more information on the replicates we will choose
# a more detailed workflow.
```

```r
## === Step 2: Binning ===
# We use bin size 1000bp and chromosome 12 to keep the example quick
binned.data <- list()
for (file in files) {
  binned.data[[basename(file)]] <- binReads(file, binsize=1000, stepsizes=500,
                                    assembly=rn4_chrominfo, chromosomes='chr12',
                                    experiment.table=exp)
}
```

```r
## === Step 3: Univariate peak calling ===
# The univariate fit is obtained for each replicate
models <- list()
for (i1 in 1:length(binned.data)) {
  models[[i1]] <- callPeaksUnivariate(binned.data[[i1]], max.time=60)
  plotHistogram(models[[i1]])
}
```

!! It is important that the distributions are fitted correctly !! Please check section 7.4 for examples of how this plot should *not* look like and what can be done to get a correct fit.

```r
## === Step 4: Check replicate correlation ===
# We run a multivariate peak calling on all 4 replicates
# A warning is issued because replicate 4 is very different from the others
multi.model <- callPeaksReplicates(models, max.time=60, eps=1)

## HMM: number of states = 16
## HMM: number of bins = 46782
## HMM: maximum number of iterations = none
## HMM: maximum running time = 60 sec
## HMM: epsilon = 1
## HMM: number of experiments = 4
## Iteration              log(P)          dlog(P)    Time in sec
##         0               -inf                -              0
## HMM: Precomputing densities ...
```

```
## Iteration            log(P)         dlog(P)    Time in sec
##        0              -inf               -              0
##        1     -542989.146422             inf             0
##        2     -538374.740061      4614.406361            0
##        3     -538271.399633       103.340428            0
##        4     -538250.213731        21.185902            0
##        5     -538244.134929         6.078802            0
##        6     -538241.935145         2.199784            0
##        7     -538240.963921         0.971224            1
## HMM: Convergence reached!
## HMM: Recoding posteriors ...

## Warning in callPeaksReplicates(models, max.time = 60, eps = 1):  Your replicates cluster in 2 groups.  Consider
redoing your analysis with only the group with the highest average coverage:
## H3K27me3-SHR-rep1
## H3K27me3-SHR-rep2
## H3K27me3-SHR-rep3
## Replicates from groups with lower coverage are:
## H3K27me3-SHR-rep4

# Checking the correlation confirms that Rep4 is very different from the others
print(multi.model$replicateInfo$correlation)

##                   H3K27me3-SHR-rep1 H3K27me3-SHR-rep2 H3K27me3-SHR-rep3 H3K27me3-SHR-rep4
## H3K27me3-SHR-rep1         1.0000000         0.9999358         0.9997432        -0.3718157
## H3K27me3-SHR-rep2         0.9999358         1.0000000         0.9997217        -0.3717750
## H3K27me3-SHR-rep3         0.9997432         0.9997217         1.0000000        -0.3716530
## H3K27me3-SHR-rep4        -0.3718157        -0.3717750        -0.3716530         1.0000000
```

```
## === Step 5: Peak calling with replicates ===
# We redo the previous step without the "bad" replicate
# Also, we force all replicates to agree in their peak calls
multi.model <- callPeaksReplicates(models[1:3], force.equal=TRUE, max.time=60)
```

```
## === Step 6: Export to genome browser ===
# Finally, we can export the results as BED file
exportPeaks(multi.model, filename=tempfile())
exportCounts(multi.model, filename=tempfile())
```

5.2 Task 2: Detecting differentially modified regions

chromstaR is extremely powerful in detecting differentially modified regions in two or more samples. The following example illustrates this on ChIP-seq data for H4K20me1 in brown norway (BN) and spontaneous hypertensive rat (SHR) in left-ventricle (lv) heart tissue. The mode of analysis is called *differential*.

```
library(chromstaR)
```

```
#=== Step 1: Preparation ===
## Prepare the file paths. Exchange this with your input and output directories.
inputfolder <- system.file("extdata","euratrans", package="chromstaRData")
outputfolder <- file.path(tempdir(), 'H4K20me1-example')

## Define experiment structure, put NA if you don't have controls
data(experiment_table_H4K20me1)
print(experiment_table_H4K20me1)

##                                file    mark condition replicate pairedEndReads
## 1  lv-H4K20me1-BN-male-bio1-tech1.bam H4K20me1        BN         1          FALSE
## 2  lv-H4K20me1-BN-male-bio2-tech1.bam H4K20me1        BN         2          FALSE
## 3 lv-H4K20me1-SHR-male-bio1-tech1.bam H4K20me1       SHR         1          FALSE
##                                                        controlFiles
## 1 lv-input-BN-male-bio1-tech1.bam|lv-input-BN-male-bio1-tech2.bam
## 2 lv-input-BN-male-bio1-tech1.bam|lv-input-BN-male-bio1-tech2.bam
## 3                         lv-input-SHR-male-bio1-tech1.bam

## Define assembly
# This is only necessary if you have BED files, BAM files are handled automatically.
# For common assemblies you can also specify them as 'hg19' for example.
data(rn4_chrominfo)
head(rn4_chrominfo)

##   chromosome   length
## 1       chrM    16300
```

```
## 2        chr12 46782294
## 3        chr20 55268282
## 4        chr19 59218465
## 5        chr18 87265094
## 6        chr11 87759784
```

```
#=== Step 2: Run Chromstar ===
## Run ChromstaR
Chromstar(inputfolder, experiment.table=experiment_table_H4K20me1,
          outputfolder=outputfolder, numCPU=4, binsize=1000, stepsize=500,
          assembly=rn4_chrominfo, prefit.on.chr='chr12', chromosomes='chr12',
          mode='differential')
```

```
## Results are stored in 'outputfolder' and can be loaded for further processing
list.files(outputfolder)

## [1] "BROWSERFILES"      "PLOTS"             "README.txt"        "binned"
## [5] "chrominfo.tsv"     "chromstaR.config"  "combined"          "experiment_table.tsv"
## [9] "multivariate"      "univariate"

model <- get(load(file.path(outputfolder,'multivariate',
      'multivariate_mode-differential_mark-H4K20me1_binsize1000_stepsize500.RData')))
```

!! It is important that the distributions in folder outputfolder/PLOTS/univariate-distributions are fitted correctly !!
Please check section 7.4 for examples of how this plot should *not* look like and what can be done to get a correct fit.

```
## === Step 3: Construct differential and common states ===
diff.states <- stateBrewer(experiment_table_H4K20me1, mode='differential',
                           differential.states=TRUE)
print(diff.states)

##    combination state
## 1        [SHR]     1
## 2         [BN]     6

common.states <- stateBrewer(experiment_table_H4K20me1, mode='differential',
                             common.states=TRUE)
print(common.states)

##    combination state
## 1          []     0
## 2     [BN+SHR]    7
```

```
## === Step 5: Export to genome browser ===
# Export only differential states
exportPeaks(model, filename=tempfile())
exportCounts(model, filename=tempfile())
exportCombinations(model, filename=tempfile(), include.states=diff.states)
```

5.3   Task 3: Finding combinatorial chromatin states

Most experimental studies that probe several histone modifications are interested in combinatorial chromatin states. An example of a simple combinatorial state would be [H3K4me3+H3K27me3], which is also frequently called "bivalent promoter", due to the simultaneous occurrence of the promoter marking H3K4me3 and the repressive H3K27me3. Finding combinatorial states with *chromstaR* is equivalent to a multivariate peak calling. The following code chunks demonstrate how to find bivalent promoters and do some simple analysis. The mode of analysis is called *combinatorial*.

```
library(chromstaR)
```

```
#=== Step 1: Preparation ===
## Prepare the file paths. Exchange this with your input and output directories.
inputfolder <- system.file("extdata","euratrans", package="chromstaRData")
outputfolder <- file.path(tempdir(), 'SHR-example')

## Define experiment structure, put NA if you don't have controls
# (SHR = spontaneous hypertensive rat)
data(experiment_table_SHR)
print(experiment_table_SHR)

##                                        file     mark condition replicate pairedEndReads
## 1 lv-H3K27me3-SHR-male-bio2-tech1.bam H3K27me3       SHR         1          FALSE
```

```
## 2  lv-H3K27me3-SHR-male-bio2-tech2.bam  H3K27me3     SHR       2        FALSE
## 3   lv-H3K4me3-SHR-male-bio2-tech1.bam   H3K4me3     SHR       1        FALSE
## 4   lv-H3K4me3-SHR-male-bio3-tech1.bam   H3K4me3     SHR       2        FALSE
##                    controlFiles
## 1 lv-input-SHR-male-bio1-tech1.bam
## 2 lv-input-SHR-male-bio1-tech1.bam
## 3 lv-input-SHR-male-bio1-tech1.bam
## 4 lv-input-SHR-male-bio1-tech1.bam

## Define assembly
# This is only necessary if you have BED files, BAM files are handled automatically.
# For common assemblies you can also specify them as 'hg19' for example.
data(rn4_chrominfo)
head(rn4_chrominfo)

##   chromosome   length
## 1       chrM    16300
## 2      chr12 46782294
## 3      chr20 55268282
## 4      chr19 59218465
## 5      chr18 87265094
## 6      chr11 87759784
```

```
#=== Step 2: Run Chromstar ===
## Run ChromstaR
Chromstar(inputfolder, experiment.table=experiment_table_SHR,
         outputfolder=outputfolder, numCPU=4, binsize=1000, stepsize=500,
         assembly=rn4_chrominfo, prefit.on.chr='chr12', chromosomes='chr12',
         mode='combinatorial')
```

!! It is important that the distributions in folder outputfolder/PLOTS/univariate-distributions are fitted correctly !!
Please check section 7.4 for examples of how this plot should *not* look like and what can be done to get a correct fit.

```
## Results are stored in 'outputfolder' and can be loaded for further processing
list.files(outputfolder)

##  [1] "BROWSERFILES"        "PLOTS"             "README.txt"          "binned"
##  [5] "chrominfo.tsv"       "chromstaR.config"  "combined"            "experiment_table.tsv"
##  [9] "multivariate"        "univariate"

model <- get(load(file.path(outputfolder,'multivariate',
      'multivariate_mode-combinatorial_condition-SHR_binsize1000_stepsize500.RData')))
# Obtain genomic frequencies for combinatorial states
genomicFrequencies(model)

## $frequency
##
##             []        [H3K4me3]         [H3K27me3] [H3K27me3+H3K4me3]
##       0.41193194       0.09392501         0.43134111         0.06280193
##
## $domains
##
##             []        [H3K4me3]         [H3K27me3] [H3K27me3+H3K4me3]
##           1188              683               1166                893

# Plot transition probabilities and read count correlation
heatmapTransitionProbs(model) + ggtitle('Transition probabilities')
```
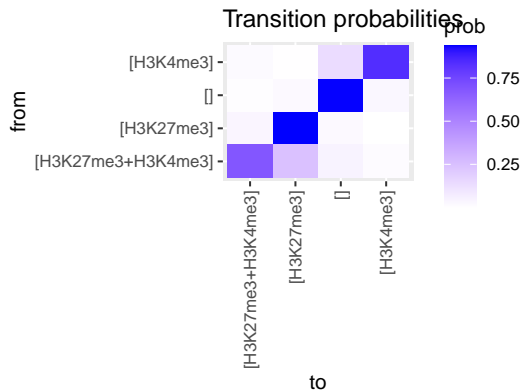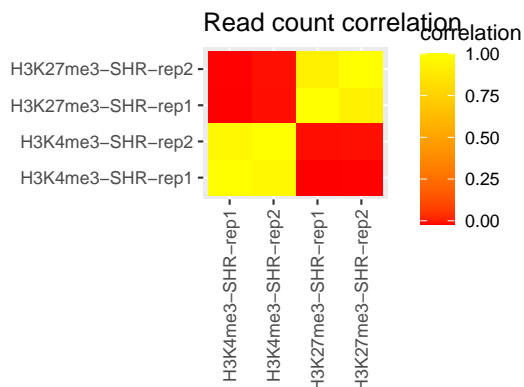
```
heatmapCountCorrelation(model) + ggtitle('Read count correlation')
```



```
## === Step 3: Enrichment analysis ===
# Annotations can easily be obtained with the biomaRt package. Of course, you can
# also load them from file if you already have annotation files available.
library(biomaRt)
ensembl <- useEnsembl(biomart='ENSEMBL_MART_ENSEMBL', dataset='rnorvegicus_gene_ensembl')
genes <- getBM(attributes=c('ensembl_gene_id', 'chromosome_name', 'start_position',
                            'end_position', 'strand', 'external_gene_name',
                            'gene_biotype'),
               mart=ensembl)
# Transform to GRanges for easier handling
genes <- GRanges(seqnames=paste0('chr',genes$chromosome_name),
                 ranges=IRanges(start=genes$start, end=genes$end),
                 strand=genes$strand,
                 name=genes$external_gene_name, biotype=genes$gene_biotype)
# Rename chrMT to chrM to avoid warnings
seqlevels(genes)[seqlevels(genes)=='chrMT'] <- 'chrM'
# Select only chr12 to avoid warnings
genes <- keepSeqlevels(genes, 'chr12', pruning.mode = 'coarse')
print(genes)

## GRanges object with 919 ranges and 2 metadata columns:
##         seqnames              ranges strand |         name        biotype
##            <Rle>           <IRanges>  <Rle> |  <character>    <character>
##     [1]    chr12 24473981-24476295      + |         Fzd9 protein_coding
##     [2]    chr12   9693779-9727489      - |       Polr1d protein_coding
##     [3]    chr12 23587222-23592774      - |       Polr2j protein_coding
##     [4]    chr12 48292273-48351084      + |         Svop protein_coding
##     [5]    chr12 47074200-47076572      + |       Dynll1 protein_coding
##     ...      ...                 ...    ... .          ...            ...
##   [915]    chr12   6016884-6078411      - |         Rxfp2 protein_coding
##   [916]    chr12   2938358-2940393      - | AABR07034991.1        lincRNA
##   [917]    chr12 25036605-25070539      + |         Limk1 protein_coding
##   [918]    chr12   9728486-9793428      + |         Lnx2 protein_coding
##   [919]    chr12 20667601-20677689      + |    RGD1560281 protein_coding
##   -------
##   seqinfo: 1 sequence from an unspecified genome; no seqlengths
```

```
# We expect promoter [H3K4me3] and bivalent-promoter signatures [H3K4me3+H3K27me3]
# to be enriched at transcription start sites.
plotEnrichment(hmm = model, annotation = genes, bp.around.annotation = 15000) +
  ggtitle('Fold enrichment around genes') +
  xlab('distance from gene body')
```

Fold enrichment around genes


```
# Plot enrichment only at TSS. We make use of the fact that TSS is the start of a gene.
plotEnrichment(model, genes, region = 'start') +
  ggtitle('Fold enrichment around TSS') +
  xlab('distance from TSS in [bp]')
```

**Fold enrichment around TSS**
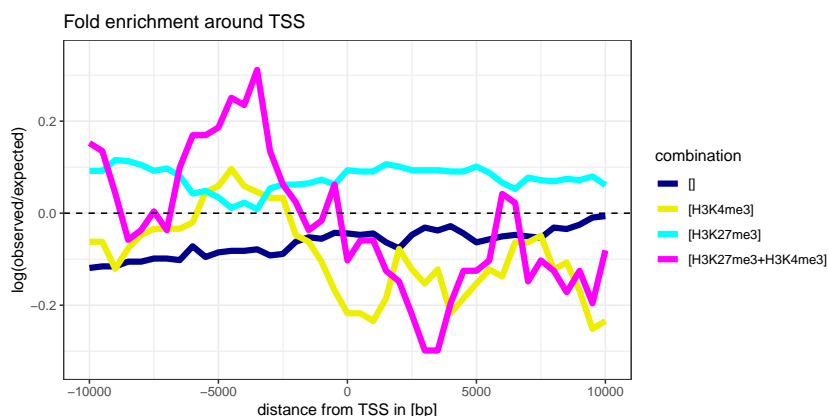

```
# Note: If you want to facet the plot because you have many combinatorial states you
# can do that with
plotEnrichment(model, genes, region = 'start') +
  facet_wrap(~ combination) + ggtitle('Fold enrichment around TSS')
```

Fold enrichment around TSS
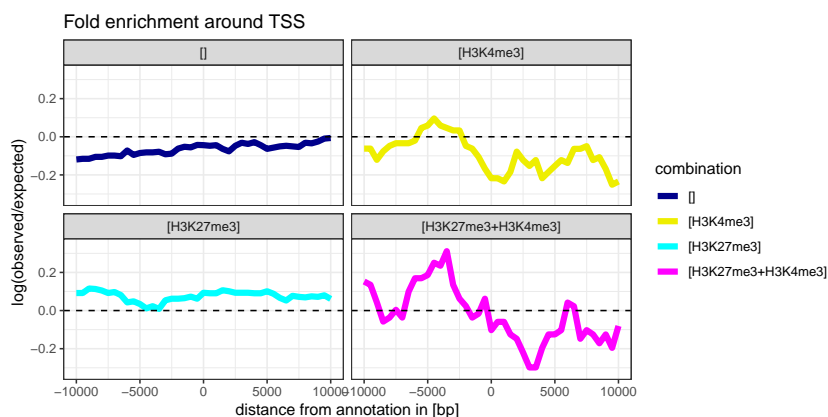

```
# Another form of visualization that shows every TSS in a heatmap
tss <- resize(genes, width = 3, fix = 'start')
plotEnrichCountHeatmap(model, tss, bp.around.annotation = 15000) +
  theme(strip.text.x = element_text(size=6)) +
  scale_x_continuous(breaks=c(-10000,0,10000)) +
  ggtitle('Read count around TSS')
```
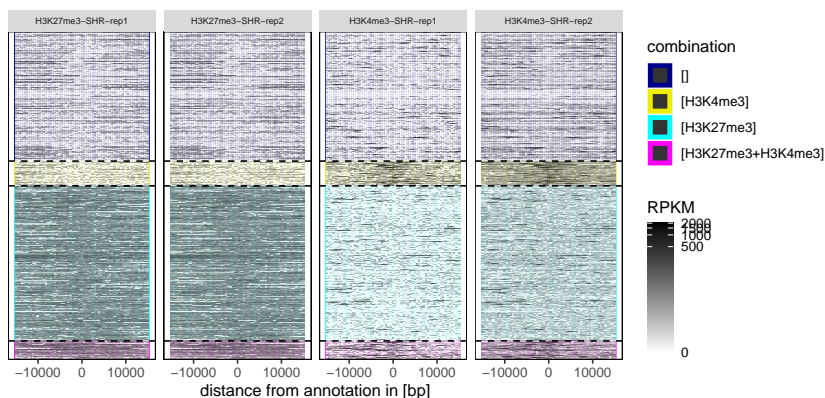
Read count around TSS


```
# Fold enrichment with different biotypes, showing that protein coding genes are
# enriched with (bivalent) promoter combinations [H3K4me3] and [H3K4me3+H3K27me3],
# while rRNA is enriched with the empty [] combination.
biotypes <- split(tss, tss$biotype)
plotFoldEnrichHeatmap(model, annotations=biotypes) + coord_flip() +
  ggtitle('Fold enrichment with different biotypes')

## Warning in regularize.values(x, y, ties, missing(ties), na.rm = na.rm):  collapsing to unique 'x' values

## Warning in regularize.values(x, y, ties, missing(ties), na.rm = na.rm):  collapsing to unique 'x' values

## Warning in regularize.values(x, y, ties, missing(ties), na.rm = na.rm):  collapsing to unique 'x' values
```

Fold enrichment with different biotypes
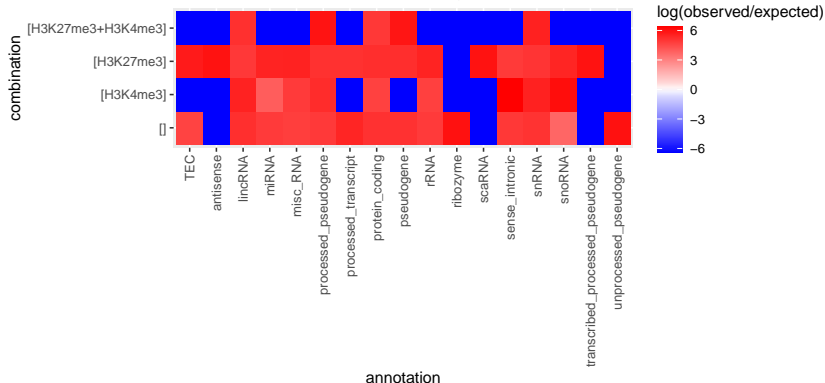

```
# === Step 4: Expression analysis ===
# Suppose you have expression data as well for your experiment. The following code
# shows you how to get the expression values for each combinatorial state.
data(expression_lv)
head(expression_lv)

##       ensembl_gene_id expression_BN expression_SHR
## 1 ENSRNOG00000000001           8.8            7.4
## 2 ENSRNOG00000000007          20.0           13.0
## 3 ENSRNOG00000000008           1.8            3.4
## 4 ENSRNOG00000000010           6.2          506.8
## 5 ENSRNOG00000000012          48.0           36.4
## 6 ENSRNOG00000000014          18.2           15.2

# We need to get coordinates for each of the genes
library(biomaRt)
ensembl <- useEnsembl(biomart='ENSEMBL_MART_ENSEMBL', dataset='rnorvegicus_gene_ensembl')
genes <- getBM(attributes=c('ensembl_gene_id', 'chromosome_name', 'start_position',
                            'end_position', 'strand', 'external_gene_name',
                            'gene_biotype'),
              mart=ensembl)
expr <- merge(genes, expression_lv, by='ensembl_gene_id')
# Transform to GRanges
expression.SHR <- GRanges(seqnames=paste0('chr',expr$chromosome_name),
                          ranges=IRanges(start=expr$start, end=expr$end),
                          strand=expr$strand, name=expr$external_gene_name,
```

```
                        biotype=expr$gene_biotype,
                        expression=expr$expression_SHR)
# Rename chrMT to chrM to avoid warnings
seqlevels(expression.SHR)[seqlevels(expression.SHR)=='chrMT'] <- 'chrM'
# We apply an asinh transformation to reduce the effect of outliers
expression.SHR$expression <- asinh(expression.SHR$expression)

## Plot
plotExpression(model, expression.SHR) +
  theme(axis.text.x=element_text(angle=0, hjust=0.5)) +
  ggtitle('Expression of genes overlapping combinatorial states')
```
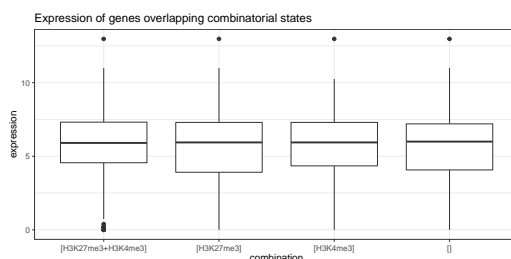


```
plotExpression(model, expression.SHR, return.marks=TRUE) +
  ggtitle('Expression of marks overlapping combinatorial states')
```
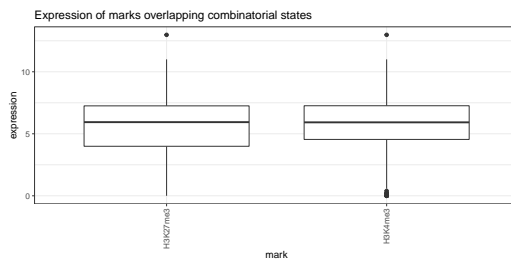

5.4 Task 4: Finding differences between combinatorial chromatin states

Consider bivalent promoters defined by [H3K4me3+H3K27me3] at two different developmental stages, or in two different strains or tissues. This is an example where one is interested in *differences* between *combinatorial states*. The following example demonstrates how such an analysis can be done with *chromstaR*. We use a data set from the Euratrans project (downsampled to chr12) to find differences in bivalent promoters between brown norway (BN) and spontaneous hypertensive rat (SHR) in left-ventricle (lv) heart tissue.

Chromstar can be run in 4 different modes:

- *full*: Recommended mode if your (number of marks) * (number of conditions) is less or equal to 8. With 8 ChIP-seq experiments there are already $2^8 = 256$ combinatorial states which is the maximum that most computers can handle computationally for a human-sized genome at bin size 1000bp.

- **DEFAULT** *differential*: Choose this mode if you are interested in highly significant differences between conditions. The computational limit for the number of conditions is ~ 8 for a human-sized genome. Combinatorial states are not as accurate as in mode *combinatorial* or *full*.

- *combinatorial*: This mode will yield good combinatorial chromatin state calls for any number of marks and conditions. However, differences between conditions have more false positives than in mode *differential* or *full*.

- *separate*: Only replicates are processed in a multivariate manner. Combinatorial states are constructed by a simple post-hoc combination of peak calls.

```
library(chromstaR)

#=== Step 1: Preparation ===
## Prepare the file paths. Exchange this with your input and output directories.
inputfolder <- system.file("extdata","euratrans", package="chromstaRData")
outputfolder <- file.path(tempdir(), 'SHR-BN-example')

## Define experiment structure, put NA if you don't have controls
data(experiment_table)
print(experiment_table)

##                                  file     mark condition replicate pairedEndReads
## 1  lv-H3K27me3-BN-male-bio2-tech1.bam H3K27me3        BN         1          FALSE
## 2  lv-H3K27me3-BN-male-bio2-tech2.bam H3K27me3        BN         2          FALSE
## 3 lv-H3K27me3-SHR-male-bio2-tech1.bam H3K27me3       SHR         1          FALSE
## 4 lv-H3K27me3-SHR-male-bio2-tech2.bam H3K27me3       SHR         2          FALSE
## 5 lv-H3K4me3-BN-female-bio1-tech1.bam  H3K4me3        BN         1          FALSE
## 6   lv-H3K4me3-BN-male-bio2-tech1.bam  H3K4me3        BN         2          FALSE
## 7  lv-H3K4me3-SHR-male-bio2-tech1.bam  H3K4me3       SHR         1          FALSE
## 8  lv-H3K4me3-SHR-male-bio3-tech1.bam  H3K4me3       SHR         2          FALSE
##                                                          controlFiles
## 1 lv-input-BN-male-bio1-tech1.bam|lv-input-BN-male-bio1-tech2.bam
## 2 lv-input-BN-male-bio1-tech1.bam|lv-input-BN-male-bio1-tech2.bam
## 3                                lv-input-SHR-male-bio1-tech1.bam
## 4                                lv-input-SHR-male-bio1-tech1.bam
## 5                                                            <NA>
## 6                                                            <NA>
## 7                                                            <NA>
## 8                                                            <NA>

## Define assembly
# This is only necessary if you have BED files, BAM files are handled automatically.
# For common assemblies you can also specify them as 'hg19' for example.
data(rn4_chrominfo)
head(rn4_chrominfo)

##    chromosome   length
## 1        chrM    16300
## 2       chr12 46782294
## 3       chr20 55268282
## 4       chr19 59218465
## 5       chr18 87265094
## 6       chr11 87759784

#=== Step 2: Run Chromstar ===
## Run ChromstaR
Chromstar(inputfolder, experiment.table=experiment_table,
          outputfolder=outputfolder, numCPU=4, binsize=1000, stepsize=500,
          assembly=rn4_chrominfo, prefit.on.chr='chr12', chromosomes='chr12',
          mode='differential')

## Results are stored in 'outputfolder' and can be loaded for further processing
list.files(outputfolder)

##  [1] "BROWSERFILES"       "PLOTS"             "README.txt"        "binned"
##  [5] "chrominfo.tsv"      "chromstaR.config"  "combined"          "experiment_table.tsv"
##  [9] "multivariate"       "univariate"

model <- get(load(file.path(outputfolder,'combined',
      'combined_mode-differential_binsize1000_stepsize500.RData')))
```

!! It is important that the distributions in folder outputfolder/PLOTS/univariate-distributions are fitted correctly !!
Please check section 7.4 for examples of how this plot should *not* look like and what can be done to get a correct fit.

```
#=== Step 3: Analysis and export ===
## Obtain all genomic regions where the two tissues have different states
segments <- model$segments
diff.segments <- segments[segments$combination.SHR != segments$combination.BN]
# Let's be strict with the differences and get only those where both marks are different
diff.segments <- diff.segments[diff.segments$differential.score >= 1.9]
exportGRangesAsBedFile(diff.segments, trackname='differential_chromatin_states',
             filename=tempfile(), scorecol='differential.score')

## Warning in exportGRangesAsBedFile(diff.segments, trackname = "differential_chromatin_states", :  Column 'differential.score'
## should contain integer values between 0 and 1000 for compatibility with the UCSC convention.
```

```
## Obtain all genomic regions where we find a bivalent promoter in BN but not in SHR
bivalent.BN <- segments[segments$combination.BN == '[H3K27me3+H3K4me3]' &
                        segments$combination.SHR != '[H3K27me3+H3K4me3]']
## Obtain all genomic regions where BN and SHR have promoter signatures
promoter.BN <- segments[segments$transition.group == 'constant [H3K4me3]']

## Get transition frequencies
print(model$frequencies)

##         combination.BN   combination.SHR domains    frequency cumulative.frequency
## 1            [H3K27me3]        [H3K27me3]    1326 4.367492e-01            0.4367492
## 2                    []                []    1324 4.197982e-01            0.8565474
## 3             [H3K4me3]         [H3K4me3]     849 8.564191e-02            0.9421893
## 4  [H3K27me3+H3K4me3] [H3K27me3+H3K4me3]     849 5.279808e-02            0.9949874
## 5            [H3K27me3]                []      19 1.934505e-03            0.9969219
## 6                    []        [H3K27me3]      17 1.752811e-03            0.9986747
## 7            [H3K27me3] [H3K27me3+H3K4me3]      16 4.702663e-04            0.9991450
## 8                    []         [H3K4me3]      12 2.992604e-04            0.9994442
## 9  [H3K27me3+H3K4me3]        [H3K27me3]       5 2.030696e-04            0.9996473
## 10            [H3K4me3] [H3K27me3+H3K4me3]      5 1.175666e-04            0.9997649
## 11            [H3K4me3]                []      3 7.481510e-05            0.9998397
## 12 [H3K27me3+H3K4me3]                []       1 6.412723e-05            0.9999038
## 13                   [] [H3K27me3+H3K4me3]      1 6.412723e-05            0.9999679
## 14           [H3K27me3]         [H3K4me3]      1 3.206361e-05            1.0000000
##                            group
## 1              constant [H3K27me3]
## 2                  zero transition
## 3               constant [H3K4me3]
## 4       constant [H3K27me3+H3K4me3]
## 5             stage-specific [H3K27me3]
## 6             stage-specific [H3K27me3]
## 7                            other
## 8             stage-specific [H3K4me3]
## 9                            other
## 10                           other
## 11            stage-specific [H3K4me3]
## 12 stage-specific [H3K27me3+H3K4me3]
## 13 stage-specific [H3K27me3+H3K4me3]
## 14                           other
```

```
## === Step 4: Enrichment analysis ===
# Annotations can easily be obtained with the biomaRt package. Of course, you can
# also load them from file if you already have annotation files available.
library(biomaRt)
ensembl <- useEnsembl(biomart='ENSEMBL_MART_ENSEMBL', dataset='rnorvegicus_gene_ensembl')
genes <- getBM(attributes=c('ensembl_gene_id', 'chromosome_name', 'start_position',
                            'end_position', 'strand', 'external_gene_name',
                            'gene_biotype'),
               mart=ensembl)
# Transform to GRanges for easier handling
genes <- GRanges(seqnames=paste0('chr',genes$chromosome_name),
                 ranges=IRanges(start=genes$start, end=genes$end),
                 strand=genes$strand,
                 name=genes$external_gene_name, biotype=genes$gene_biotype)
# Rename chrMT to chrM to avoid warnings
seqlevels(genes)[seqlevels(genes)=='chrMT'] <- 'chrM'
# Select only chr12 to avoid warnings
genes <- keepSeqlevels(genes, 'chr12', pruning.mode = 'coarse')
print(genes)

## GRanges object with 919 ranges and 2 metadata columns:
##         seqnames              ranges strand |           name        biotype
##            <Rle>           <IRanges>  <Rle> |    <character>    <character>
##    [1]     chr12 24473981-24476295      + |           Fzd9 protein_coding
##    [2]     chr12   9693779-9727489      - |         Polr1d protein_coding
##    [3]     chr12 23587222-23592774      - |         Polr2j protein_coding
##    [4]     chr12 48292273-48351084      + |           Svop protein_coding
##    [5]     chr12 47074200-47076572      + |         Dynll1 protein_coding
##    ...       ...                 ...    ... .            ...            ...
##  [915]     chr12   6016884-6078411      - |           Rxfp2 protein_coding
##  [916]     chr12   2938358-2940393      - | AABR07034991.1        lincRNA
##  [917]     chr12 25036605-25070539      + |           Limk1 protein_coding
##  [918]     chr12   9728486-9793428      + |           Lnx2 protein_coding
##  [919]     chr12 20667601-20677689      + |     RGD1560281 protein_coding
##   -------
##   seqinfo: 1 sequence from an unspecified genome; no seqlengths
```

```
# We expect promoter [H3K4me3] and bivalent-promoter signatures [H3K4me3+H3K27me3]
# to be enriched at transcription start sites.
plots <- plotEnrichment(hmm=model, annotation=genes, region='start')
plots[['BN']] + facet_wrap(~ combination) +
  ggtitle('Fold enrichment around TSS') +
  xlab('distance from TSS')
```


Fold enrichment around TSS
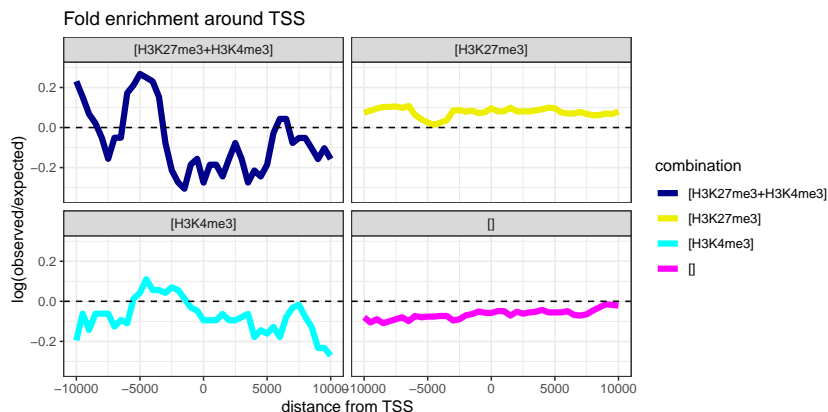
```
plots <- plotEnrichment(hmm=model, annotation=genes, region='start', what='peaks')
plots[['BN']] + facet_wrap(~ mark) +
  ggtitle('Fold enrichment around TSS') +
  xlab('distance from TSS')
```
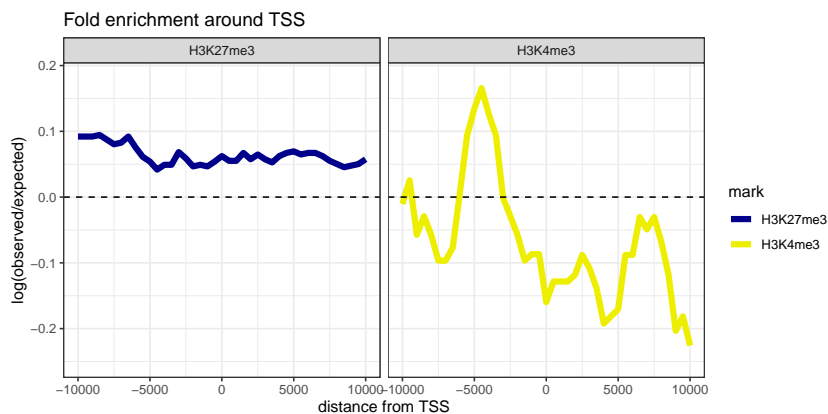

Fold enrichment around TSS

```
# Fold enrichment with different biotypes, showing that protein coding genes are
# enriched with (bivalent) promoter combinations [H3K4me3] and [H3K4me3+H3K27me3],
# while rRNA is enriched with the empty [] combination.
tss <- resize(genes, width = 3, fix = 'start')
biotypes <- split(tss, tss$biotype)
plots <- plotFoldEnrichHeatmap(model, annotations=biotypes)
plots[['BN']] + coord_flip() +
  ggtitle('Fold enrichment with different biotypes')

## Warning in regularize.values(x, y, ties, missing(ties), na.rm = na.rm):  collapsing to unique 'x' values

## Warning in regularize.values(x, y, ties, missing(ties), na.rm = na.rm):  collapsing to unique 'x' values

## Warning in regularize.values(x, y, ties, missing(ties), na.rm = na.rm):  collapsing to unique 'x' values
```

Fold enrichment with different biotypes
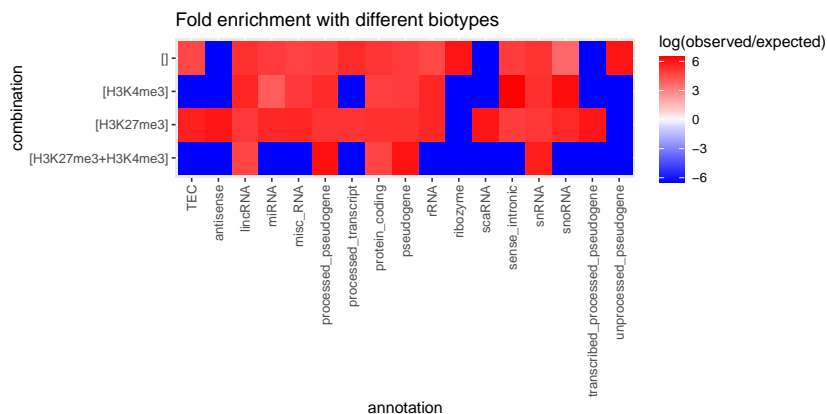
6 Output of function `Chromstar()`

`Chromstar()` is the workhorse of the *chromstaR* package and produces all the files that are necessary for downstream analyses. Here is an explanation of the *files* and **folders** you will find in your **outputfolder**:

- *chrominfo.tsv*:
  A tab-separated file with chromosome lengths.

- *chromstaR.config*:
  A text file with all the parameters that were used to run function `Chromstar()`.

- *experiment_table.tsv*:
  A tab-separated file of your experiment setup.

- **binned**:
  RData files with the results of the binnig step. Contains *GRanges* objects with binned genomic coordinates and read counts.

- **BROWSERFILES**:
  Bed files for upload to the UCSC genome browser. It contains files with combinatorial states ("_combinations.bed.gz") and underlying peak calls ("_peaks.bed.gz"). !!Always check the "_peaks.bed.gz" files if you are satisfied with the peak calls. If not, there are ways to make the calls stricter (see section 7.1).

- →**combined**←:
  RData files with the combined results of the uni- and multivariate peak calling steps. This is what you want to use for downstream analyses. Contains *combinedMultiHMM* objects.

  - *combined_mode-separate.RData*: Simple combination of peak calls (replicates considered) without multivariate analysis.

  - *combined_mode-combinatorial.RData*: Combination of multivariate results for mode='combinatorial'.

  - *combined_mode-differential.RData*: Combination of multivariate results for mode='differential'.

  - *combined_mode-full.RData*: Combination of multivariate results for mode='full'.

- **multivariate**:
  RData files with the results of the multivariate peak calling step. Contains *multiHMM* objects.

- **PLOTS**:
  Several plots that are produced by default. Please check the plots in subfolder **univariate-distributions** for irregularities (see section 4).

- **replicates**:
  RData files with the result of the replicate peak calling step. Contains *multiHMM* objects.

- **univariate**:
  RData files with the result of the univariate peak calling step. Contains *uniHMM* objects.

7 FAQ

7.1 The peak calls are too lenient. Can I adjust the strictness of the peak calling?

The strictness of the peak calling can be controlled with a cutoff on the posterior probability. The Hidden Markov Model gives posterior probabilities for each peak, and based on these probabilites the model decides if a peak is present or not by picking the state with the highest probability. This way of peak calling leads to very lenient peak calls, and for some applications it may be desirable to obtain only very clear peaks. This can be achieved by using `changePostCutoff` and `changeMaxPostCutoff`. `changePostCutoff` applies a cutoff on the posteriors in each bin, which will make peaks narrower but might also lead to fragmented peaks in the case of broad peaks. `changeMaxPostCutoff` applies a cutoff on the maximum posterior within each peak, which will preserve broad peaks. To follow the below example, please first run step 1 and 2 from section 5.4.

```r
model <- get(load(file.path(outputfolder,'combined',
    'combined_mode-differential_binsize1000_stepsize500.RData')))
# Try a strict cutoff close to 1
model2 <- changeMaxPostCutoff(model, maxPost.cutoff=0.99999)
model3 <- changePostCutoff(model, post.cutoff=0.99999)
# Check the peaks before and after adjustment
plots <- plotGenomeBrowser(model, chr='chr12', start=1, end=3e5)
plots2 <- plotGenomeBrowser(model2, chr='chr12', start=1, end=3e5)
plots3 <- plotGenomeBrowser(model3, chr='chr12', start=1, end=3e5)
plots$`H3K27me3-BN-rep1` + ggtitle('H3K27me3 original')
```
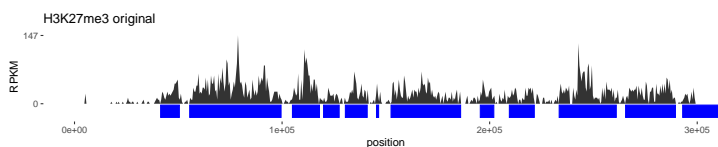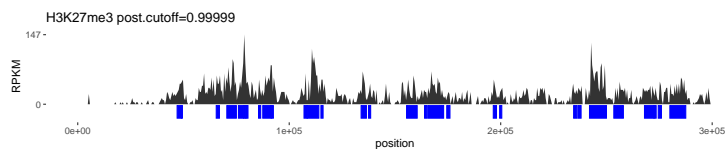


```r
plots2$`H3K27me3-BN-rep1` + ggtitle('H3K27me3 maxPost.cutoff=0.99999')
```
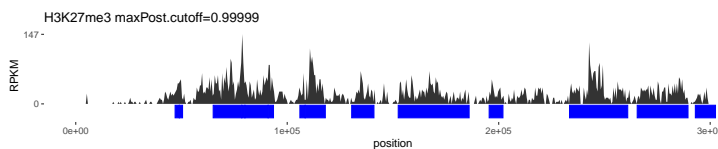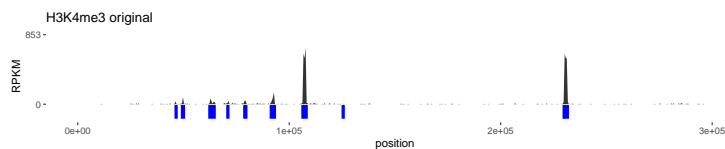


```r
plots3$`H3K27me3-BN-rep1` + ggtitle('H3K27me3 post.cutoff=0.99999')
```
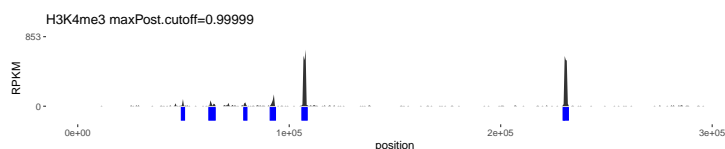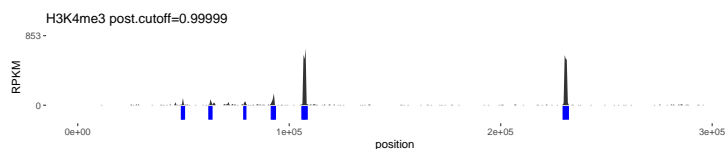
H3K27me3 post.cutoff=0.99999


```
plots$`H3K4me3-BN-rep1` + ggtitle('H3K4me3 original')
```

H3K4me3 original


```
plots2$`H3K4me3-BN-rep1` + ggtitle('H3K4me3 maxPost.cutoff=0.99999')
```

H3K4me3 maxPost.cutoff=0.99999


```
plots3$`H3K4me3-BN-rep1` + ggtitle('H3K4me3 post.cutoff=0.99999')
```

H3K4me3 post.cutoff=0.99999

It is even possible to adjust the sensitivity differently for the different marks or conditions.

```
# Set a stricter cutoff for H3K4me3 than for H3K27me3
cutoffs <- c(0.9, 0.9, 0.9, 0.9, 0.99, 0.99, 0.99, 0.99)
names(cutoffs) <- model$info$ID
print(cutoffs)

##  H3K27me3-BN-rep1  H3K27me3-BN-rep2   H3K4me3-BN-rep1   H3K4me3-BN-rep2 H3K27me3-SHR-rep1 H3K27me3-SHR-rep2
##              0.90              0.90              0.90              0.90              0.99              0.99
##  H3K4me3-SHR-rep1  H3K4me3-SHR-rep2
##              0.99              0.99

model2 <- changeMaxPostCutoff(model, maxPost.cutoff=cutoffs)
```

7.2 The combinatorial differences that chromstaR gives me are not convincing. Is there a way to restrict the results to a more convincing set?

You were interested in combinatorial state differences as in section 5.4 and checked the results in a genome browser. You found that some differences are convincing by eye and some are not. There are several possibilities to explore:

1. Run `Chromstar` in mode='differential' (instead of mode='combinatorial') and see if the results improve.

2. You can play with the "differential.score" (see section 5.4, step 3) and export only differences with a high score. A differential score close to 1 means that one modification is different, a score close to 2 means that two modifications are different etc. The score is calculated as the sum of differences in posterior probabilities between marks.

3. Use `changePostCutoff` or `changeMaxPostCutoff` to obtain only high confidence peaks.

4. Check for bad replicates that are very different from the rest and exclude them prior to the analysis.

## 7.3 How do I plot a simple heatmap with the combinations?

```
heatmapCombinations(marks=c('H3K4me3', 'H3K27me3', 'H3K36me3', 'H3K27Ac'))

## Warning: It is deprecated to specify 'guide = FALSE' to remove a guide.  Please use 'guide = "none"' instead.
```
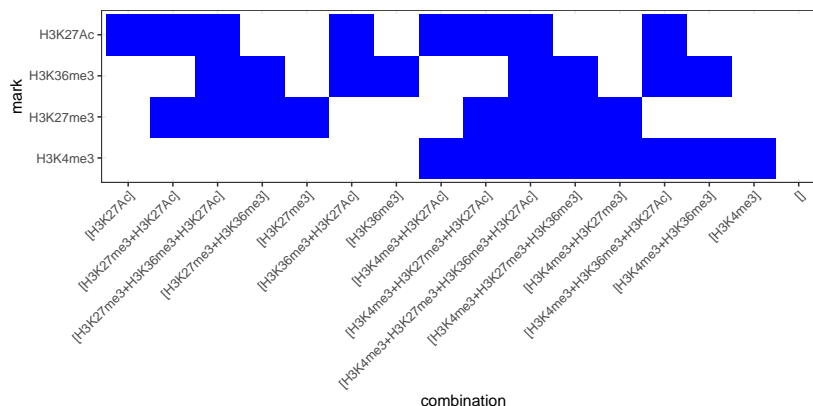

7.4 Examples of problematic distributions.

For the chromstaR peak calling to work correctly it is essential that the Baum-Welch algorithm correctly identifies unmodified (background) and modified (signal/peak) components in the data. Therefore, you should always check the plots in folder **PLOTS/univariate-distributions** for correct convergence. Here are some plots that indicate failed and succesful fitting procedures:
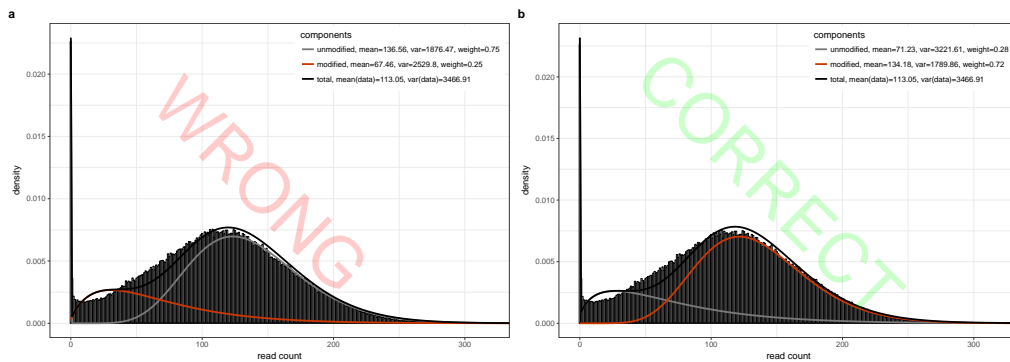
The plot shows data for H3K27me3 at binsize 1000bp. (a) Incorrectly converged fit, where the **modified** component (red) has lower read counts than the **unmodified** component (gray). (b) Correctly converged fit. Even here, the fit could be improved by reducing the average number of reads per bin, either by selecting a smaller binsize or by downsampling the data before using chromstaR.
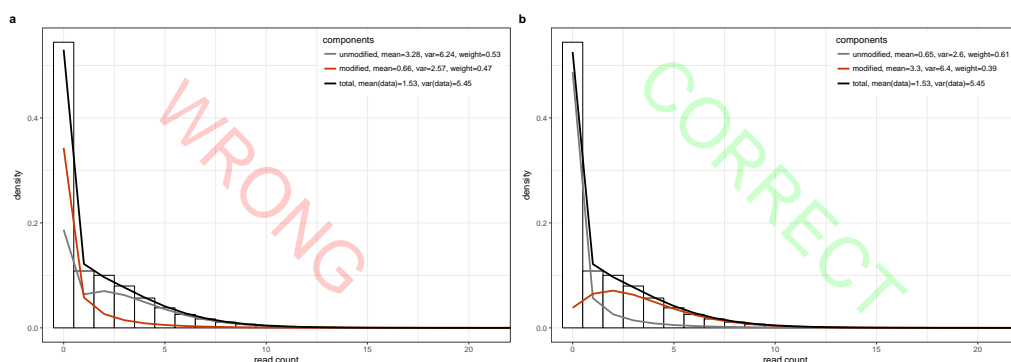
The plot shows data for H3K27me3 at binsize 150bp. (a) Incorrectly converged fit, where the **modified** component (red) has a higher density at zero reads than the **unmodified** component (gray). (b) Correctly converged fit.

8   Session Info

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

- R version 4.1.1 (2021-08-10), x86_64-pc-linux-gnu

- Locale: LC_CTYPE=en_US.UTF-8, LC_NUMERIC=C, LC_TIME=en_GB, LC_COLLATE=C, LC_MONETARY=en_US.UTF-8, 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

- Running under: Ubuntu 20.04.3 LTS

- Matrix products: default

- BLAS: /home/biocbuild/bbs-3.14-bioc/R/lib/libRblas.so

- LAPACK: /home/biocbuild/bbs-3.14-bioc/R/lib/libRlapack.so

- Base packages: base, datasets, grDevices, graphics, methods, stats, stats4, utils

- Other packages: BiocGenerics 0.40.0, GenomeInfoDb 1.30.0, GenomicRanges 1.46.0, IRanges 2.28.0, S4Vectors 0.32.2, biomaRt 2.50.0, chromstaR 1.20.2, chromstaRData 1.20.0, ggplot2 3.3.5

- Loaded via a namespace (and not attached): AnnotationDbi 1.56.1, Biobase 2.54.0, BiocFileCache 2.2.0, BiocManager 1.30.16, BiocParallel 1.28.0, BiocStyle 2.22.0, Biostrings 2.62.0, DBI 1.1.1, DelayedArray 0.20.0, GenomeInfoDbData 1.2.7, GenomicAlignments 1.30.0, KEGGREST 1.34.0, Matrix 1.3-4, MatrixGenerics 1.6.0, R6 2.5.1, RCurl 1.98-1.5, RSQLite 2.2.8, Rcpp 1.0.7, Rsamtools 2.10.0, SummarizedExperiment 1.24.0, XML 3.99-0.8, XVector 0.34.0, assertthat 0.2.1, bamsignals 1.26.0, bit 4.0.4, bit64 4.0.5, bitops 1.0-7, blob 1.2.2, cachem 1.0.6, codetools 0.2-18, colorspace 2.0-2, compiler 4.1.1, crayon 1.4.2, curl 4.3.2, dbplyr 2.1.1, digest 0.6.28, doParallel 1.0.16, dplyr 1.0.7, ellipsis 0.3.2, evaluate 0.14, fansi 0.5.0, farver 2.1.0, fastmap 1.1.0, filelock 1.0.2, foreach 1.5.1, generics 0.1.1, glue 1.5.0, grid 4.1.1, gtable 0.3.0, highr 0.9, hms 1.1.1, htmltools 0.5.2, httr 1.4.2, iterators 1.0.13, knitr 1.36, labeling 0.4.2, lattice 0.20-45, lifecycle 1.0.1, magrittr 2.0.1, matrixStats 0.61.0, memoise 2.0.0, munsell 0.5.0, mvtnorm 1.1-3, parallel 4.1.1, pillar 1.6.4, pkgconfig 2.0.3, plyr 1.8.6, png 0.1-7, prettyunits 1.1.1, progress 1.2.2, purrr 0.3.4, ragg 1.2.0, rappdirs 0.3.3, reshape2 1.4.4, rlang 0.4.12, rmarkdown 2.11, scales 1.1.1, stringi 1.7.5, stringr 1.4.0, systemfonts 1.0.3, textshaping 0.3.6, tibble 3.1.6, tidyselect 1.1.1, tools 4.1.1, utf8 1.2.2, vctrs 0.3.8, withr 2.4.2, xfun 0.28, xml2 1.3.2, yaml 2.2.1, zlibbioc 1.40.0

References

[1] Aaron Taudt, Minh Anh Nguyen, Matthias Heinig, Frank Johannes, and Maria Colome-Tatche. chromstaR: Tracking combinatorial chromatin state dynamics in space and time. Technical report, feb 2016. URL: http://biorxiv.org/content/early/2016/02/04/038612.abstract, doi:10.1101/038612.

[2] Courtney W. Hanna, Aaron Taudt, Jiahao Huang, Lenka Gahurova, Andrea Kranz, Simon Andrews, Wendy Dean, A. Francis Stewart, Maria Colomé-Tatché, and Gavin Kelsey. MLL2 conveys transcription-independent H3K4 trimethylation in oocytes. *Nature Structural & Molecular Biology*, page 1, jan 2018. URL: http://www.nature.com/articles/s41594-017-0013-5, doi:10.1038/s41594-017-0013-5.