

RiboProfiling - Analysis and quality assessment of Ribosome-profiling data

Alexandra Popa

April 26, 2022

Contents

1	Introduction	2
2	Data Input	2
3	Quick start.	3
4	Histogram of read match length	5
5	Read start coverage plot around the TSS	6
6	Count reads on features	7
7	Count reads on codon motifs	9

1 Introduction

Ribosome profiling, the recently developed high throughput sequencing technique, enabled the mapping of translated regions genome-wide. This technique takes advantage of the fact that ribosomes actively engaged in translation can protect their associated mRNA fragments against RNase digestion. The sequencing of these protected fragments can reveal potentially translated sequences. *RiboProfiling* is a Bioconductor package that provides multiple types of ribo-seq data analysis, starting from BAM files alone:

- Quality assessment of read match size distribution.
- Read coverage around the TSS for the definition of an offset (shift value) of ribosome position.
- Calibrate reads by applying an offset to the read start positions along the transcript.
- Table of read start counts (shifted if specified) on the specified region (CDS, 5pUTR, 3pUTR).
- A graphical quality function for ribo-seq data, based on the previously obtained tables.
- The quantification of the frequency and coverage matrices for motifs of codons (1, 2 or 3 codons) in ORFs.
- Principal component analysis of codon motif coverage in ribosome-profiling.

The package also provides the `riboSeqFromBAM` function, that assembles in a global framework some of the above analyses, allowing a quick overview of the data.

2 Data Input

As input data, the *RiboProfiling* package can start with a path to a ribo-seq (or other type of) BAM to be analyzed. A list of multiple BAM files can also be provided. BAM files from either bowtie/tophat, STAR, and Lifescope (Solid), single- and paired-end have been validated. Many intermediate functions of the *RiboProfiling* package can be called for additional modifications on the data objects.

Example data files provided with the *RiboProfiling* package are based on chromosome 1 reads from human primary BJ fibroblasts data (PMID: 23594524, SRA: <http://www.ebi.ac.uk/ena/data/view/SRP020544>):

- BAM file Ctrl "<http://genomique.info/data/public/RiboProfiling/ctrl.bam>".
- BAM file Nutlin2h "<http://genomique.info/data/public/RiboProfiling/nutlin2h.bam>".

The data provided with the *RiboProfiling* package and accessible through the `data` function consists of:

- **ctrlGAlignments** - a GAlignments object corresponding to the "<http://genomique.info/data/public/RiboProfiling/ctrl.bam>" BAM.
- **codonIndexCovCtrl** - a list containing the number of reads for each codon in each ORF.
- **codonDataCtrl** - a list of 2 data.frames containing the frequency of codons for each ORF and, respectively, the read coverage for the same codons and ORFs.

3 Quick start

The fastest way to analyze a list of BAM files with the *RiboProfiling* package is to call the `riboSeqFromBAM` function. The only input needed are the paths to the BAM files and the genome version on which the mapping was done. Here is an example on how to use this wrapper function on 2 BAM files.

```
library(RiboProfiling)
listInputBam <- c(
  BamFile("http://genomique.info/data/public/RiboProfiling/ctrl.bam"),
  BamFile("http://genomique.info/data/public/RiboProfiling/nutlin2h.bam")
)
covData <- riboSeqFromBAM(listInputBam, genomeName="hg19")
```

The following analyses and results are returned:

- a histogram of read match lengths: figures 1 and 2.
- a read start coverage plot around the TSS: figures 1 and 2.
- an automatic estimation and application of the ribosome offset position (if no offset value is provided (`listShiftValue` parameter missing))
- a data.frame containing CDSs annotation and counts on the 5pUTR, CDS, and 3pUTR once the offset has been applied
- pairs and boxplots of the read counts: figures 3 and 4.
- a list of per ORF per codon coverage

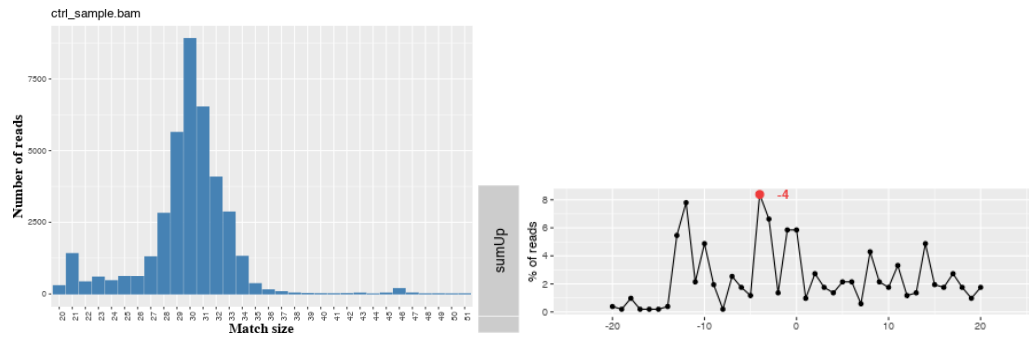


Figure 1: Histogram of read match size and read start offset to ribosome P-site for ctrl.bam dataset.

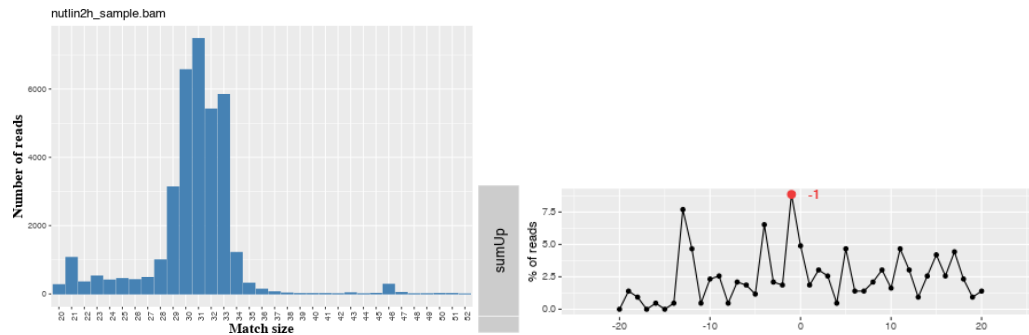


Figure 2: Histogram of read match size and read start offset to ribosome P-site for nutlin2h.bam dataset.

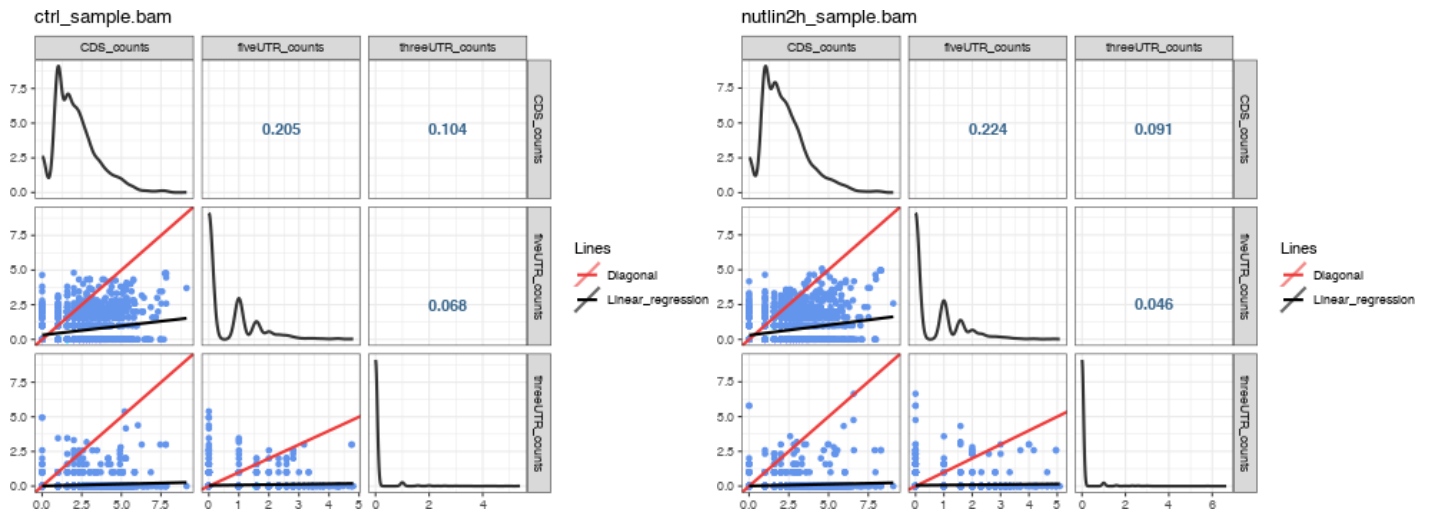


Figure 3: Pairs of shifted read counts on 5pUTR, CDS, and 3pUTR for ctrl.bam and nutlin2h.bam dataset.

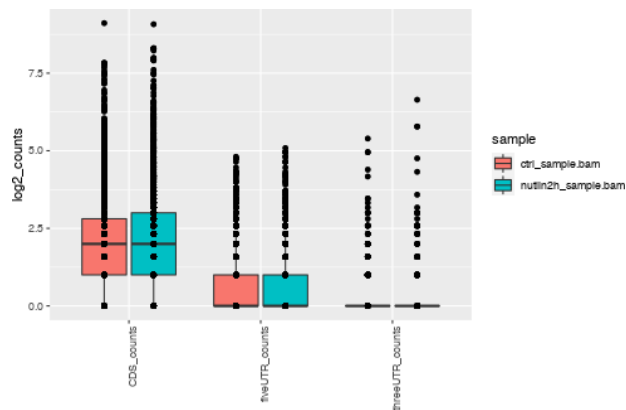


Figure 4: Boxplot of shifted read counts on 5pUTR, CDS, and 3pUTR for ctrl.bam and nutlin2h.bam dataset.

All these results will be explained in detail in the following sections.

This general function only works if your species model has an ensembl table in the UCSC database or if you provide the corresponding txdb object containing the annotations.

4 Histogram of read match length

This histogram represents both a quality control of the sequencing and an important tool to define the match sizes of reads corresponding to ribosome footprints (around 30bp). The *histMatchLength* function in the *RiboProfiling* package produces this histogram starting from a *GAlignment* object (the loaded records of a BAM) such as the *ctrlAlignments* data example (figure 5).

One can create a *GAlignments* from a BAM using the `readGAlignments` function from the *GenomicAlignments* package:

```
aln <- readGAlignments(
  BamFile("http://genomique.info/data/public/RiboProfiling/ctrl.bam")
)
```

Or based on an already existing *GAlignments* object:

```
data(ctrlAlignments)
aln <- ctrlAlignments
matchLenDistr <- histMatchLength(aln, 0)
matchLenDistr[[2]]
```

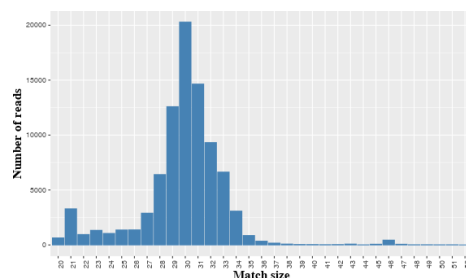


Figure 5: Histogram of read match length for example data: ctrlAlignments.

5 Read start coverage plot around the TSS

The initiation of protein synthesis starts with the first codon in the P-site. In ribosome profiling experiments, the location of the P-site codon must be inferred in order to recalibrate the read start positions relative to the transcript. An offset is usually observed in Ribo-seq experiments between the starting of the reads and the AUG codons of protein coding sequences. Three functions allow the estimation of the read start position relative to the P-site codon:

- `aroundPromoter`: returns the genomic positions flanking the transcript start site (TSS) for the x% (3% default value) best expressed CDSs.
- `readStartCov`: returns the read start coverage around the TSS on the predefined CDSs.
- `plotSummarizedCov`: plots the summarized coverage in a specified range (e.g. around TSS) for the specified match sizes.

```
#transform the GAlignments object into a GRanges object
#(faster processing of the object)
alnGRanges <- readsToStartOrEnd(aln, what="start")
#txdb object with annotations
txdb <- TxDb.Hsapiens.UCSC.hg19.knownGene
oneBinRanges <- aroundPromoter(txdb, alnGRanges, percBestExpressed=0.001)
#the coverage in the TSS flanking region for the reads with match sizes 29:31
listPromoterCov <-
  readStartCov(
    alnGRanges,
    oneBinRanges,
    matchSize=c(29:31),
    fixedInterval=c(-20, 20),
    renameChr="aroundTSS",
    charPerc="perc"
  )
plotSummarizedCov(listPromoterCov)
```

```
## Warning in !vapply(ggl, fixed, logical(1L)) & !vapply(PlotList, is, "Ideogram",
: longer object length is not a multiple of shorter object length
```

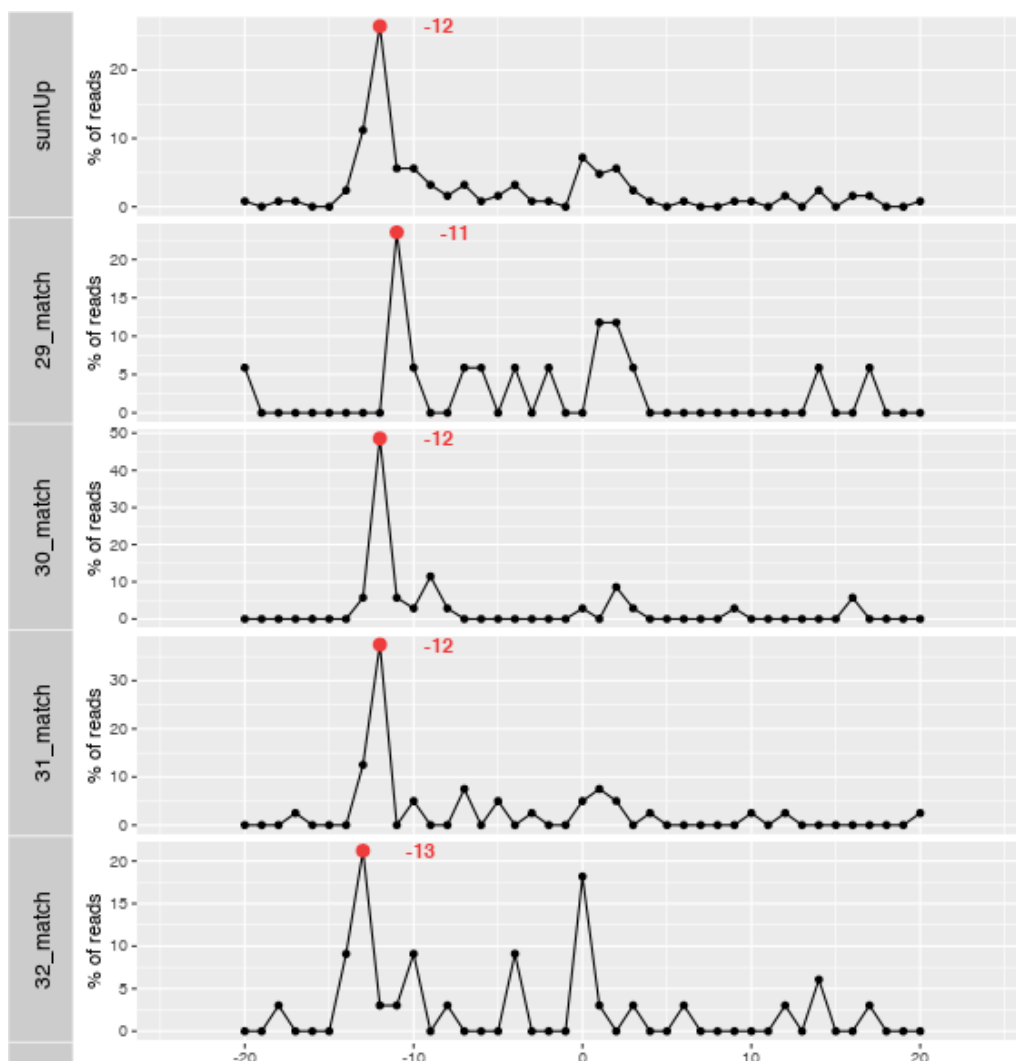


Figure 6: Offset around TSS for the best expressed CDSs on a subset of match sizes.

The first time you analyse a ribo-seq BAM file, it is advisable to make an offset graph for all the previously selected read match sizes, as the offset might be different depending on the read match size .

6 Count reads on features

The purpose of Ribosome Profiling experiments is mainly to identify RNase resistant regions, which might indicate that these sequences are likely to be translated. The `countShiftReads` function quantifies the read start coverage on different sequence features. This function integrates the specificity of ribo-seq data of having shifted reads starts from the P-site codon. The `countShiftReads` function recalibrates the read start positions along the transcript with the specified offset (parameter `shiftValue`, default 0) and returns the coverage on the 5pUTR, CDS, 3pUTR, and a matrix of codon coverage for each ORF.

The counts can be used globally to see what is RNAse protected and what is not, but also for differential analysis between conditions. The `countsPlot` function provides some quality graphs for the ribo-seq data: pairs between counts on features and boxplots of counts between conditions.

As an example, we compute the read coverage on the **ctrlGAlignments** example data:

```
#keep only the match read sizes 30-33
alnGRanges <- alnGRanges[which(!is.na(match(alnGRanges$score,30:33)))]
#get all CDSs by transcript
cds <- cdsBy(txdb, by="tx", use.names=TRUE)
#get all exons by transcript
exonGRanges <- exonsBy(txdb, by="tx", use.names=TRUE)
#get the per transcript relative position of start and end codons
cdsPosTransc <- orfRelativePos(cds, exonGRanges)
#compute the counts on the different features
#after applying the specified shift value on the read start along the transcript
countsDataCtrl1 <-
  countShiftReads(
    exonGRanges=exonGRanges[names(cdsPosTransc)],
    cdsPosTransc=cdsPosTransc,
    alnGRanges=alnGRanges,
    shiftValue=-14
  )
head(countsDataCtrl1[[1]])
listCountsPlots <- countsPlot(
  list(countsDataCtrl1[[1]]),
  grep("_counts$", colnames(countsDataCtrl1[[1]])),
  1
)
listCountsPlots
```

```
## Warning in countShiftReads(exonGRanges[names(cdsPosTransc)], cdsPosTransc, :
Param motifSize should be an integer! Accepted values 3, 6 or 9. Default value
is 3.

## Warning: Ignoring unknown aesthetics: fill
```

##	gene	chr	strand	transc_genomic_start	transc_genomic_end	transc_length	orf_start	orf_end	orf_length	CDS_counts	fiveUTR_counts	threeUTR_counts
##	uc001abw.1	uc001abw.1	chr1	+	861121	879961						
##	uc031pjl.1	uc031pjl.1	chr1	+	861302	879533						
##	uc031pjm.1	uc031pjm.1	chr1	+	861302	879533						
##	uc001abw.1					2554	81	2126	2046	2	0	
##	uc031pjl.1					2120	21	2120	2100	2	0	
##	uc031pjm.1					2084	21	2084	2064	2	0	
##	uc001abw.1											0
##	uc031pjl.1											0
##	uc031pjm.1											0

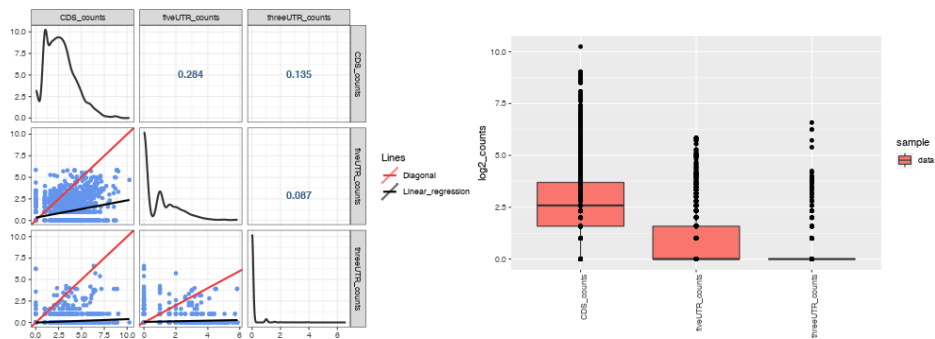


Figure 7: Pairs and boxplots of read coverage on genomic features.

7 Count reads on codon motifs

The previous `countShiftReads` function produces a second element in the list: the counts of reads per motifs of 1, 2, or 3 codons. For each ORF, for each motif in the ORF, the Ribo-seq coverage of the motif is reported as follows:

- for motifs of 3 nucleotides (1 codon) - the sum of read starts on the 3 nucleotides in the codon is reported.
- consecutive motifs of 6 nucleotides (2 consecutive codons) overlap on 3 nucleotides. The Ribo-seq coverage is reported as the coverage on the 1st codon in the motif considered as being in the P-site.
- consecutive motifs of 9 nucleotides (3 consecutive codons) overlap on 6 nucleotides. The Ribo-seq coverage is reported as the coverage on the 2nd codon in the motif considered as being in the P-site.

Codon motifs in each ORF are described as the index of their position in the ORF.

Here is an example:

```
data(codonIndexCovCtrl)
head(codonIndexCovCtrl[[1]], n=3)

##   codonID nbrReads
## 1       1        0
## 2       2        0
## 3       3        0
```

The `codonInfo` function associates the read coverage on codons with their corresponding codon type for each ORF. It returns a list of 2 matrices: one containing the frequency of each codon motif type in each ORF and the second containing the coverage of each codon motif type in each ORF.

```
listReadsCodon <- countsDataCtrl1[[2]]
#get the names of the expressed ORFs grouped by transcript
cds <- cdsBy(txdb, use.names=TRUE)
orfCoord <- cds[names(cds) %in% names(listReadsCodon)]
#chromosome names should correspond between the BAM,
#the annotations, and the genome
genomeSeq <- BSgenome.Hsapiens.UCSC.hg19
#codon frequency, coverage, and annotation
```

```
codonData <- codonInfo(listReadsCodon, genomeSeq, orfCoord)
```

```
## Warning in codonInfo(listReadsCodon, genomeSeq, orfCoord): Param motifSize should
be an integer! Accepted values 3, 6 or 9. Default value is 3.
```

```
## Warning in result_fetch(res@ptr, n = n): SQL statements must be issued with
dbExecute() or dbSendStatement() instead of dbGetQuery() or dbSendQuery().
```

```
data("codonDataCtrl")
head(codonDataCtrl[[1]], n=3)
```

	AAA	AAC	AAG	AAT	ACA	ACC	ACG	ACT	AGA	AGC	AGG	AGT	ATA	ATC	ATG	ATT	CAA
## uc010nxq.1	1	2	1	0	1	2	2	2	2	3	6	6	0	3	3	0	0
## uc001abv.1	0	4	6	0	1	3	0	1	3	8	1	1	2	5	3	1	0
## uc001abw.1	1	8	18	2	4	13	7	4	3	21	11	2	2	8	8	1	4

	CAC	CAG	CAT	CCA	CCC	CCG	CCT	CGA	CGC	CGG	CGT	CTA	CTC	CTG	CTT	GAA	GAC
## uc010nxq.1	7	5	2	3	9	0	4	0	0	0	0	1	3	8	2	1	3
## uc001abv.1	5	5	2	1	5	4	2	4	3	5	1	0	2	6	1	1	6
## uc001abw.1	16	27	6	14	45	13	12	8	10	21	3	1	13	56	5	5	22

	GAG	GAT	GCA	GCC	GCG	GCT	GGA	GGC	GGG	GGT	GTA	GTC	GTG	GTT	TAA	TAC	TAG
## uc010nxq.1	6	1	1	2	1	2	3	2	1	3	0	1	5	2	0	1	1
## uc001abv.1	5	2	1	5	1	0	0	6	4	2	0	3	3	1	0	0	0
## uc001abw.1	41	7	7	38	11	11	7	30	26	7	1	9	12	3	0	6	0

	TAT	TCA	TCC	TCG	TCT	TGA	TGC	TGG	TGT	TTA	TTC	TTG	TTT
## uc010nxq.1	0	0	3	0	3	0	3	0	3	0	3	2	3
## uc001abv.1	0	0	5	0	2	0	4	1	2	0	3	1	0
## uc001abw.1	2	4	17	6	4	1	8	5	4	0	14	5	2

The codon coverage can be used to study the ribosome translation dynamics. One can test if codon motifs accumulating ribo-seq reads represent a stalling in the ribosome progression. In the *RiboProfiling* package we have implemented the `codonPCA` function that performs a PCA analysis on a matrix of read density on codons (figures 8).

```
codonUsage <- codonData[[1]]
codonCovMatrix <- codonData[[2]]
#keep only genes with a minimum number of reads
nbrReadsGene <- apply(codonCovMatrix, 1, sum)
ixExpGenes <- which(nbrReadsGene >= 50)
codonCovMatrix <- codonCovMatrix[ixExpGenes, ]
#get the PCA on the codon coverage
codonCovMatrixTransp <- t(codonCovMatrix)
rownames(codonCovMatrixTransp) <- colnames(codonCovMatrix)
colnames(codonCovMatrixTransp) <- rownames(codonCovMatrix)
listPCACodonCoverage <- codonPCA(codonCovMatrixTransp, "codonCoverage")
listPCACodonCoverage[[2]]

## $`PC_1-2`
##
## $`PC_1-3`
##
## $`PC_1-4`
##
```

```
## $`PC_2-3`
##
## $`PC_2-4`
```

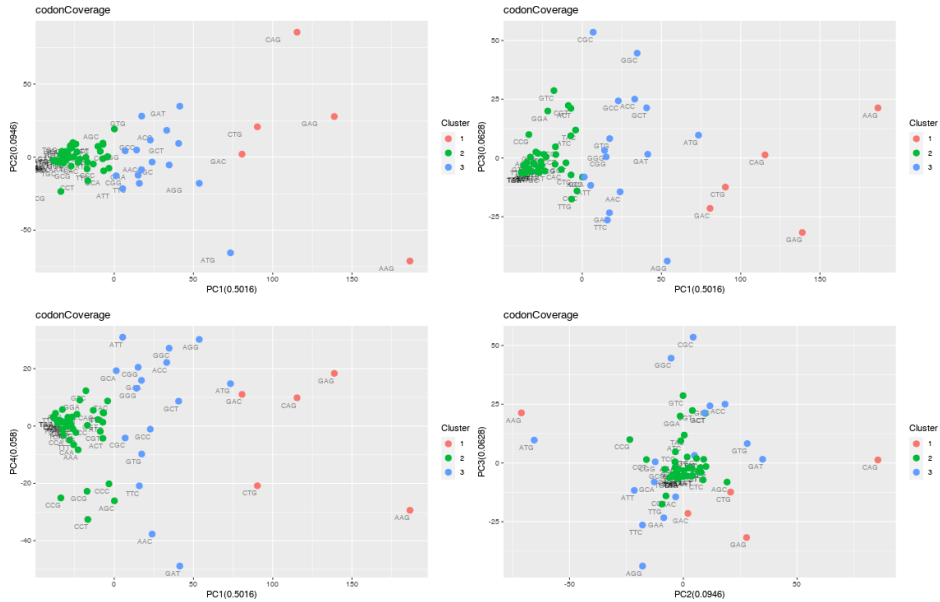


Figure 8: PCA on codon coverage for the ctrlAlignments example data.

```
sessionInfo()

## R version 4.2.0 RC (2022-04-19 r82224)
## Platform: x86_64-apple-darwin17.0 (64-bit)
## Running under: macOS Mojave 10.14.6
##
## Matrix products: default
## BLAS: /Library/Frameworks/R.framework/Versions/4.2/Resources/lib/libRblas.0.dylib
## LAPACK: /Library/Frameworks/R.framework/Versions/4.2/Resources/lib/libRlapack.dylib
##
## locale:
## [1] C/en_US.UTF-8/en_US.UTF-8/C/en_US.UTF-8/en_US.UTF-8
##
## attached base packages:
## [1] stats4 stats graphics grDevices utils datasets methods
## [8] base
##
## other attached packages:
## [1] RSQLite_2.2.12 RiboProfiling_1.26.0 Biostrings_2.64.0
## [4] GenomeInfoDb_1.32.0 XVector_0.36.0 IRanges_2.30.0
## [7] S4Vectors_0.34.0 BiocGenerics_0.42.0 knitr_1.38
##
## loaded via a namespace (and not attached):
## [1] colorspace_2.0-3
## [2] rjson_0.2.21
## [3] ellipsis_0.3.2
## [4] biovizBase_1.44.0
```

```
## [5] htmlTable_2.4.0
## [6] GenomicRanges_1.48.0
## [7] base64enc_0.1-3
## [8] dichromat_2.0-0
## [9] rstudioapi_0.13
## [10] farver_2.1.0
## [11] bit64_4.0.5
## [12] AnnotationDbi_1.58.0
## [13] fansi_1.0.3
## [14] sqldf_0.4-11
## [15] xml2_1.3.3
## [16] splines_4.2.0
## [17] ggbio_1.44.0
## [18] cachem_1.0.6
## [19] Formula_1.2-4
## [20] Rsamtools_2.12.0
## [21] cluster_2.1.3
## [22] dbplyr_2.1.1
## [23] png_0.1-7
## [24] graph_1.74.0
## [25] BiocManager_1.30.17
## [26] compiler_4.2.0
## [27] httr_1.4.2
## [28] backports_1.4.1
## [29] lazyeval_0.2.2
## [30] assertthat_0.2.1
## [31] Matrix_1.4-1
## [32] fastmap_1.1.0
## [33] TxDb.Hsapiens.UCSC.hg19.knownGene_3.2.2
## [34] cli_3.3.0
## [35] htmltools_0.5.2
## [36] prettyunits_1.1.1
## [37] tools_4.2.0
## [38] gtable_0.3.0
## [39] glue_1.6.2
## [40] GenomeInfoDbData_1.2.8
## [41] reshape2_1.4.4
## [42] dplyr_1.0.8
## [43] rappdirs_0.3.3
## [44] Rcpp_1.0.8.3
## [45] Biobase_2.56.0
## [46] vctrs_0.4.1
## [47] nlme_3.1-157
## [48] rtracklayer_1.56.0
## [49] xfun_0.30
## [50] stringr_1.4.0
## [51] proto_1.0.0
## [52] lifecycle_1.0.1
## [53] ensemblDb_2.20.0
## [54] restfulr_0.0.13
## [55] XML_3.99-0.9
```

```
## [56] zlibbioc_1.42.0
## [57] scales_1.2.0
## [58] BSgenome_1.64.0
## [59] BiocStyle_2.24.0
## [60] VariantAnnotation_1.42.0
## [61] ProtGenerics_1.28.0
## [62] hms_1.1.1
## [63] MatrixGenerics_1.8.0
## [64] RBGL_1.72.0
## [65] parallel_4.2.0
## [66] SummarizedExperiment_1.26.0
## [67] AnnotationFilter_1.20.0
## [68] RColorBrewer_1.1-3
## [69] yaml_2.3.5
## [70] curl_4.3.2
## [71] memoise_2.0.1
## [72] gridExtra_2.3
## [73] ggplot2_3.3.5
## [74] biomaRt_2.52.0
## [75] rpart_4.1.16
## [76] reshape_0.8.9
## [77] latticeExtra_0.6-29
## [78] stringi_1.7.6
## [79] highr_0.9
## [80] BiocIO_1.6.0
## [81] checkmate_2.1.0
## [82] GenomicFeatures_1.48.0
## [83] filelock_1.0.2
## [84] BiocParallel_1.30.0
## [85] chron_2.3-56
## [86] rlang_1.0.2
## [87] pkgconfig_2.0.3
## [88] matrixStats_0.62.0
## [89] bitops_1.0-7
## [90] evaluate_0.15
## [91] lattice_0.20-45
## [92] purrr_0.3.4
## [93] labeling_0.4.2
## [94] GenomicAlignments_1.32.0
## [95] htmlwidgets_1.5.4
## [96] bit_4.0.4
## [97] tidyselect_1.1.2
## [98] BSgenome.Hsapiens.UCSC.hg19_1.4.3
## [99] GGally_2.1.2
## [100] plyr_1.8.7
## [101] magrittr_2.0.3
## [102] R6_2.5.1
## [103] magick_2.7.3
## [104] generics_0.1.2
## [105] Hmisc_4.7-0
## [106] DelayedArray_0.22.0
```

```
## [107] DBI_1.1.2
## [108] mgcv_1.8-40
## [109] gsubfn_0.7
## [110] pillar_1.7.0
## [111] foreign_0.8-82
## [112] survival_3.3-1
## [113] KEGGREST_1.36.0
## [114] RCurl_1.98-1.6
## [115] nnet_7.3-17
## [116] tibble_3.1.6
## [117] crayon_1.5.1
## [118] utf8_1.2.2
## [119] OrganismDbi_1.38.0
## [120] BiocFileCache_2.4.0
## [121] rmarkdown_2.14
## [122] jpeg_0.1-9
## [123] progress_1.2.2
## [124] grid_4.2.0
## [125] data.table_1.14.2
## [126] blob_1.2.3
## [127] digest_0.6.29
## [128] munsell_0.5.0
## [129] tcltk_4.2.0
```