

nearBynding Vignette

Veronica Busa

2025-10-30

Contents

Introduction	2
Installation	2
External Software Dependencies	2
bedtools	2
CapR	2
StereoGene	2
1. Concatenate the Transcriptome	3
2. Fold RNA via CapR	3
3. Map to Transcriptome	4
4. Calculate Cross-correlation via StereoGene	4
5. Visualize Results	5
6. Calculate Distance	6

Introduction

nearBynding is a package designed to discern annotated RNA structures proximal to protein binding. nearBynding allows users to annotate RNA structure contexts via CapR or input their own annotations in BED/bedGraph format. It accomodates protein binding information from CLIP-seq experiments as either aligned CLIP-seq reads or peak-called intervals. This vignette will walk you through:

- * The external software necessary to support this pipeline
- * Creating a concatenated transcriptome
- * Extracting and folding RNA from the transcriptome via CapR
- * Mapping protein-binding and RNA structure information onto a transcriptome
- * Running StereoGene to identify RNA structure proximal to protein binding
- * Visualizing binding results
- * Determining the distance between binding contexts

Before running any of these examples, it is highly recommended that the user establishes a new empty directory and uses `setwd()` to make certain that all outputs are deposited there. Some of the functions below create multiple output files.

Installation

```
if(!requireNamespace("BiocManager", quietly = TRUE))
  install.packages("BiocManager")
BiocManager::install("nearBynding")
```

External Software Dependencies

Add all dependency directories to your PATH after installation.

bedtools

bedtools is available for installation [here](#).

Installation instructions will vary by operating system.

CapR

Download the zip file from the [github repository](#), unzip the file, and move it to a directory where you want to permanently store the function.

In the command line, access the folder where the unzipped file is stored.

```
cd CapR-master
make
./CapR
```

If installation is successful, the final line will output

```
Error: The number of argument is invalid.
```

StereoGene

Download the zip file from the [github repository](#), unzip the file, and move it to a directory where you want to permanently store the function.

In the command line, access the folder where the unzipped file is stored.

```
cd stereogene-master
cd src
make
./stereogene -h
```

If installation is successful, the final line will output a menu of argument options.

1. Concatenate the Transcriptome

Although nearBynding is designed to support whole-genome analyses, we will exclusively be evaluating protein-coding genes of chromosomes 4 and 5 through this vignette.

First, a list of transcripts must be identified for analysis. A recommended criterium for selection is that the transcripts be expressed in the cell type used for CLIP-seq experiments. For this vignette, 50 random transcripts have been selected, and the 3'UTR structure of each transcript will be used for analysis, though any region of a transcript such as 5'UTR or CDS could be assessed instead.

This step creates a chain file that will be used to map the selected regions of transcripts end-to-end, excluding the intergenic regions and undesired transcripts that comprise the majority of the genome.

```
# load transcript list
load(system.file("extdata/transcript_list.Rda", package="nearBynding"))
# get GTF file
gtf<-system.file("extdata/Homo_sapiens.GRCh38.chr4&5.gtf",
                 package="nearBynding")

GenomeMappingToChainFile(genome_gtf = gtf,
                        out_chain_name = "test.chain",
                        RNA_fragment = "three_prime_utr",
                        transcript_list = transcript_list,
                        alignment = "hg38")
```

A file containing the sizes of each concatenated chromosome in the chain file will be required for downstream analysis.

```
getChainChrSize(chain = "test.chain",
               out_chr = "chr4and5_3UTR.size")
```

2. Fold RNA via CapR

In order to fold the 3'UTRs, the sequences must first be extracted. This is achieved with the following code:

```
ExtractTranscriptomeSequence(transcript_list = transcript_list,
                             ref_genome = "Homo_sapiens.GRCh38.dna.primary_assembly.fa",
                             genome_gtf = gtf,
                             RNA_fragment = "three_prime_utr",
                             exome_prefix = "chr4and5_3UTR")
```

The reference genome can be found through [Ensembl](#), but for users who do not want to download that 3.2GB file for the sake of this vignette, the FASTA output of the above code is available via:

```
chr4and5_3UTR.fa <- system.file("extdata/chr4and5_3UTR.fa",
                               package="nearBynding")
```

These sequences can then be submitted to CapR for folding.

```
runCapR(in_file = chr4and5_3UTR.fa)
```

Warning: This step can take hours or even days depending on how many transcripts are submitted, how long the RNA fragments are, and the maximum distance between base-paired nucleotides submitted to the CapR algorithm.

3. Map to Transcriptome

The nearBynding pipeline can accomodate either a BAM file of aligned CLIP-seq reads or a BED file of peak intervals. BAM files must be sorted and converted to a bedGraph file first, whereas BED files can be read-in directly.

```
bam <- system.file("extdata/chr4and5.bam", package="nearBynding")
sorted_bam<-sortBam(bam, "chr4and5_sorted")

CleanBAMtoBG(in_bam = sorted_bam)
```

BAM file input

Map Protein Binding and RNA Structure to the Transcriptome BED or bedGraph files can then be mapped onto the concatenated transcriptome using the chain file created by `GenomeMappingToChainFile()`. This way, only the protein binding from transcriptomic regions of interest will be considered in the final protein binding analysis.

```
liftOverToExomicBG(input = "chr4and5_sorted.bedGraph",
                   chain = "test.chain",
                   chrom_size = "chr4and5_3UTR.size",
                   output_bg = "chr4and5_liftOver.bedGraph")
```

For BED file inputs, use the additional argument `format = "BED"`.

The RNA structure information from the CapR output also needs to be mapped onto the concatenated transcriptome. There are six different binding contexts calculated by CapR – *stem*, *hairpin*, *multibranch*, *exterior*, *internal*, and *bulge*. `processCapRout()` parses the CapR output, converts it into six separate bedGraph files, and then performs `liftOverToExomic()` for all the files.

For this sake of this vignette, the CapR outfile is available:

```
processCapRout(CapR_outfile = system.file("extdata/chr4and5_3UTR.out",
                                           package="nearBynding"),
               chain = "test.chain",
               output_prefix = "chr4and5_3UTR",
               chrom_size = "chr4and5_3UTR.size",
               genome_gtf = gtf,
               RNA_fragment = "three_prime_utr")
```

It is possible for users to input their own RNA structure information rather than using CapR. The information should be in BED file format and can be input into `liftOverToExomicBG()` to map the RNA structure data to the same transcriptome as the protein binding data.

4. Calculate Cross-correlation via StereoGene

This is the process that directly answers the question, “What does RNA structure look like around where the protein is binding?” StereoGene is used to calculate the cross-correlation between RNA structure and protein binding in order to visualize the RNA structure landscape surrounding protein binding.

If CapR is used to determine RNA structure, `runStereoGeneOnCapR()` initiates StereoGene for a given protein against all CapR-generated RNA structure contexts.

For the sake of this vignette, use `outfiles()` to pull the StereoGene output files to your local directory if you do not want to run StereoGene.

```
runStereoGeneOnCapR(protein_file = "chr4and5_liftOver.bedGraph",
                    chrom_size = "chr4and5_3UTR.size",
                    name_config = "chr4and5_3UTR.cfg",
                    input_prefix = "chr4and5_3UTR")
```

If external RNA structure data is being tested, `runStereoGene()` initiates analysis for a given protein and a single RNA structure context.

Note: The input track file order matters! The correct order is:

- 1) RNA structure
- 2) protein binding

Otherwise, data visualization will be inverted and all downstream analysis will be backwards.

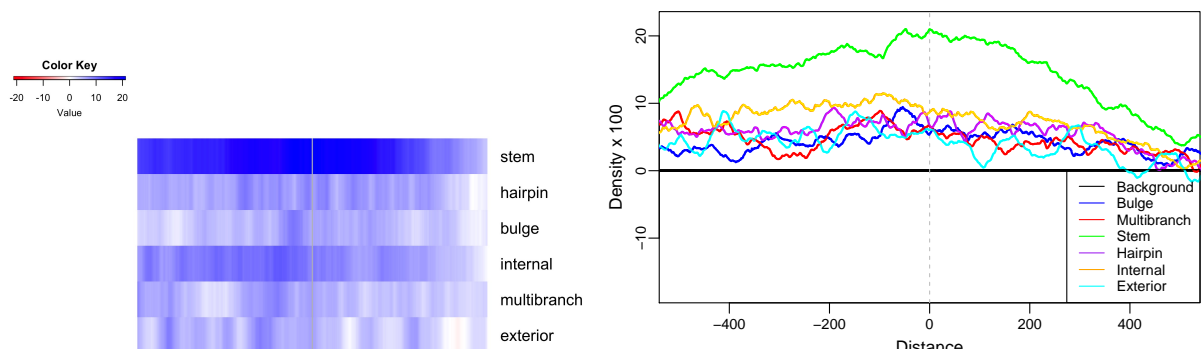
```
runStereoGene(track_files = c("chr4and5_3UTR_stem_liftOver.bedGraph",
                              "chr4and5_liftOver.bedGraph"),
              name_config = "chr4and5_3UTR.cfg")
```

5. Visualize Results

The cross-correlation output of StereoGene can be visualized as either a heatmap or a line plot. Examples of both are below.

For CapR-derived RNA structure, all six contexts can be viewed simultaneously.

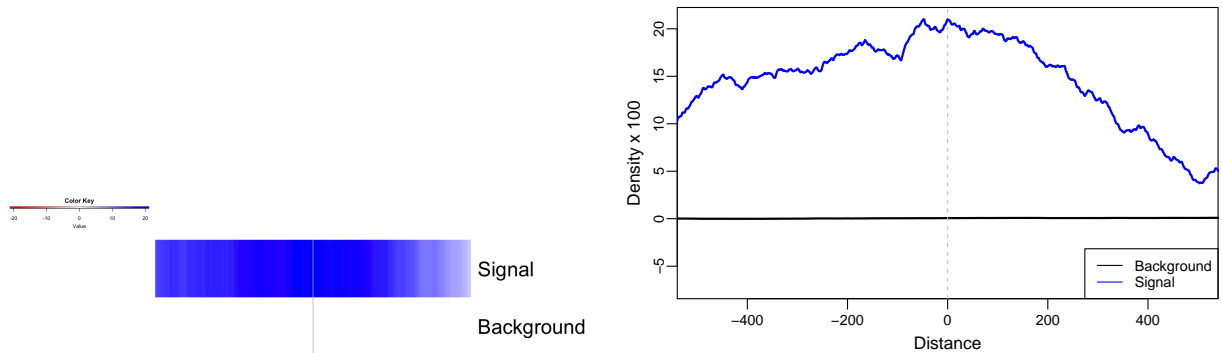
```
visualizeCapRStereoGene(CapR_prefix = "chr4and5_3UTR",
                      protein_file = "chr4and5_liftOver",
                      heatmap = TRUE,
                      out_file = "all_contexts_heatmap",
                      x_lim = c(-500, 500))
visualizeCapRStereoGene(CapR_prefix = "chr4and5_3UTR",
                      protein_file = "chr4and5_liftOver",
                      x_lim = c(-500, 500),
                      out_file = "all_contexts_line",
                      y_lim = c(-18, 22))
```



Warning: This step may take up to an hour for a full transcriptome.

Alternatively, a single context can be viewed at a time.

```
visualizeStereogene(context_file = "chr4and5_3UTR_stem_liftOver",
  protein_file = "chr4and5_liftOver",
  out_file = "stem_line",
  x_lim = c(-500, 500))
visualizeStereogene(context_file = "chr4and5_3UTR_stem_liftOver",
  protein_file = "chr4and5_liftOver",
  heatmap = TRUE,
  out_file = "stem_heatmap",
  x_lim = c(-500, 500))
```



Although this specific, limited example does not provide a particularly visually stimulating image, larger data sets (which provide many more StereoGene windows) result in narrower peaks and less noise.

6. Calculate Distance

In order to determine the similarity of two binding contexts, we can calculate the Wasserstein distance between curves. A small value suggests two binding contexts are very similar, whereas larger values suggest substantial differences.

For example, calculate the distance between the stem and hairpin contexts visualized above.

```
bindingContextDistance(RNA_context = "chr4and5_3UTR_stem_liftOver",
  protein_file = "chr4and5_liftOver",
  RNA_context_2 = "chr4and5_3UTR_hairpin_liftOver")
#> [1] 11.75237
```

Now compare it to the distance between internal and hairpin contexts.

```
bindingContextDistance(RNA_context = "chr4and5_3UTR_internal_liftOver",
  protein_file = "chr4and5_liftOver",
  RNA_context_2 = "chr4and5_3UTR_hairpin_liftOver")
#> [1] 1.788653
```

We can see that the stem context is less similar to the hairpin context than the internal context, and this is reflected in the calculated distances.

Questions? Comments? Please email Veronica Busa at vbusa1@jhmi.edu

```
sessionInfo()
#> R version 4.5.1 Patched (2025-08-23 r88802)
#> Platform: x86_64-pc-linux-gnu
#> Running under: Ubuntu 24.04.3 LTS
#>
```

```

#> Matrix products: default
#> BLAS: /home/biocbuild/bbs-3.22-bioc/R/lib/libRblas.so
#> LAPACK: /usr/lib/x86_64-linux-gnu/lapack/liblapack.so.3.12.0 LAPACK version 3.12.0
#>
#> locale:
#> [1] LC_CTYPE=en_US.UTF-8 LC_NUMERIC=C
#> [3] LC_TIME=en_GB LC_COLLATE=C
#> [5] LC_MONETARY=en_US.UTF-8 LC_MESSAGES=en_US.UTF-8
#> [7] LC_PAPER=en_US.UTF-8 LC_NAME=C
#> [9] LC_ADDRESS=C LC_TELEPHONE=C
#> [11] LC_MEASUREMENT=en_US.UTF-8 LC_IDENTIFICATION=C
#>
#> time zone: America/New_York
#> tzcode source: system (glibc)
#>
#> attached base packages:
#> [1] stats4 stats graphics grDevices utils datasets methods
#> [8] base
#>
#> other attached packages:
#> [1] Rsamtools_2.26.0 Biostrings_2.78.0 XVector_0.50.0
#> [4] GenomicRanges_1.62.0 IRanges_2.44.0 S4Vectors_0.48.0
#> [7] Seqinfo_1.0.0 BiocGenerics_0.56.0 generics_0.1.4
#> [10] nearBynding_1.20.0
#>
#> loaded via a namespace (and not attached):
#> [1] transport_0.15-4
#> [2] KEGGREST_1.50.0
#> [3] SummarizedExperiment_1.40.0
#> [4] gtable_0.3.6
#> [5] ggplot2_4.0.0
#> [6] rjson_0.2.23
#> [7] xfun_0.53
#> [8] plyranges_1.30.0
#> [9] caTools_1.18.3
#> [10] Biobase_2.70.0
#> [11] lattice_0.22-7
#> [12] vctrs_0.6.5
#> [13] tools_4.5.1
#> [14] bitops_1.0-9
#> [15] curl_7.0.0
#> [16] parallel_4.5.1
#> [17] tibble_3.3.0
#> [18] AnnotationDbi_1.72.0
#> [19] RSQLite_2.4.3
#> [20] blob_1.2.4
#> [21] pkgconfig_2.0.3
#> [22] R.oo_1.27.1
#> [23] KernSmooth_2.23-26
#> [24] Matrix_1.7-4
#> [25] data.table_1.17.8
#> [26] RColorBrewer_1.1-3
#> [27] S7_0.2.0

```

```

#> [28] cigarillo_1.0.0
#> [29] lifecycle_1.0.4
#> [30] farver_2.1.2
#> [31] compiler_4.5.1
#> [32] gplots_3.2.0
#> [33] codetools_0.2-20
#> [34] htmltools_0.5.8.1
#> [35] RCurl_1.98-1.17
#> [36] yaml_2.3.10
#> [37] TxDb.Hsapiens.UCSC.hg19.knownGene_3.22.1
#> [38] pillar_1.11.1
#> [39] crayon_1.5.3
#> [40] R.utils_2.13.0
#> [41] BiocParallel_1.44.0
#> [42] cachem_1.1.0
#> [43] DelayedArray_0.36.0
#> [44] abind_1.4-8
#> [45] gtools_3.9.5
#> [46] tidyselect_1.2.1
#> [47] digest_0.6.37
#> [48] dplyr_1.1.4
#> [49] restfulr_0.0.16
#> [50] fastmap_1.2.0
#> [51] grid_4.5.1
#> [52] TxDb.Hsapiens.UCSC.hg38.knownGene_3.22.0
#> [53] cli_3.6.5
#> [54] SparseArray_1.10.0
#> [55] magrittr_2.0.4
#> [56] S4Arrays_1.10.0
#> [57] GenomicFeatures_1.62.0
#> [58] dichromat_2.0-0.1
#> [59] XML_3.99-0.19
#> [60] scales_1.4.0
#> [61] bit64_4.6.0-1
#> [62] rmarkdown_2.30
#> [63] httr_1.4.7
#> [64] matrixStats_1.5.0
#> [65] bit_4.6.0
#> [66] png_0.1-8
#> [67] R.methodsS3_1.8.2
#> [68] memoise_2.0.1
#> [69] evaluate_1.0.5
#> [70] knitr_1.50
#> [71] BiocIO_1.20.0
#> [72] rtracklayer_1.70.0
#> [73] rlang_1.1.6
#> [74] Rcpp_1.1.0
#> [75] glue_1.8.0
#> [76] DBI_1.2.3
#> [77] R6_2.6.1
#> [78] MatrixGenerics_1.22.0
#> [79] GenomicAlignments_1.46.0

```