

# Detection of de novo copy number alterations in case-parent trios using the R package `MinimumDistance`

Rob Scharpf

May 2, 2019

## Abstract

For the analysis of case-parent trio genotyping arrays, copy number variants (CNV) appearing in the offspring that differ from the parental copy numbers are often of interest (de novo CNV). This package defines a statistic, referred to as the minimum distance, for identifying de novo copy number alterations in the offspring. We smooth the minimum distance using the circular binary segmentation algorithm implemented in the Bioconductor package `DNAcopy`. Trio copy number states are inferred from the maximum a posteriori probability of the segmented data, incorporating information from the log R ratios and B allele frequencies. As both log R ratios and B allele frequencies are estimable from Illumina and Affymetrix arrays, this package supports de novo copy number inference in both platforms.

## 1 Introduction

There are numerous R packages available from Bioconductor for smoothing copy number alterations. For example, the biocview `CopyNumberVariants` in the 2.9 release of Bioconductor lists 27 packages. For the analysis of germline diseases, hidden Markov models have emerged as a popular tool as inference regarding the latent copy number state incorporates information from both the estimates of copy number (or relative copy number) and the allelic frequencies, such as the B allele frequency [4] or genotype calls [1, 7, 5]. For the analysis of somatic cell diseases such as cancer, algorithms that segment the genome into regions of constant copy number (referred to here as segmentation algorithms) may be more preferable for the detection of copy number aberrations (CNA) as a mixture of cells with different copy numbers give rise to mosaic (non-integer) copy numbers. Examples include circular binary segmentation implemented in the R package `DNAcopy` and the `GLAD`, both of which were originally developed for array CGH platforms [3, 2, 6]. One disadvantage of segmentation algorithms is that inference regarding duplication and deletions is not directly available.

More recently, HMMs and segmentation algorithms have been developed for inferring common regions of copy number alterations in multiple samples. However, relatively few algorithms are available for inferring copy number alterations, especially de novo copy number alterations, in family-based study designs involving case-parent trios. Instead, a common strategy has been a two-step approach of identifying copy number alterations in the individual samples and then comparing the results across samples to infer whether an alteration observed in the offspring is de novo. A disadvantage of the two-step approach is that unadjusted sources of technical variation, such as waves, can contribute to false positives. The joint HMM implemented in the `PennCNV` software is one of the few algorithms to date that provides direct inference regarding de novo alterations in case parent study designs.

This package develops an alternative framework for inferring regions of de novo copy number alterations in case-parent trios. Like the `PennCNV` joint HMM, inference regarding de novo alterations integrates information from both the log R ratios and B allele frequencies. Differences in the two approaches are most apparent in non-HapMap experimental datasets in which technical and experimental sources of variation appear to contribute to a large number of false positives. This vignette describes the analysis pipeline from preprocessed and normalized estimates of copy number and allele frequencies to inference of de novo copy number alterations. The workflow is illustrated using a publicly available HapMap trio and a publicly available oral cleft trio, each assayed on the Illumina 610quad array.

## 2 Data input

```
R> library(oligoClasses)
R> library(VanillaICE)
R> library(MinimumDistance)
R> foreach::registerDoSEQ()
```

### 2.1 Reading and organizing the annotation on the markers

We require that marker-level annotation is represented as a *GRanges*-derived class. To read the plain text annotation file, we use the `fread` function provided in the `data.table` package. In addition, we define an indicator for whether the marker is polymorphic using the 'Intensity Only' flag. The *SnpGRanges* class created from the `fgr` object in the following code-chunk ensures that this binary indicator is created and can be reliably accessed.

```
R> library(data.table)
R> extdir <- system.file("extdata", package="VanillaICE")
R> features <- suppressWarnings(fread(file.path(extdir, "SNP_info.csv")))
R> fgr <- GRanges(paste0("chr", features$Chr), IRanges(features$Position, width=1),
  isSnp=features[["Intensity Only"]==0)
R> fgr <- SnpGRanges(fgr)
R> names(fgr) <- features[["Name"]]
```

Ideally, one should include the genome build and the chromosome lengths appropriate to the build. Here, we extract the metadata on the chromosomes using the *BSgenome* Bioconductor package for the hg18 build. Finally, we sort the `fgr` object such that the chromosomes are ordered by their `seqlevels` and the markers are ordered by their genomic position along the chromosome.

```
R> library(BSgenome.Hsapiens.UCSC.hg18)
R> sl <- seqlevels(BSgenome.Hsapiens.UCSC.hg18)
R> seqlevels(fgr) <- sl[sl %in% seqlevels(fgr)]
R> seqinfo(fgr) <- seqinfo(BSgenome.Hsapiens.UCSC.hg18)[seqlevels(fgr),]
R> fgr <- sort(fgr)
```

### 2.2 Organizing the marker-level summaries

The abbreviated plain text files included with this package that contain the log R ratios and B allele frequencies from 2 trios are listed below.

```
R> files <- list.files(extdir, full.names=TRUE, recursive=TRUE, pattern="FinalReport")
```

We parse these files into a specific format such that all downstream steps for copy number estimation no longer depend on the format of the source files. At this point, we encapsulate the names of the source files ('sourcePaths'), the location of where we intend to store the parsed data ('parsedPath'), and the genomic marker annotation created in the preceding section ('rowRanges') in a single object called an *ArrayViews*.

```
R> ##
R> ## Where to keep parsed files
R> ##
R> parsedDir <- "ParsedFiles"
R> if(!file.exists(parsedDir)) dir.create(parsedDir)
R> views <- ArrayViews(rowRanges=fgr, sourcePaths=files, parsedPath=parsedDir)
R> show(views)
```

```
class 'ArrayViews'
  No. files : 6
  No. markers: 11780
```

Because the format of the source files depends on upstream software, we read one file and store information regarding its parsing so that subsequent files can be parsed in a similar fashion. We use the `fread` to read in the first file.

```
R> ## read the first file
R> dat <- fread(files[1], skip="[Data]")
R> head(dat,n=3)

      SNP Name Allele1 - AB Allele2 - AB B Allele Freq
1: cnvi0005126      -      -      0.0085
2: cnvi0005127      -      -      0.0081
3: cnvi0005128      -      -      0.0063
Log R Ratio
1:   -0.7214
2:   -0.8589
3:   -0.1640
```

Next, we select which columns we plan to keep. Again, the required data for downstream processing is the name of the SNP identifier, the log R ratios, and B allele frequencies.

```
R> ## information to store on the markers
R> select_columns <- match(c("SNP Name", "Allele1 - AB", "Allele2 - AB",
                           "Log R Ratio", "B Allele Freq"), names(dat))
```

We also specify the order in which we will store the marker-level summaries by matching the `rownames` of the `views` object with the names of the markers in the source file:

```
R> index_genome <- match(names(fgr), dat[["SNP Name"]])
```

Similar to the parameter classes defined in `Rsamtools`, we encapsulate the information for parsing the columns and rows of the source files in a class. In addition, we specify which variable names in the source file refers to log R ratios (`'cnvar'`), B allele frequencies (`'bafvar'`), and genotypes (`'gtvar'`).

```
R> scan_params <- CopyNumScanParams(index_genome=index_genome,
                                   select=select_columns,
                                   cnvar="Log R Ratio",
                                   bafvar="B Allele Freq",
                                   gtvar=c("Allele1 - AB", "Allele2 - AB"))
```

The `parseSourceFile` will parse a single file in the `views` object (by default, the first file) according to the parameters for reading the data in the `scan_params` object and write the output to the `parsedPath` directory. In particular, the `parseSourceFile` returns `NULL`.

```
R> parsedPath(views)
[1] "ParsedFiles"
R> parseSourceFile(views[, 1], scan_params)
```

To apply the function `parseSourceFile` to all arrays in the `views` object, one can use the functional `sapply`.

```
R> invisible(sapply(views, parseSourceFile, param=scan_params))
```

Apart from confirming their existence, the user should not have a need to directly access the parsed files. Utilities for querying these files are provided through the `views` object.

```
R> head(list.files(parsedPath(views)), n=3)
[1] "FinalReport1664_baf.rds" "FinalReport1664_gt.rds"
[3] "FinalReport1664_lrr.rds"
```

## 2.3 Accessors for the parsed data

The preference for writing the parsed data to disk rather than keeping the data in RAM is simply that the latter does not scale to projects involving thousands of samples. For the former, slices of the parsed data easily be accessed from the `parsedPath` directory via methods defined for the `ArrayViews` class. For example, one can use accessors for the low-level summaries directly: `lrr`, `baf`, and `genotypes` for log R ratios, B allele frequencies, and genotypes, respectively. The user has the option of either subsetting the views object or subsetting the matrix returned by the accessor to extract the appropriate data slice. In the following examples, we access data on the first 2 markers and sample indices 2-4.

```
R> lrr(views)[1:2, 2:4]
```

```
          FinalReport1675.txt FinalReport1686.txt
rs12789205          -0.169           0.191
rs2114088           -0.182           -0.114
          FinalReport6841.txt
rs12789205          -0.084
rs2114088            0.253
```

```
R> ## or
```

```
R> lrr(views[1:2, 2:4])
```

```
          FinalReport1675.txt FinalReport1686.txt
rs12789205          -0.169           0.191
rs2114088           -0.182           -0.114
          FinalReport6841.txt
rs12789205          -0.084
rs2114088            0.253
```

```
R> ## B allele frequencies
```

```
R> baf(views[1:2, 2:4])
```

```
          FinalReport1675.txt FinalReport1686.txt
rs12789205           0.529           0.549
rs2114088            0.000           0.000
          FinalReport6841.txt
rs12789205           0.481
rs2114088            0.000
```

```
R> ## potentially masked by function of the same name in crlmm
```

```
R> VanillaICE::genotypes(views)[1:2, 2:4]
```

```
          FinalReport1675.txt FinalReport1686.txt
rs12789205                2                2
rs2114088                 1                1
          FinalReport6841.txt
rs12789205                2
rs2114088                 1
```

More often, it is useful to extract the low-level data in a `RangedSummarizedExperiment`-derived class such that meta-data on the samples remains bound to the columns of the assay data (log R ratios / B allele frequencies) and meta-data on the rows remains bound to the rows of the assay data. This is accomplished by applying the `MinDistExperiment` function to a `views` object, as described in the following section.

## 3 Binding the genomic, marker, and individual-level data

### 3.1 Pedigree annotation

The container for a single pedigree is the *ParentOffspring* class.

```
R> ped_hapmap <- ParentOffspring(id = "hapmap", father="12287_03",
                                mother="12287_02",
                                offspring="12287_01",
                                parsedPath=parsedPath(views))
```

For families with multiple affected offspring, the argument to offspring can be a character-vector with length greater than 1. We can store any number of *ParentOffspring* objects in a list using the *ParentOffspringList* class. In particular, for the 2 trios provided in the *VanillaICE* package

```
R> ped_list <- ParentOffspringList(pedigrees=list(
  ParentOffspring(id = "hapmap", father="12287_03",
                  mother="12287_02",
                  offspring="12287_01",
                  parsedPath=parsedPath(views)),
  ParentOffspring(id = "cleft",
                  father="22169_03",
                  mother="22169_02",
                  offspring="22169_01",
                  parsedPath=parsedPath(views))))
```

```
R> pedigreeName(ped_list)
```

```
[1] "hapmap" "cleft"
```

For a single pedigree, the *MinDistExperiment* encapsulates the annotation on the pedigree, the genomic annotation of the markers, and the marker-level summary statistics. Note, however, that the sample identifiers in the pedigree are not the same as the file identifiers used by default to create the R object *views*. To resolve the naming issue, we read in separate file provided in the *VanillaICE* package:

```
R> sample_info <- read.csv(file.path(extdir, "sample_data.csv"), stringsAsFactors=FALSE)
R> ind_id <- setNames(gsub(" ", "", sample_info$IndividualID), sample_info$File)
R> colnames(views) <- ind_id[gsub(".txt", "", colnames(views))]
```

### 3.2 MinDistExperiment class

The constructor function for *MinDistExperiment* extracts from the *views* object only the data relevant for the provided pedigree.

```
R> me <- MinDistExperiment(views, pedigree=ped_list[[2]])
```

```
R> colnames(me)
```

```
[1] "22169_03" "22169_02" "22169_01"
```

```
R> me
```

```
class: MinDistExperiment
dim: 11780 3
metadata(0):
assays(2): cn baf
rownames(11780): rs12789205 rs2114088 ... rs5994329
               rs1055232
rowData names(1): isSnp
colnames(3): 22169_03 22169_02 22169_01
colData names(0):
```

## 4 Detection of de novo copy number variants

A container for the various parameters used to segment and call de novo copy number variants is provided by the *MinDistParam*. A constructor function of the same name provides a default parametrization that works well in most instances:

```
R> params <- MinDistParam()
R> show(params)
```

```
An object of class 'MinDistParam'
call segments with |seg.mean|/MAD > 0.75
Setting nMAD() to smaller values will increase the number of segments that are called.
DNACopy settings:
  alpha: 0.01
  min.width: 2
  undo.splits: none
  undo.SD: 3
  See segment() for description of DNACopy parameters
```

The parameters in this class are organized according to function. For example, the argument `dnacopy` takes an argument of class *DNACopyParam* that contains settings that are passed to the `segment` function for the implementation of circular binary segmentation in the R package *DNACopy*. Changing the parameters for the segmentation is most easily accomplished by a call to the constructor. E.g.,

```
R> segment_params <- DNACopyParam(alpha=0.01)
R> params <- MinDistParam(dnacopy=segment_params)
```

Several of the parameters relate to a priori assumptions of the probability of a non-Mendelian transmission. These a priori assumptions are defined in *PennCNV* [7] and encapsulated in the *PennParam* class.

```
R> penn_param <- PennParam()
R> show(penn_param)
```

```
Object of class `PennParam'
121 trio states
state names: 000,100,200,300,400,010 ...
reference state: 222
PennCNV Table 1: 121 vector
PennCNV Table 3: 5 x 5 x 5 x 5 x 5 array
probability non-Mendelian: 1.5e-06
initial state probabilities: 0.2,0.2,0.2,0.2,0.2
transition prob: 5 x 5 matrix
See table1(), table3(), state(), stateNames(), referenceState()
```

Finally, parameters for the emission probabilities computed by the R package *VanillaICE* are encapsulated in the *EmissionParam* class. Again, default values for the emission probabilities will be created automatically as part of the `params` object if none is specified.

**Segmentation and posterior calls.** For a given trio, the signed minimum absolute difference of the offspring and parental  $\log_2$  R ratios ( $\mathbf{r}$ ) is defined as

$$\mathbf{d} \equiv (\mathbf{r}_O - \mathbf{r}_M) \times \mathbb{I}_{|\mathbf{r}_O - \mathbf{r}_F| > |\mathbf{r}_O - \mathbf{r}_M|} + (\mathbf{r}_O - \mathbf{r}_F) \times \mathbb{I}_{|\mathbf{r}_O - \mathbf{r}_F| \leq |\mathbf{r}_O - \mathbf{r}_M|}. \quad (1)$$

If the offspring copy number changes within a minimum distance segment (as determined by the segmentation of the offspring copy number), the start and stop position of the minimum distance segments may be edited. The approach currently implemented is to define a new start and stop position if a breakpoint for the offspring segmentation occurs in the minimum distance interval. To illustrate, the following diagram uses vertical dashes (|) to denote breakpoints:

```

1 ...--|-----|---...    ## minimum distance segment (before editing)
2 ...---|-----|-----... ## segmenation of log R ratios for offspring

->

3 ...--|-----|---|---...    ## after editing

```

In the above illustration, posterior calls are provided for the 3 segments indicated in line 3 instead of the single segment in line 1. Two additional steps are therefore required: (1) segmentation of the offspring log R ratios and (2) editing of the minimum distance breakpoints when appropriate. The following codechunk, segments the log R ratios for all chromosomes in the `me` object and edits the ranges for the minimum distance when conflicts with the offspring segmentation boundaries arise (as illustrated above).

```
R> mdgr <- segment2(me, params)
```

```

Analyzing: X22169_03
Analyzing: X22169_02
Analyzing: X22169_01
Analyzing: md_22169_01

```

For each minimum distance segment in which the mean minimum distance is above a user-specified cutoff in absolute value (specified in terms of the median absolute deviations of the minimum distance), a trio copy number state is assigned from the maximum a posteriori estimate.

```

R> ## the threshold in terms of the number of median absolute deviations from zero
R> nMAD(params)
R> md_g <- MAP2(me, mdgr, params)
R> show(md_g)

```

## 5 Inspecting, Filtering, and plotting de novo CNV inference

### 5.1 Filtering

There are several meta-data columns stored in the *GRanges*-derived summary useful for filtering the set of genomic intervals. All the available parameters for filtering the `fit` object are stored in the parameter class *FilterParam* of which an instance can be created by its constructor of the same name without any arguments.

```

R> filter_param <- FilterParamMD()
R> show(filter_param)

```

```

An object of class 'FilterParamMD'
 min. posterior probability of trio CNV call: 0.99
 min. no. of markers spanned by segment      : 10
 min. width of segment                       : 1
 121 selected trio CN states                 : 000, 100, 200, 300, 400, 010, ...
 selected seqnames                           : chr1, chr2, chr3, chr4, chr5, chr6 ...

```

To apply the default filter parameters to the `fit` object, we use the `cnvFilter` function. The `cnvFilter` function returns only the set of genomic ranges satisfying the filters. For example,

```
R> cnvFilter(md_g, filter_param)
```

```

MDRanges object with 5 ranges and 9 metadata columns:
  seqnames      ranges strand |      sample  seg.mean
  <Rle>         <IRanges> <Rle> | <character> <numeric>
[1] chr11 55139733-55204003 * | md_22169_01 -0.4712

```

```

[2] chr22 17060279-17270615 * | md_22169_01 -0.097
[3] chr22 17281004-17633332 * | md_22169_01 -0.3178
[4] chr22 17670969-18670691 * | md_22169_01 -0.354
[5] chr22 19066315-19792353 * | md_22169_01 -0.3359
  log_RR log_odds prob_MAP prob_222 prob_221
  <numeric> <numeric> <numeric> <numeric> <numeric>
[1] 72.737 -82.214 1 0 0
[2] 0 -91.046 1 1 0
[3] 152.429 Inf 1 0 1
[4] 697.008 Inf 1 0 1
[5] 430.425 Inf 1 0 1
  calls number_probes
  <character> <integer>
[1] 302 25
[2] 222 25
[3] 221 88
[4] 221 278
[5] 221 208

```

-----  
seqinfo: 3 sequences from hg18 genome

The helper functions `denovoHemizygous` and `denovoHomozygous` create commonly used filters.

R> `denovoHemizygous(md_g)`

MDRanges object with 3 ranges and 9 metadata columns:

```

  seqnames      ranges strand |      sample seg.mean
  <Rle>         <IRanges> <Rle> | <character> <numeric>
[1] chr22 17281004-17633332 * | md_22169_01 -0.3178
[2] chr22 17670969-18670691 * | md_22169_01 -0.354
[3] chr22 19066315-19792353 * | md_22169_01 -0.3359
  log_RR log_odds prob_MAP prob_222 prob_221
  <numeric> <numeric> <numeric> <numeric> <numeric>
[1] 152.429 Inf 1 0 1
[2] 697.008 Inf 1 0 1
[3] 430.425 Inf 1 0 1
  calls number_probes
  <character> <integer>
[1] 221 88
[2] 221 278
[3] 221 208

```

-----  
seqinfo: 3 sequences from hg18 genome

R> `denovoHomozygous(md_g)`

MDRanges object with 0 ranges and 9 metadata columns:

```

  seqnames      ranges strand |      sample seg.mean log_RR
  <Rle> <IRanges> <Rle> | <character> <numeric> <numeric>
  log_odds prob_MAP prob_222 prob_221 calls
  <numeric> <numeric> <numeric> <numeric> <character>
  number_probes
  <integer>

```

-----  
seqinfo: 3 sequences from hg18 genome



Equivalently, one could customize the filter parameters to achieve the same result:

```
R> select_cnv <- FilterParamMD(state=c("220", "221", "223"), seqnames="chr22")
R> cnvs <- cnvFilter(md_g, select_cnv)
R> cnvs
```

MDRanges object with 3 ranges and 9 metadata columns:

	seqnames	ranges	strand	sample	seg.mean
	<Rle>	<IRanges>	<Rle>	<character>	<numeric>
[1]	chr22	17281004-17633332	*	md_22169_01	-0.3178
[2]	chr22	17670969-18670691	*	md_22169_01	-0.354
[3]	chr22	19066315-19792353	*	md_22169_01	-0.3359
	log_RR	log_odds	prob_MAP	prob_222	prob_221
	<numeric>	<numeric>	<numeric>	<numeric>	<numeric>
[1]	152.429	Inf	1	0	1
[2]	697.008	Inf	1	0	1
[3]	430.425	Inf	1	0	1
	calls	number_probes			
	<character>	<integer>			
[1]	221	88			
[2]	221	278			
[3]	221	208			

-----  
seqinfo: 3 sequences from hg18 genome

## 5.2 Visualization

This produces a somewhat unsatisfactory result in that there are 2 additional *denovo* hemizygous deletions nearby that likely comprise one *denovo* segment.

```
R> denovoHemizygous(md_g)
```

MDRanges object with 3 ranges and 9 metadata columns:

	seqnames	ranges	strand	sample	seg.mean
	<Rle>	<IRanges>	<Rle>	<character>	<numeric>
[1]	chr22	17281004-17633332	*	md_22169_01	-0.3178
[2]	chr22	17670969-18670691	*	md_22169_01	-0.354
[3]	chr22	19066315-19792353	*	md_22169_01	-0.3359
	log_RR	log_odds	prob_MAP	prob_222	prob_221
	<numeric>	<numeric>	<numeric>	<numeric>	<numeric>
[1]	152.429	Inf	1	0	1
[2]	697.008	Inf	1	0	1
[3]	430.425	Inf	1	0	1
	calls	number_probes			
	<character>	<integer>			
[1]	221	88			
[2]	221	278			
[3]	221	208			

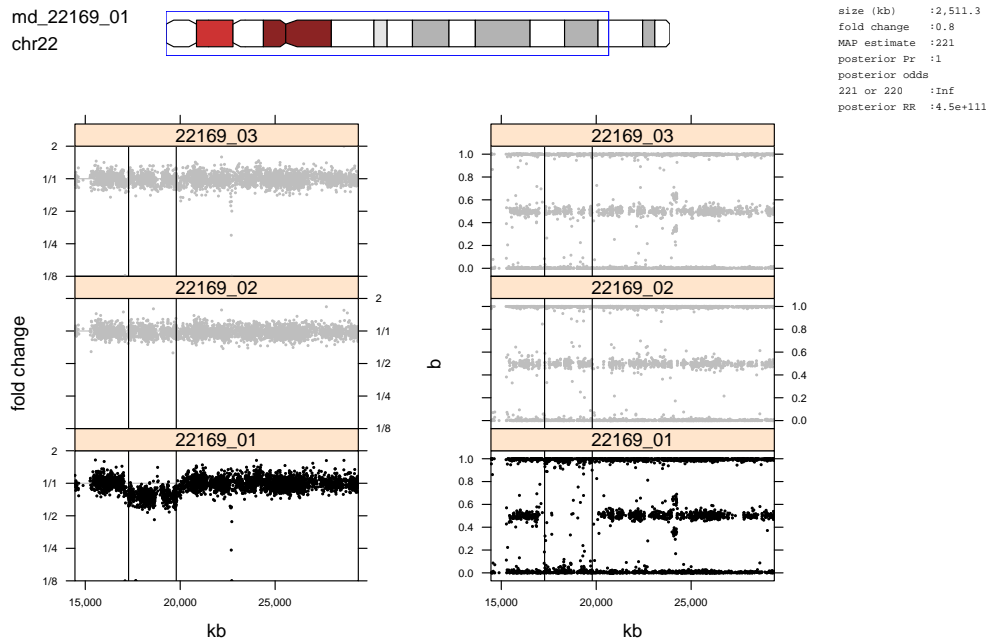
-----  
seqinfo: 3 sequences from hg18 genome

In the following code-chunk, we reduce the *denovo* hemizygous deletions and update the posterior probabilities for the MAP estimates. We use a combination of lattice and grid to visualize the marker-level summaries and copy number inference for the trio.

```

R> library(grid)
R> g2 <- reduce(denovoHemizygous(md_g), min.gapwidth=500e3)
R> post <- MAP2(me, g2, params)
R> g2 <- denovoHemizygous(post)
R> vps <- pedigreeViewports()
R> grid.params <- HmmTrellisParam()
R> p <- plotDenovo(me, g2, grid.params)
R> pedigreeGrid(g=g2, vps=vps, figs=p)

```



## 6 Acknowledgements

Moiz Bootwalla contributed to early versions of this vignette.

## 7 Session information

```
R> toLatex(sessionInfo())
```

- R version 3.6.0 (2019-04-26), x86\_64-pc-linux-gnu
- Locale: LC\_CTYPE=en\_US.UTF-8, LC\_NUMERIC=C, LC\_TIME=en\_US.UTF-8, 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 18.04.2 LTS
- Matrix products: default
- BLAS: /home/biocbuild/bbs-3.9-bioc/R/lib/libRblas.so
- LAPACK: /home/biocbuild/bbs-3.9-bioc/R/lib/libRlapack.so
- Base packages: base, datasets, grDevices, graphics, grid, methods, parallel, stats, stats4, utils

- Other packages: BSgenome 1.52.0, BSgenome.Hsapiens.UCSC.hg18 1.3.1000, Biobase 2.44.0, BiocGenerics 0.30.0, BiocParallel 1.18.0, Biostrings 2.52.0, DelayedArray 0.10.0, GenomeInfoDb 1.20.0, GenomicRanges 1.36.0, IRanges 2.18.0, MinimumDistance 1.28.0, S4Vectors 0.22.0, SummarizedExperiment 1.14.0, VanillaICE 1.46.0, XVector 0.24.0, data.table 1.12.2, matrixStats 0.54.0, oligoClasses 1.46.0, rtracklayer 1.44.0
- Loaded via a namespace (and not attached): BiocManager 1.30.4, DBI 1.0.0, DNACopy 1.58.0, GenomeInfoDbData 1.2.1, GenomicAlignments 1.20.0, Matrix 1.2-17, RCurl 1.95-4.12, Rcpp 1.0.1, RcppEigen 0.3.3.5.0, Rsamtools 2.0.0, VGAM 1.1-1, XML 3.98-1.19, affyio 1.54.0, askpass 1.1, base64 2.0, beanplot 1.2, bit 1.1-14, bitops 1.0-6, codetools 0.2-16, compiler 3.6.0, crlmm 1.42.0, ellipse 0.4.1, ff 2.2-14, foreach 1.4.4, illuminaio 0.26.0, iterators 1.0.10, lattice 0.20-38, limma 3.40.0, mvtnorm 1.0-10, openssl 1.3, preprocessCore 1.46.0, splines 3.6.0, tools 3.6.0, zlibbioc 1.30.0

## References

- [1] Stefano Colella, Christopher Yau, Jennifer M Taylor, Ghazala Mirza, Helen Butler, Penny Clouston, Anne S Bassett, Anneke Seller, Christopher C Holmes, and Jiannis Ragoussis. QuantisNP: an Objective Bayes Hidden-Markov Model to detect and accurately map copy number variation using SNP genotyping data. *Nucleic Acids Res*, 35(6):2013–2025, 2007.
- [2] Philippe Hupe, Nicolas Stransky, Jean-Paul Thiery, Francois Radvanyi, and Emmanuel Barillot. Analysis of array cgh data: from signal ratio to gain and loss of dna regions. *Bioinformatics*, 20(18):3413–3422, Dec 2004.
- [3] Adam B Olshen, E. S. Venkatraman, Robert Lucito, and Michael Wigler. Circular binary segmentation for the analysis of array-based DNA copy number data. *Biostatistics*, 5(4):557–72, Oct 2004.
- [4] Daniel A Peiffer, Jennie M Le, Frank J Steemers, Weihua Chang, Tony Jenniges, Francisco Garcia, Kirt Haden, Jiangzhen Li, Chad A Shaw, John Belmont, Sau Wai Cheung, Richard M Shen, David L Barker, and Kevin L Gunderson. High-resolution genomic profiling of chromosomal aberrations using infinium whole-genome genotyping. *Genome Res*, 16(9):1136–1148, Sep 2006.
- [5] Robert B Scharpf, Giovanni Parmigiani, Jonathan Pevsner, and Ingo Ruczinski. Hidden Markov models for the assessment of chromosomal alterations using high-throughput SNP arrays. *Annals of Applied Statistics*, 2(2):687–713, 2008.
- [6] E. S. Venkatraman and Adam B Olshen. A faster circular binary segmentation algorithm for the analysis of array cgh data. *Bioinformatics*, 23(6):657–663, Mar 2007.
- [7] Kai Wang, Zhen Chen, Mahlet G Tadesse, Joseph Glessner, Struan F A Grant, Hakon Hakonarson, Maja Bucan, and Mingyao Li. Modeling genetic inheritance of copy number variations. *Nucleic Acids Res*, 36(21):e138, Dec 2008.