

Generation of transcript counts from pasilla dataset with kallisto

Malgorzata Nowicka¹, Mark Robinson

October 10, 2026

This vignette describes version 1.41.0 of the *PasillaTranscriptExpr* package.

Contents

1	Description of pasilla dataset	1
2	Required software	2
3	Downloading the pasilla data	2
4	Downloading the reference genome	3
5	Transcript quantification with kallisto	3
	APPENDIX	7
A	Session information	7
B	References	8

1 Description of pasilla dataset

The *pasilla* dataset was produced by Brooks et al. [1]. The aim of their study was to identify exons that are regulated by pasilla protein, the *Drosophila melanogaster* ortholog of mammalian NOVA1 and NOVA2 (well studied splicing factors). In their RNA-seq experiment, the libraries were prepared from 7 biologically independent samples: 4 control samples and 3 samples in which pasilla was knocked-down. The libraries were sequenced on the Illumina Genome Analyzer II using single-end and paired-end sequencing and different read lengths. The RNA-seq data can be downloaded from the NCBI's Gene Expression Omnibus (GEO) under the accession number GSE18508.

¹gosia.nowicka@uzh.ch

2 Required software

This work-flow can be run on a Unix-like operating system, i.e., Linux or MacOS X with bash shell. All commands, including the one that could be run from terminal window, are run from within *R* using `system()` function. The downloaded and generated files will be saved in the current working directory.

Brooks et al. deposited their data in the Short Read Archive. In order to convert SRA data into fastq format, you need to install the *SRA toolkit* available on <http://www.ncbi.nlm.nih.gov/Traces/sra/sra.cgi?view=software>.

For the transcript quantification, we use *kallisto* version 0.42.1 [2], which is an extremely fast program that quantifies abundances of transcripts. *kallisto* is based on the novel idea of pseudoalignment to rapidly determine the compatibility of reads with transcripts, without the need for alignment. Thus it works directly on fastq files. The quantification is available in transcripts per million (TPM) and in expected counts. In this package, we make available the expected counts. *kallisto* can be downloaded from <http://pachterlab.github.io/kallisto/>.

3 Downloading the pasilla data

We use an automated process to download the SRA files that correspond to 4 control (Untreated) samples and 3 pasilla knocked-down (CG8144_RNAi) samples. All the information about the pasilla assay can be found in the metadata file *SraRunInfo.csv*, which can be downloaded from <http://www.ncbi.nlm.nih.gov/sra?term=SRP001537> under *Send to:* → *File* → *RunInfo* → *Create File*. The same file is also available within this package in the `extdata` directory.

```
library(PasillaTranscriptExpr)

data_dir <- system.file("extdata", package = "PasillaTranscriptExpr")

sri <- read.table(paste0(data_dir, "/SraRunInfo.csv"), stringsAsFactors = FALSE,
  sep = ",", header = TRUE)
keep <- grep("CG8144|Untreated-", sri$LibraryName)
sri <- sri[keep, ]

sra_files <- basename(sri$download_path)

for(i in 1:nrow(sri))
  download.file(sri$download_path[i], sra_files[i])
```

To convert the SRA files to fastq format, we use the *fastq-dump* command from the *SRA toolkit*. Then, we compress the fastq files.

Generation of transcript counts from pasilla dataset with kallisto

```
cmd <- paste0("fastq-dump -O ./ --split-3 ", sra_files)

for(i in 1:length(cmd))
  system(cmd[i])

system("gzip *.fastq")
```

4 Downloading the reference genome

To run *kallisto*, you need to download a FASTA formatted file of target sequences:

```
system("wget ftp://ftp.ensembl.org/pub/release-70/fasta/drosophila_melanogaster/cdna/Drosophila_melano
system("gunzip Drosophila_melanogaster.BDGP5.70.cdna.all.fa.gz")
```

The output produced by *kallisto* contains only the transcript IDs. To add the corresponding gene IDs, we need to download the gene model annotation in GTF format:

```
system("wget ftp://ftp.ensembl.org/pub/release-70/gtf/drosophila_melanogaster/Drosophila_melanogast
system("gunzip Drosophila_melanogaster.BDGP5.70.gtf.gz")
```

5 Transcript quantification with kallisto

We create a metadata file where each row corresponds to a collection of information needed for a single call of *kallisto*. The *pasilla* data consists of paired-end and single-end samples. When you run *kallisto* on single-end reads, you have to specify an `-l` option which defines the average fragment length. It can be found in `sri$avgLength`. There is one sample (GSM461179) which was sequenced using different read lengths. Therefore, for this sample, we do the transcript quantification for each read length separately and we add the resulting transcript counts in another step.

```
sri$LibraryName <- gsub("S2_DRSC_", "", sri$LibraryName)
metadata <- unique(sri[,c("LibraryName", "LibraryLayout", "SampleName",
  "avgLength")])

for(i in seq_len(nrow(metadata))){
  indx <- sri$LibraryName == metadata$LibraryName[i]

  if(metadata$LibraryLayout[i] == "PAIRED"){
    metadata$fastq[i] <- paste0(sri$Run[indx], "_1.fastq.gz ",
      sri$Run[indx], "_2.fastq.gz", collapse = " ")
  }else{
    metadata$fastq[i] <- paste0(sri$Run[indx], ".fastq.gz", collapse = " ")
  }
}
```

Generation of transcript counts from pasilla dataset with kallisto

```
}

metadata$condition <- ifelse(grepl("CG8144_RNAi", metadata$LibraryName),
  "KD", "CTL")
metadata$UniqueName <- paste0(1:nrow(metadata), "_", metadata$SampleName)
```

In the first step of *kallisto* work-flow, we build an index with *kallisto index*:

```
cDNA_fasta <- "Drosophila_melanogaster.BDGP5.70.cdna.all.fa"
index <- "Drosophila_melanogaster.BDGP5.70.cdna.all.idx"

cmd <- paste("kallisto index -i", index, cDNA_fasta, sep = " ")
cmd
## [1] "kallisto index -i Drosophila_melanogaster.BDGP5.70.cdna.all.idx Drosophila_melanogaster.BDGP5.70.cdna.all.fa"
```

```
system(cmd)
```

The quantification is done with *kallisto quant* command:

```
out_dir <- metadata$UniqueName

cmd <- paste("kallisto quant -i", index, "-o", out_dir, "-b 0 -t 5",
  ifelse(metadata$LibraryLayout == "SINGLE",
    paste("--single -l", metadata$avgLength), ""), metadata$fastq)
cmd
## [1] "kallisto quant -i Drosophila_melanogaster.BDGP5.70.cdna.all.idx -o 1_GSM461176 -b 0 -t 5 --single -l 1000"
## [2] "kallisto quant -i Drosophila_melanogaster.BDGP5.70.cdna.all.idx -o 2_GSM461177 -b 0 -t 5 --single -l 1000"
## [3] "kallisto quant -i Drosophila_melanogaster.BDGP5.70.cdna.all.idx -o 3_GSM461178 -b 0 -t 5 --single -l 1000"
## [4] "kallisto quant -i Drosophila_melanogaster.BDGP5.70.cdna.all.idx -o 4_GSM461179 -b 0 -t 5 --single -l 1000"
## [5] "kallisto quant -i Drosophila_melanogaster.BDGP5.70.cdna.all.idx -o 5_GSM461179 -b 0 -t 5 --single -l 1000"
## [6] "kallisto quant -i Drosophila_melanogaster.BDGP5.70.cdna.all.idx -o 6_GSM461179 -b 0 -t 5 --single -l 1000"
## [7] "kallisto quant -i Drosophila_melanogaster.BDGP5.70.cdna.all.idx -o 7_GSM461180 -b 0 -t 5 --single -l 1000"
## [8] "kallisto quant -i Drosophila_melanogaster.BDGP5.70.cdna.all.idx -o 8_GSM461181 -b 0 -t 5 --single -l 1000"
## [9] "kallisto quant -i Drosophila_melanogaster.BDGP5.70.cdna.all.idx -o 9_GSM461182 -b 0 -t 5 --single -l 1000"
```

```
for(i in 1:length(cmd))
  system(cmd[i])
```

We want to add the gene information and merge the expected transcript counts from different samples into one table.

```
library(rtracklayer)
```

```
gtf_dir <- "Drosophila_melanogaster.BDGP5.70.gtf"
```

Generation of transcript counts from pasilla dataset with kallisto

```
gtf <- import(gtf_dir)

gt <- unique(mcols(gtf)[, c("gene_id", "transcript_id")])
rownames(gt) <- gt$transcript_id

samples <- unique(metadata$SampleName)

counts_list <- lapply(1:length(samples), function(i){
  indx <- which(metadata$SampleName == samples[i])

  if(length(indx) == 1){
    abundance <- read.table(file.path(metadata$UniqueName[indx],
      "abundance.txt"), header = TRUE, sep = "\t", as.is = TRUE)
  }else{
    abundance <- lapply(indx, function(j){
      abundance_tmp <- read.table(file.path(metadata$UniqueName[j],
        "abundance.txt"), header = TRUE, sep = "\t", as.is = TRUE)
      abundance_tmp <- abundance_tmp[, c("target_id", "est_counts")]
      abundance_tmp
    })
    abundance <- Reduce(function(...) merge(..., by = "target_id", all = TRUE,
      sort = FALSE), abundance)
    est_counts <- rowSums(abundance[, -1])
    abundance <- data.frame(target_id = abundance$target_id,
      est_counts = est_counts, stringsAsFactors = FALSE)
  }

  counts <- data.frame(abundance$target_id, abundance$est_counts,
    stringsAsFactors = FALSE)
  colnames(counts) <- c("feature_id", samples[i])
  return(counts)
})

counts <- Reduce(function(...) merge(..., by = "feature_id", all = TRUE,
  sort = FALSE), counts_list)

### Add gene IDs
counts$gene_id <- gt[counts$feature_id, "gene_id"]
```

At the end, we keep only the unique samples in our metadata file.

```
metadata <- unique(metadata[, c("LibraryName", "LibraryLayout", "SampleName",
  "condition")])
metadata
##      LibraryName LibraryLayout SampleName condition
```

Generation of transcript counts from pasilla dataset with kallisto

## 156	Untreated-1	SINGLE	GSM461176	CTL
## 162	Untreated-3	PAIRED	GSM461177	CTL
## 164	Untreated-4	PAIRED	GSM461178	CTL
## 166	CG8144_RNAi-1	SINGLE	GSM461179	KD
## 172	CG8144_RNAi-3	PAIRED	GSM461180	KD
## 174	CG8144_RNAi-4	PAIRED	GSM461181	KD
## 176	Untreated-6	SINGLE	GSM461182	CTL

```
write.table(metadata, "metadata.txt", quote = FALSE, sep = "\t",  
            row.names = FALSE)
```

```
### Final counts with columns sorted as in metadata  
counts <- counts[, c("feature_id", "gene_id", metadata$SampleName)]  
write.table(counts, "counts.txt", quote = FALSE, sep = "\t", row.names = FALSE)
```

APPENDIX

A Session information

```
sessionInfo()
## R version 4.6.1 (2026-06-24)
## Platform: x86_64-pc-linux-gnu
## Running under: Ubuntu 26.04.1 LTS
##
## Matrix products: default
## BLAS:   /usr/lib/x86_64-linux-gnu/openblas-pthread/libblas.so.3
## LAPACK: /usr/lib/x86_64-linux-gnu/openblas-pthread/libopenblas-p-r0.3.32.so; LAPACK version 3.12
##
## locale:
##  [1] LC_CTYPE=en_US.UTF-8      LC_NUMERIC=C
##  [3] LC_TIME=en_US.UTF-8      LC_COLLATE=en_US.UTF-8
##  [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: Etc/UTC
## tzcode source: system (glibc)
##
## attached base packages:
## [1] stats4      stats      graphics  grDevices  utils      datasets  methods
## [8] base
##
## other attached packages:
## [1] rtracklayer_1.73.0      GenomicRanges_1.65.4
## [3] Seqinfo_1.3.2           IRanges_2.47.5
## [5] S4Vectors_0.51.10      BiocGenerics_0.59.12
## [7] generics_0.1.4         PasillaTranscriptExpr_1.41.0
## [9] knitr_1.52
##
## loaded via a namespace (and not attached):
## [1] BiocStyle_2.41.0      crayon_1.5.3
## [3] compiler_4.6.1       BiocManager_1.30.27
## [5] rjson_0.2.23         highr_0.12
## [7] BiocBaseUtils_1.15.1  GenomicAlignments_1.49.2
## [9] Rsamtools_2.29.0     Biostrings_2.81.9
## [11] bitops_1.1-0         parallel_4.6.1
## [13] BiocParallel_1.47.0   yaml_2.3.12
## [15] fastmap_1.2.0        R6_2.6.1
## [17] XVector_0.53.0       curl_8.0.0
```

Generation of transcript counts from pasilla dataset with kallisto

```
## [19] XML_3.99-0.25      maketools_1.3.2
## [21] rlang_1.3.0        xfun_0.61
## [23] sys_3.4.3          otel_0.2.0
## [25] cli_3.6.6          digest_0.6.39
## [27] evaluate_1.0.5     cigarillo_1.3.1
## [29] codetools_0.2-20   buildtools_1.0.0
## [31] RCurl_1.98-1.20    restfulr_0.0.17
## [33] rmarkdown_2.32     httr_1.4.9
## [35] tools_4.6.1        BiocIO_1.23.3
## [37] htmltools_0.5.9
```

B References

References

- [1] A. N. Brooks, L. Yang, M. O. Duff, K. D. Hansen, J. W. Park, S. Dudoit, S. E. Brenner, and B. R. Graveley, "Conservation of an RNA regulatory map between *Drosophila* and mammals.," *Genome research*, vol. 21, no. 2, pp. 193–202, 2011.
- [2] N. L. Bray, H. Pimentel, P. Melsted, and L. Pachter, "Near-optimal RNA-Seq quantification."