

Preparation of transcript expression and genotypes from the GEUVADIS project for sQTL analysis

Malgorzata Nowicka¹, Mark Robinson

¹gosia.nowicka@uzh.ch

October 31, 2017

This vignette describes version 1.6.0 of the *GeuvadisTranscriptExpr* package.

Contents

1	Description of the GEUVADIS project	1
2	Downloading the GEUVADIS data	2
3	Data preprocessing	2
3.1	Annotation and metadata.	2
3.2	Transcript counts	4
3.3	Genotypes	4
	APPENDIX	7
A	Session information	7
B	References	8

1 Description of the GEUVADIS project

In the GEUVADIS project [1], 462 RNA-Seq samples from lymphoblastoid cell lines were obtained. The genome sequencing data of the same individuals is provided by the 1000 Genomes Project. The samples in this project come from five populations: CEPH (CEU), Finns (FIN), British (GBR), Toscani (TSI) and Yoruba (YRI).

In the sQTL analysis, we want to identify genetic variants (here, bi-allelic SNPs) that are associated with changes in splicing. Such SNPs are then called splicing quantitative trait locies (sQTLs).

Preparation of transcript expression and genotypes from the GEUVADIS project for sQTL analysis

Ideally, we would like to test the associations of every SNP with every gene. However, such an approach would be very costly computationally and in terms of multiple testing correction. Under the assumption that SNPs that directly affect splicing are likely to be placed in the close surrounding of genes, we test only SNPs that are located within the gene body and within some range upstream and downstream of the gene.

2 Downloading the GEUVADIS data

In the *GeuvadisTranscriptExpr* package, we use transcript quantification (expected counts from FluxCapacitor) and genotypes available on the GEUVADIS project website <http://www.ebi.ac.uk/Tools/geuvadis-das/>.

As a reference genome, we use the Gencode v12 gene annotation available on <http://www.gencodegenes.org/releases/12.html>, which we save in a `gencode.v12.annotation.gtf` file.

From http://www.ebi.ac.uk/arrayexpress/files/E-GEUV-1/analysis_results/, we can download a file with information about samples `E-GEUV-1.sdrf.txt` and a file with transcript quantification `GD660.TrQuantCount.txt`. Let's save them under their original names.

The VCF files with genotypes are available on <http://www.ebi.ac.uk/arrayexpress/files/E-GEUV-1/genotypes/> in the compressed format.

The *R* code in this vignette assumes that all the files are saved in a current working directory.

3 Data preprocessing

In the following sections, we show how you can prepare the GEUVADIS data for the sQTL analysis. We preprocess counts and genotypes for all the populations and chromosomes of the human genome, but the package makes available only the preprocessed data for chromosome 19 from CEU population.

The *R* code in this section does not run when the vignette is generated, mainly because it takes some time to read in the files with transcript counts and genotypes.

3.1 Annotation and metadata

To find out which SNPs should be tested with which genes, we need to know the location of genes, which we get from the GTF file. We can import it into *R* using the *rtracklayer* package. We save the extracted locations of protein coding genes as BED files (about BED format, see <https://genome.ucsc.edu/FAQ/FAQformat.html>) for

Preparation of transcript expression and genotypes from the GEUVADIS project for sQTL analysis

every chromosome separately. We also remove the *'chr'* string from the chromosome names. Otherwise, we could not read in the VCF files with genotypes, which originally contain only a chromosome number.

In the *GeuvadisTranscriptExpr* package, you can access the `"genes_chr19.bed"` file.

```
library(GenomicRanges)
library(rtracklayer)

gtf_dir <- "gencode.v12.annotation.gtf"

gtf <- import(gtf_dir)

### Keep protein coding genes
keep_index <- mcols(gtf)$gene_type == "protein_coding" &
  mcols(gtf)$type == "gene"
gtf <- gtf[keep_index]
### Remove 'chr' from chromosome names
seqlevels(gtf) <- gsub(pattern = "chr", replacement = "", x = seqlevels(gtf))

genes_bed <- data.frame(chr = seqnames(gtf), start = start(gtf),
  end = end(gtf), geneId = mcols(gtf)$gene_id,
  stringsAsFactors = FALSE)

for(i in as.character(1:22)){
  genes_bed_sub <- genes_bed[genes_bed$chr == i, ]
  write.table(genes_bed_sub, "genes_chr", i, ".bed", quote = FALSE,
    sep = "\t", row.names = FALSE, col.names = FALSE)
}
```

The metadata information about sample names and the population they originate from can be accessed from the `E-GEUV-1.sdrf.txt` file, and we save it in `samples` data frame. Additionally, we create a variable with shorter sample names because such names are used in the genotype files.

```
library(limma)

metadata_dir <- "E-GEUV-1.sdrf.txt"

samples <- read.table(metadata_dir, header = TRUE, sep = "\t", as.is = TRUE)

samples <- unique(samples[c("Assay.Name", "Characteristics.population.")])
colnames(samples) <- c("sample_id", "population")

samples$sample_id_short <- strsplit2(samples$sample_id, "\\."),[,1]
```

3.2 Transcript counts

The `GD660.TrQuantCount.txt` file contains quantification for all the populations. We split this table by population and chromosome, and use the short sample names. Then the tables are saved in separate files.

In the `Geu vadisTranscriptExpr` package, you can access the `"TrQuantCount_CEU_chr19.tsv"` file.

```
expr_dir <- "GD660.TrQuantCount.txt"

expr_all <- read.table(expr_dir, header = TRUE, sep="\t", as.is = TRUE)

expr_all <- expr_all[, c("TargetID", "Gene_Symbol", "Chr",
  samples$sample_id)]
colnames(expr_all) <- c("TargetID", "Gene_Symbol", "Chr",
  samples$sample_id_short)

for(j in "CEU"){
  for(i in 1:22){
    expr <- expr_all[expr_all$Chr == i, c("TargetID", "Gene_Symbol",
      samples$sample_id_short[samples$population == j])]
    write.table(expr, paste0("TrQuantCount_", j, "_chr", i, ".tsv"),
      quote = FALSE, sep = "\t", row.names = FALSE, col.names = TRUE)
  }
}
```

3.3 Genotypes

VCF files can be loaded into *R* using the `VariantAnnotation` package (see the Annotating Genomic Variants Bioconductor workflow <https://bioconductor.org/help/workflows/variants/>). To do so, we have to first `bgzip` and create indexes with `indexTabix`. Both functions are from the `Rsamtools` package.

```
library(Rsamtools)

files <- list.files(path = ".", pattern = "genotypes.vcf.gz",
  full.names = TRUE, include.dirs = FALSE)

### bgzip and index the vcf files
for(i in 1:length(files)){
  zipped <- bgzip(files[i])
  idx <- indexTabix(zipped, format = "vcf")
}
```

Preparation of transcript expression and genotypes from the GEUVADIS project for sQTL analysis

We are interested in bi-allelic SNPs that lay within 5000 bases of a gene. Additionally, we want to convert the genotype information into 0 for ref/ref, 1 for ref/not ref, 2 for not ref/not ref, -1 or NA for missing values. Newly encoded genotypes are saved as text files, one file for each chromosome and population.

In the *GeuvadisTranscriptExpr* package, you can access the "genotypes_CEU_chr19.tsv" file.

```
library(VariantAnnotation)
library(tools)
```

```
### Extended gene ranges
window <- 5000
gene_ranges <- resize(gtf, GenomicRanges::width(gtf) + 2 * window,
  fix = "center")

chr <- gsub("chr", "", strsplit2(files, split = "\\.[.]", 2))

for(j in "CEU"){
  for(i in 1:length(files)){
    cat(j, chr[i], fill = TRUE)

    zipped <- paste0(file_path_sans_ext(files[i]), ".bgz")
    idx <- paste0(file_path_sans_ext(files[i]), ".bgz.tbi")
    tab <- TabixFile(zipped, idx)

    ### Explore the file header with scanVcfHeader
    hdr <- scanVcfHeader(tab)
    print(all(samples$sample_id_short %in% samples(hdr)))

    ### Read a subset of VCF file
    gene_ranges_tmp <- gene_ranges[seqnames(gene_ranges) == chr[i]]
    param <- ScanVcfParam(which = gene_ranges_tmp, samples =
      samples$sample_id_short[samples$population == j])
    vcf <- readVcf(tab, "hg19", param)

    ### Keep only the bi-allelic SNPs
    # width of ref seq
    rw <- width(ref(vcf))
    # width of first alt seq
    aw <- unlist(lapply(alt(vcf), function(x) {width(x[1])}))
    # number of alternate genotypes
    nalt <- elementLengths(alt(vcf))
    # select only bi-allelic SNPs (monomorphic OK, so aw can be 0 or 1)
    snp <- rw == 1 & aw <= 1 & nalt == 1
    # subset vcf
```

Preparation of transcript expression and genotypes from the GEUVADIS project for sQTL analysis

```
vcfbi <- vcf[snp,]

rowdata <- rowData(vcfbi)

### Convert genotype into number of alleles different from reference
geno <- geno(vcfbi)$GT
geno01 <- geno
geno01[, ] <- -1
geno01[geno %in% c("0/0", "0|0")] <- 0 # REF/REF
geno01[geno %in% c("0/1", "0|1", "1/0", "1|0")] <- 1 # REF/ALT
geno01[geno %in% c("1/1", "1|1")] <- 2 # ALT/ALT
mode(geno01) <- "integer"

genotypes <- unique(data.frame(chr = seqnames(rowdata),
  start = start(rowdata), end = end(rowdata), snpId = rownames(geno01),
  geno01, stringsAsFactors = FALSE))

### Sort SNPs by position
genotypes <- genotypes[order(genotypes[,2]), ]

write.table(genotypes, file = paste0("genotypes_", j, "_chr",
  chr[i], ".tsv"), quote = FALSE, sep = "\t", row.names = FALSE,
  col.names = TRUE)

}
}
```

APPENDIX

A Session information

```

sessionInfo()
## R version 3.4.2 (2017-09-28)
## Platform: x86_64-pc-linux-gnu (64-bit)
## Running under: Ubuntu 16.04.3 LTS
##
## Matrix products: default
## BLAS: /home/biocbuild/bbs-3.6-bioc/R/lib/libRblas.so
## LAPACK: /home/biocbuild/bbs-3.6-bioc/R/lib/libRlapack.so
##
## locale:
##  [1] LC_CTYPE=en_US.UTF-8      LC_NUMERIC=C
##  [3] LC_TIME=en_US.UTF-8      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] tools      parallel  stats4    stats     graphics  grDevices  utils
## [8] datasets  methods   base
##
## other attached packages:
## [1] VariantAnnotation_1.24.0 SummarizedExperiment_1.8.0
## [3] DelayedArray_0.4.0      matrixStats_0.52.2
## [5] Biobase_2.38.0          Rsamtools_1.30.0
## [7] Biostrings_2.46.0      XVector_0.18.0
## [9] limma_3.34.0           rtracklayer_1.38.0
## [11] GenomicRanges_1.30.0   GenomeInfoDb_1.14.0
## [13] IRanges_2.12.0         S4Vectors_0.16.0
## [15] BiocGenerics_0.24.0    knitr_1.17
##
## loaded via a namespace (and not attached):
## [1] Rcpp_0.12.13           compiler_3.4.2
## [3] highr_0.6              prettyunits_1.0.2
## [5] progress_1.1.2         GenomicFeatures_1.30.0
## [7] bitops_1.0-6          zlibbioc_1.24.0
## [9] biomaRt_2.34.0        bit_1.1-12
## [11] digest_0.6.12         BSgenome_1.46.0
## [13] memoise_1.1.0         tibble_1.3.4
## [15] RSQLite_2.0           evaluate_0.10.1

```

Preparation of transcript expression and genotypes from the GEUVADIS project for sQTL analysis

```
## [17] lattice_0.20-35      rlang_0.1.2
## [19] Matrix_1.2-11        DBI_0.7
## [21] yaml_2.1.14          GenomeInfoDbData_0.99.1
## [23] stringr_1.2.0        bit64_0.9-7
## [25] rprojroot_1.2        grid_3.4.2
## [27] R6_2.2.2             AnnotationDbi_1.40.0
## [29] RMySQL_0.10.13       XML_3.98-1.9
## [31] BiocParallel_1.12.0  rmarkdown_1.6
## [33] blob_1.1.0           magrittr_1.5
## [35] backports_1.1.1      codetools_0.2-15
## [37] htmltools_0.3.6     GenomicAlignments_1.14.0
## [39] assertthat_0.2.0    BiocStyle_2.6.0
## [41] stringi_1.1.5       RCurl_1.95-4.8
```

B References

References

- [1] T. Lappalainen, M. Sammeth, M. R. Friedländer, P. A. C. 't Hoen, J. Monlong, M. A. Rivas, M. González-Porta, N. Kurbatova, T. Griebel, P. G. Ferreira, M. Barann, T. Wieland, L. Greger, M. van Iterson, J. Almlöf, P. Ribeca, I. Pulyakhina, D. Esser, T. Giger, A. Tikhonov, M. Sultan, G. Bertier, D. G. MacArthur, M. Lek, E. Lizano, H. P. J. Buermans, I. Padioleau, T. Schwarzmayr, O. Karlberg, H. Ongen, H. Kilpinen, S. Beltran, M. Gut, K. Kahlem, V. Amstislavskiy, O. Stegle, M. Pirinen, S. B. Montgomery, P. Donnelly, M. I. McCarthy, P. Flicek, T. M. Strom, H. Lehrach, S. Schreiber, R. Sudbrak, A. Carracedo, S. E. Antonarakis, R. Häsler, A.-C. Syvänen, G.-J. van Ommen, A. Brazma, T. Meitinger, P. Rosenstiel, R. Guigó, I. G. Gut, X. Estivill, and E. T. Dermitzakis, "Transcriptome and genome sequencing uncovers functional variation in humans.," *Nature*, vol. 501, no. 7468, pp. 506–11, 2013.