

gmapR: Use the GMAP Suite of Tools in R

Michael Lawrence, Cory Barr

May 19, 2021

Contents

1	Introduction	2
2	What is GMAP, GSNAP, and bam_tally?	2
3	Create a <i>GmapGenome</i> Object	2
4	Aligning with GSNAP	3
5	Using bam_tally	5
6	Creating a <i>GmapGenome</i> Package	5

1 Introduction

The *gmapR* packages provides users with a way to access GSNAP, *bam_tally*, and other utilities from the GMAP suite of tools within an R session. In this vignette, we briefly look at the GMAP suite of tools available through the *gmapR* package and work through an example.

2 What is GMAP, GSNAP, and bam_tally?

The GMAP suite offers useful tools for the following:

- Genomic mapping: Given a cDNA, find where it best aligns to an entire genome
- Genomic alignment: Given a cDNA and a genomic segment, provide a nucleotide-level correspondence for the exons of the cDNA to the genomic segment
- Summarization via coverage plus reference and allele nucleotide polymorphism counts for an aligned set of sequencing reads over a given genomic location

GMAP (Genomic Mapping and Alignment Program) is particularly suited to relatively long mRNA and EST sequences such as those that are obtained from Roche 454 or Pacific Biosciences sequencing technologies. (At present, only GSNAP is available through the *gmapR*. GMAP integration is scheduled for the near future.)

GSNAP (Genomic Short-read Nucleotide Alignment Program) also provides users with genomic mapping and alignment capabilities, but is optimized to handle issues that arise when dealing with the alignment of short reads generated from sequencing technologies such as those from Illumina/Solexa or ABI/SOLiD. GSNAP offers the following functionality, as mentioned in Fast and SNP-tolerant detection of complex variants and splicing in short reads by Thomas D. Wu and Serban Nacu:

- fast detection of complex variants and splicing in short reads, based on a successively constrained search process of merging and filtering position lists from a genomic index
- alignment of both single- and paired-end reads as short as 14 nt and of arbitrarily long length
- detection of short- and long-distance splicing, including interchromosomal splicing, in individual reads, using probabilistic models or a database of known splice sites
- SNP-tolerant alignment to a reference space of all possible combinations of major and minor alleles
- alignment of reads from bisulfite-treated DNA for the study of methylation state

bam_tally provides users with the coverage as well as counts of both reference alleles, alternative alleles, and indels over a genomic region.

For more detailed information on the GMAP suite of tools including a detailed explication of algorithmic specifics, see <http://research-pub.gene.com/gmap/>.

3 Create a *GmapGenome* Object

To align reads with GMAP or GSNAP, or to use *bam_tally*, you will need to either obtain or create a *GmapGenome* object. *GmapGenome* objects can be created from FASTA files or *BSgenome* objects, as the following example demonstrates:

```
> library(gmapR)
> if (!suppressWarnings(require(BSgenome.Dmelanogaster.UCSC.dm3))) {
+   if (!requireNamespace("BiocManager", quietly=TRUE))
```

```

+     install.packages("BiocManager")
+   BiocManager::install("BSgenome.Dmelanogaster.UCSC.dm3")
+   library(BSgenome.Dmelanogaster.UCSC.dm3)
+ }
> gmapGenomePath <- file.path(getwd(), "flyGenome")
> gmapGenomeDirectory <- GmapGenomeDirectory(gmapGenomePath, create = TRUE)
> ##> gmapGenomeDirectory
> ##GmapGenomeDirectory object
> ##path: /reshpcfs/home/coryba/projects/gmapR2/testGenome
>
> gmapGenome <- GmapGenome(genome=Dmelanogaster,
+                           directory=gmapGenomeDirectory,
+                           name="dm3",
+                           create=TRUE,
+                           k = 12L)
> ##> gmapGenome
> ##GmapGenome object
> ##genome: dm3
> ##directory: /reshpcfs/home/coryba/projects/gmapR2/testGenome

```

4 Aligning with GSNAP

The GSNAP algorithm incorporates biological knowledge to provide accurate alignments, particularly for RNA-seq data. In this section, we will align reads from an RNA-seq experiment provided in a fastq file to a selected region of the human genome. In this example, we will align to the region of the human genome containing TP53 plus an additional one megabase on each side of this gene.

First we need to obtain the desired region of interest from the genome:

```

> library("org.Hs.eg.db")
> library("TxDb.Hsapiens.UCSC.hg19.knownGene")
> eg <- org.Hs.eg.db::org.Hs.egSYMBOL2EG[["TP53"]]
> txdb <- TxDb.Hsapiens.UCSC.hg19.knownGene
> tx <- transcripts(txdb, filter = list(gene_id = eg))
> roi <- range(tx) + 1e6
> strand(roi) <- "*"

```

Next we get the genetic sequence and use it to create a GmapGenome object. (Please note that the TP53_demo GmapGenome object is used by many examples and tests in the gmapR and VariationTools packages. If the object has been created before, its subsequent creation will be instantaneous.)

```

> library("BSgenome.Hsapiens.UCSC.hg19")
> library("gmapR")
> p53Seq <- getSeq(BSgenome.Hsapiens.UCSC.hg19:Hsapiens, roi,
+                 as.character = FALSE)
> names(p53Seq) <- "TP53"
> gmapGenome <- GmapGenome(genome = p53Seq,
+                           name = paste0("TP53_demo_",
+                           packageVersion("TxDb.Hsapiens.UCSC.hg19.knownGene")),
+                           create = TRUE,
+                           k = 12L)

```

We add the known transcripts (splice sites) to the genome index:

```
> exons <- gmapR::subsetRegion(exonsBy(txdb), roi, "TP53")
> spliceSites(gmapGenome, "knownGene") <- exons
```

The data package LungCancerLines contains fastqs of reads obtained from sequencing H1993 and H2073 cell lines. We will use these fastqs to demonstrate GSNAP alignment with gmapR.

```
> library("LungCancerLines")
> fastqs <- LungCancerFastqFiles()
```

GSNAP is highly configurable. Users create *GsnapParam* objects to specify desired GSNAP behavior.

```
> ##specify how GSNAP should behave using a GsnapParam object
> gsnapParam <- GsnapParam(genome=gmapGenome,
+                           unique_only=FALSE,
+                           suboptimal_levels=2L,
+                           npaths=10L,
+                           novelsplicing=TRUE,
+                           splicing="knownGene",
+                           indel_penalty=1L,
+                           distant_splice_penalty=1L,
+                           clip_overlap=TRUE)
```

Now we are ready to align.

```
> gsnapOutput <- gsnap(input_a=fastqs["H1993.first"],
+                       input_b=fastqs["H1993.last"],
+                       params=gsnapParam)
>
> ##gsnapOutput
> ##An object of class "GsnapOutput"
> ##Slot "path":
> ##[1] "/local/Rtmporwsvr"
> ##
> ##Slot "param":
> ##A GsnapParams object
> ##genome: dm3 (/reshpcfs/home/coryba/projects/gmapR2/testGenome)
> ##part: NULL
> ##batch: 2
> ##max_mismatches: NULL
> ##suboptimal_levels: 0
> ##snps: NULL
> ##mode: standard
> ##nthreads: 1
> ##novelsplicing: FALSE
> ##splicing: NULL
> ##npaths: 10
> ##quiet_if_excessive: FALSE
> ##nofails: FALSE
> ##split_output: TRUE
> ##extra: list()
> ##
> ##Slot "version":
> ## [1] NA NA NA NA NA NA NA NA NA NA NA NA NA NA
> ##
```

The `gsnapOutput` object will be of the class `GsnapOutput`. It will provide access to the BAM files of alignments that GSNAP outputs along with other utilities.

```
> ##> dir(path(gsnapOutput), full.names=TRUE, pattern="\\.bam$")
> ##[1] "/local/Rtmporwsvr/file1cbc73503e9.1.nomapping.bam"
> ##[2] "/local/Rtmporwsvr/file1cbc73503e9.1.unpaired_mult.bam"
> ##[3] "/local/Rtmporwsvr/file1cbc73503e9.1.unpaired_transloc.bam"
> ##[4] "/local/Rtmporwsvr/file1cbc73503e9.1.unpaired_uniq.bam"
```

5 Using bam_tally

Running the `bam_tally` method will return a `GRanges` of information per nucleotide. Below is an example demonstrating how to find variants in the TP53 gene of the human genome. See the documentation for the `bam_tally` method for more details.

`gmapR` provides access to a demo genome for examples. This genome encompasses the TP53 gene along with a 1-megabase flanking region on each side.

```
> genome <- TP53Genome()
```

The `LungCancerLines` R package contains a BAM file of reads aligned to the TP53 region of the human genome. We'll use this file to demonstrate the use of `bam_tally` through the `gmapR` package. The resulting data structure will contain the needed information such as number of alternative alleles, quality scores, and read position for each allele.

The call to `bam_tally` returns an opaque pointer to a C-level data structure. We anticipate having multiple means of summarizing these data into R data structures. For now, there is one: `variantSummary`, which returns a `VRanges` object describing putative genetic variants in the sample.

```
> bam_file <- system.file("extdata/H1993.analyzed.bam",
+                          package="LungCancerLines", mustWork=TRUE)
> breaks <- c(0L, 15L, 60L, 75L)
> bqual <- 56L
> mapq <- 13L
> param <- BamTallyParam(genome,
+                          minimum_mapq = mapq,
+                          concordant_only = FALSE, unique_only = FALSE,
+                          primary_only = FALSE,
+                          min_depth = 0L, variant_strand = 1L,
+                          ignore_query_Ns = TRUE,
+                          indels = FALSE, include_soft_clips = 1L, xs=TRUE,
+                          min_base_quality = 23L)
> tallies <- bam_tally(bam_file,
+                      param)
> variantSummary(tallies)
```

6 Creating a GmapGenome Package

After creating a `GmapGenome` object, you might want to distribute it for collaboration or version it for reproducible research. The function `makeGmapGenomePackage` allows you to do this. Continuing on with the *D. melanogaster* example from above, here is how to archive the *D. melanogaster* `GmapGenome` object in a `GmapGenome` package:

```

> makeGmapGenomePackage(gmapGenome=gmapGenome,
+                       version="1.0.0",
+                       maintainer="<your.name@somewhere.com>",
+                       author="Your Name",
+                       destDir="myDestDir",
+                       license="Artistic-2.0",
+                       pkgName="GmapGenome.Dmelanogaster.UCSC.dm3")

```

After creating the package, you can run R CMD INSTALL myDestDir to install it, or run R CMD build myDestDir to create a distribution of the package.

Many users will be working with the human genome. Many of the examples used in `gmapR` make use of a particular build of the human genome. As such, creating a *GmapGenome* of hg19 is recommended. Here is one way to create it, using a *BSgenome* object:

```

> if (!suppressWarnings(require(BSgenome.Hsapiens.UCSC.hg19))) {
+   if (!requireNamespace("BiocManager", quietly=TRUE))
+     install.packages("BiocManager")
+   BiocManager::install("BSgenome.Hsapiens.UCSC.hg19")
+   library(BSgenome.Hsapiens.UCSC.hg19)
+ }
> gmapGenome <- GmapGenome(genome=Hsapiens,
+                          directory = "Hsapiens",
+                          name = "hg19",
+                          create = TRUE)
> destDir <- "HsapiensGmapGenome"
> pkgName <- "GmapGenome.Hsapiens.UCSC.hg19"
> makeGmapGenomePackage(gmapGenome=gmapGenome,
+                       version="1.0.0",
+                       maintainer="<your.name@somewhere.com>",
+                       author="Your Name",
+                       destDir=destDir,
+                       license="Artistic-2.0",
+                       pkgName="GmapGenome.Hsapiens.UCSC.hg19")

```

After running the above code, you should be able to run R CMD INSTALL GmapGenome.Hsapiens.UCSC.hg19 in the appropriate directory from the command line, submit GmapGenome.Hsapiens.UCSC.hg19 to a repository, etc. After installation, `library("GmapGenome.Hsapiens.UCSC.hg19")` will load a *GmapGenome* object that has the same name as the package.

```

> sessionInfo()

```

```

R version 4.1.0 beta (2021-05-03 r80259)
Platform: x86_64-pc-linux-gnu (64-bit)
Running under: Ubuntu 20.04.2 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

```

```

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
```

attached base packages:

```
[1] parallel stats4 stats graphics grDevices utils datasets
[8] methods base
```

other attached packages:

```
[1] LungCancerLines_0.29.0
[2] gmapR_1.35.0
[3] Rsamtools_2.9.0
[4] BSgenome.Hsapiens.UCSC.hg19_1.4.3
[5] BSgenome_1.61.0
[6] rtracklayer_1.53.0
[7] Biostrings_2.61.0
[8] XVector_0.33.0
[9] TxDb.Hsapiens.UCSC.hg19.knownGene_3.2.2
[10] GenomicFeatures_1.45.0
[11] GenomicRanges_1.45.0
[12] GenomeInfoDb_1.29.0
[13] org.Hs.eg.db_3.13.0
[14] AnnotationDbi_1.55.0
[15] IRanges_2.27.0
[16] S4Vectors_0.31.0
[17] Biobase_2.53.0
[18] BiocGenerics_0.39.0
```

loaded via a namespace (and not attached):

```
[1] Rcpp_1.0.6 lattice_0.20-44
[3] prettyunits_1.1.1 png_0.1-7
[5] assertthat_0.2.1 digest_0.6.27
[7] utf8_1.2.1 BiocFileCache_2.1.0
[9] R6_2.5.0 RSQLite_2.2.7
[11] httr_1.4.2 pillar_1.6.1
[13] zlibbioc_1.39.0 rlang_0.4.11
[15] progress_1.2.2 curl_4.3.1
[17] rstudioapi_0.13 blob_1.2.1
[19] Matrix_1.3-3 BiocParallel_1.27.0
[21] stringr_1.4.0 RCurl_1.98-1.3
[23] bit_4.0.4 biomaRt_2.49.0
[25] DelayedArray_0.19.0 compiler_4.1.0
[27] pkgconfig_2.0.3 tidyselect_1.1.1
[29] KEGGREST_1.33.0 SummarizedExperiment_1.23.0
[31] tibble_3.1.2 GenomeInfoDbData_1.2.6
[33] matrixStats_0.58.0 XML_3.99-0.6
[35] fansi_0.4.2 crayon_1.4.1
[37] dplyr_1.0.6 dbplyr_2.1.1
[39] GenomicAlignments_1.29.0 bitops_1.0-7
[41] rappdirs_0.3.3 grid_4.1.0
```

[43] lifecycle_1.0.0	DBI_1.1.1
[45] magrittr_2.0.1	stringi_1.6.2
[47] cachem_1.0.5	ellipsis_0.3.2
[49] filelock_1.0.2	generics_0.1.0
[51] vctrs_0.3.8	rjson_0.2.20
[53] restfulr_0.0.13	tools_4.1.0
[55] bit64_4.0.5	glue_1.4.2
[57] purrr_0.3.4	MatrixGenerics_1.5.0
[59] hms_1.1.0	fastmap_1.1.0
[61] yaml_2.2.1	memoise_2.0.0
[63] VariantAnnotation_1.39.0	BiocIO_1.3.0