

HTSeqGenie: a software package to analyse high-throughput sequencing experiments

Gregoire Pau, Cory Barr, Jens Reeder, Michael Lawrence
Jeremiah Degenhardt, Tom Wu, Melanie Huntley, Matt Brauer

April 4, 2013

Contents

1	Introduction	2
2	Example analysis of two RNA-Seq experiments of lung cell lines	2
2.1	Analysis on a genomic region centered on TP53	2
2.2	Analysis on the full human genome	4
3	Analysis steps	4
3.1	General	4
3.2	Preprocessing	6
3.3	Alignment	6
3.4	Count genomic features	6
3.5	Coverage	7
4	Configuration	8
4.1	General	8
4.2	Configuration parameters	8
5	Session Information	11

1 Introduction

The *HTSeqGenie* package is a robust and efficient software to analyze high-throughput sequencing experiments in a reproducible manner. It supports the RNA-Seq and Exome-Seq protocols and provides: quality control reporting (using the *ShortRead* package), detection of adapter contamination, read alignment versus a reference genome (using the *gmapR* package), counting reads in genomic regions (using the *GenomicRanges* package), and read-depth coverage computation.

2 Example analysis of two RNA-Seq experiments of lung cell lines

2.1 Analysis on a genomic region centered on TP53

In this section, we are analysing two RNA-Seq experiments of human lung cell lines (H1993 and H2073) on a 2 Mb genomic region centered on the TP53 gene. This region, not the full human genome, has been chosen to provide a vignette that can be run in a reasonable amount of time.

We first load the package *HTSeqGenie*. We then load the package *LungCancerLines* to get the FASTQ files of the lung cancer cell lines we are going to analyze. These files only cover the TP53 genomic region and have been subsampled to 2500 reads, due to vignette time constraints.

```
> library("HTSeqGenie")
> library("LungCancerLines")
```

To analyze the reads, the package needs a **GmapGenome** object, which contains the information needed to align the reads to a reference genome. The package also needs a genomic features object, which is a named list of **GRanges** objects containing the genomic features (gene, exon, transcript...) of the corresponding genome. For convenience, this vignette provides the function **TP53Genome**, generating a 2 Mb subset of the UCSC hg19 human genome, centered on the gene TP53, and the **TP53GenomicFeatures**, producing a RData file containing the associated genomic features. The next section explains how to build custom **GmapGenome** and genomic features objects.

```
> tp53Genome <- TP53Genome()
> tp53GenomicFeatures <- TP53GenomicFeatures()
```

The samples are analysed with a common set of configuration parameters, specific to RNA-Seq paired-end experiments. Details about these parameters are provided in the Configuration section.

```
> rtp53 <- list(
+   ## aligner
+   path.gsnap_genomes=path(directory(tp53Genome)),
+   alignReads.genome=genome(tp53Genome),
+
+   ## gene model
+   path.genomic_features=dirname(tp53GenomicFeatures),
+   countGenomicFeatures.gfeatures=basename(tp53GenomicFeatures)
+ )
```

The function **runPipeline** runs the full analysis pipeline, using a list of configuration parameters. The following commands analyse the H1993 and H2073 samples. The function returns the path to a directory that contains the full analysis: bams, QC reports and a results directory which include coverage and count data.

```

> H1993dir <- do.call(runPipeline,
+                   c(## RNASeq TP53genome parameters
+                     rtp53,
+
+                     ## input
+                     input_file=LungCancerFastqFiles()[["H1993.first"]],
+                     input_file2=LungCancerFastqFiles()[["H1993.last"]],
+                     paired_ends=TRUE,
+                     quality_encoding="illumina1.8",
+
+                     ## output
+                     save_dir="H1993",
+                     prepend_str="H1993",
+                     overwrite_save_dir="erase"
+                   ))
> H2073dir <- do.call(runPipeline,
+                   c(## RNASeq TP53genome parameters
+                     rtp53,
+
+                     ## input
+                     input_file=LungCancerFastqFiles()[["H2073.first"]],
+                     input_file2=LungCancerFastqFiles()[["H2073.last"]],
+                     paired_ends=TRUE,
+                     quality_encoding="illumina1.8",
+
+                     ## output
+                     save_dir="H2073",
+                     prepend_str="H2073",
+                     overwrite_save_dir="erase"
+                   ))

```

Several quality control reports (short reads, alignment, counting) have been generated in the `report/` directories of the paths pointed by `H1993dir` and `H2073dir`. The following example uses the function `getTabDataFromFile` to get the read counts per gene in both cell lines. Further analysis such as differential expression using replicates can be performed using the package *DESeq*.

```

> library("org.Hs.eg.db")
> gc1 <- getTabDataFromFile(H1993dir, "counts_gene")
> gc2 <- getTabDataFromFile(H2073dir, "counts_gene")
> hgnc <- as.character(org.Hs.egSYMBOL[as.character(gc1$name)])
> data.frame(H1993.count=gc1$count, H2073.count=gc2$count, row.names=hgnc)

```

	H1993.count	H2073.count
LOC100128288	0	0
LOC100506713	0	0
MIR497HG	0	0
C17orf61-PLSCR3	4	9
CLEC10A	5	11
SAT2	0	0
GABARAP	0	0
KRBA2	0	0
CLDN7	0	0

SLC16A11	0	0
DLG4	3	19
DVL2	1	0
CTDNEP1	0	0
ALOX12B	54	16
TMEM256	0	4
LOC284023	1	1
LINC00324	0	0
SLC13A5	0	0
GPS2	0	0
SLC25A35	0	0
ASGR1	16	8
ASGR2	11	6
MIR324	0	0
MYH10	0	0
YBX2	1	0
PER1	1	3
C17orf59	0	0
PLSCR3	0	0
MIR497	0	0
ZBTB4	6	8
TRAPPC1	0	4
ALOXE3	108	151
RPL26	0	0
RNF222	0	2
SOX15	0	0
VAMP2	0	0
TP53	83	70
PHF23	0	0
CTC1	0	0
TEKT1	29	13
TMEM107	0	0
LSMD1	0	0
NEURL4	0	0
HES7	20	25
KCNAB3	0	0
AURKB	0	0
FXR2	0	1

2.2 Analysis on the full human genome

This section will be completed later.

3 Analysis steps

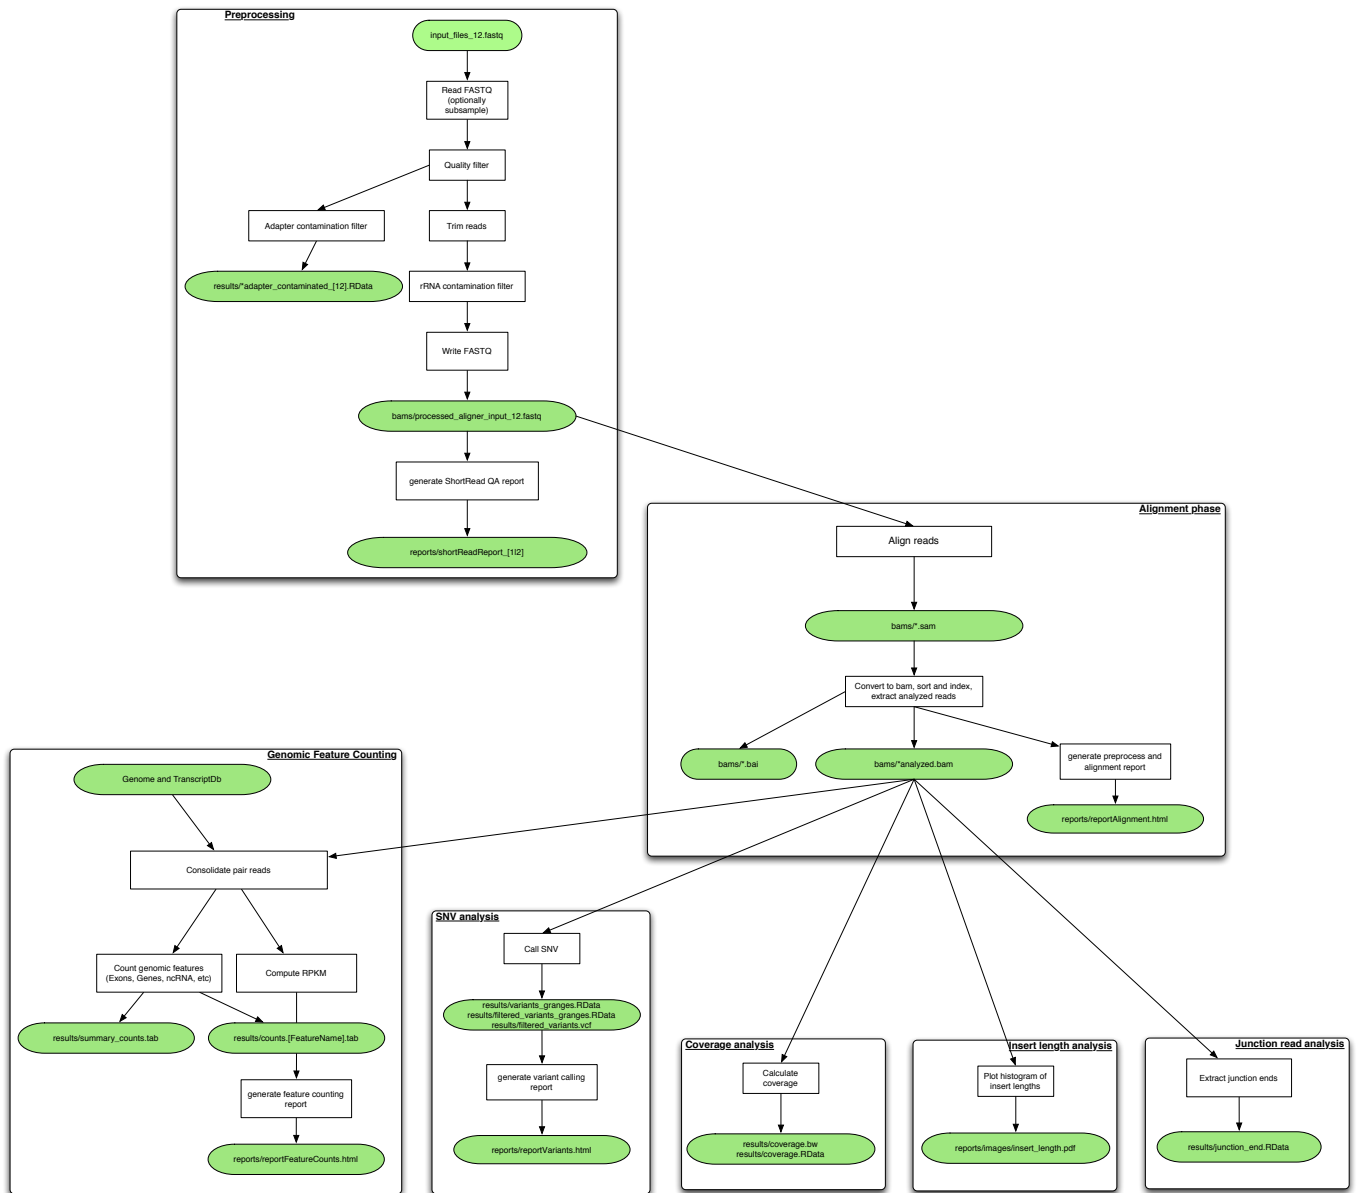
3.1 General

The pipeline consists in a set of high-level modules that are configured with a configuration file. See the Configuration section for a list of all parameters.

The modules are:

- preprocessing (preprocessReads)
- read alignment (alignReads)
- read counting (countGenomicFeatures)
- computation of coverage (calculateCoverage)

The following diagram shows a high-level overview of the performed analysis steps:



The pipeline produces the directory `logs/`, which contains log files about the analysis:

- `logs/audit.txt`: the environment used to run the analysis (pipeline version number, host, R session, PATH, gsnap version, samtools version)

- logs/config.txt: the expanded configuration file used to run the analysis
- logs/progress.log: the progress log

3.2 Preprocessing

The preprocessing module applies a sequence of optional steps before read alignment. These optional operations are:

- read subsampling (randomly sample a subset of reads from input data)
- trimming reads end to a fixed size
- filtering reads based on nucleotide qualities
- detection of (but not filter) adapter-contaminated reads
- creation of ShortRead reports (on a random subset of reads)

Output

- bams/processed.aligner.input.fastq: processed FASTQ file(s) given to the aligner
- results/adapter_contaminated.RData: RData files containing the names of the detected adapter-contaminated reads
- results/summary_preprocess.tab: a tab-separated file containing a summary of the preprocessing step
- reports/shortReadReport: HTML ShortRead quality reports

3.3 Alignment

The alignment module uses the gsnap aligner to align reads, and creates the file analyzed.bam, containing reads that are analyzed by the downstream modules. This file currently contains all uniquely mapping reads (i.e. this is the concatenation of the *uniq.bam files).

Output


- bams/*.bam: output aligned bam files
- bams/*.bai: corresponding bam indexes
- results/summary_alignment.tab: a tab-separated file containing summary alignment information
- report/reportAlignment.html: an HTML summary report about alignment

3.4 Count genomic features

This module counts how many reads overlaps with the genomic intervals defined in RData file pointed by the "countGenomicFeatures.gfeatures" parameter. On each interval, RPKM is computed using: $(1e9 * \text{count}) / (\text{nbreads} * \text{width})$, where count is the number of reads overlapping with the interval, nbreads is the total number of reads being analysed and width is the width of the interval.

The genomic features are defined as:

- A gene region is an interval (i.e. connected) and includes introns, 5' and 3' UTRs.



- A gene exonic region is an interval union of all exons of a particular gene. It includes 5' and 3' UTRs, but not introns.
- A gene coding region is a gene exonic region without the 5' and 3' UTRs.
- A transcript is an interval union of the exons of a particular transcript. It includes 5' and 3' UTRs, but not introns.
- Non-overlapping exonic regions (DEXSeq) are exons that have been split to remove overlapping regions.

Output

- results/counts_feature_type.tab: a tab-separated file containing counts, widths and RPKMs
- results/summary_counts.tab: a tab-separated file containing summary count information
- report/reportFeatureCounts.html: a HTML summary report about count genomic features

3.5 Coverage

This module computes coverage (read-depth) data.

Output

- results/coverage.RData: a RData file containing coverage data
- results/coverage.bw: bigwig file for IGB

4 Configuration

4.1 General

Configuration files are written in a DCF format, a sequence of lines containing "`<parameter>:<value>`". The format supports comments starting with a sharp (`#`) sign and is whitespace tolerant. The HTSeqGenie pipeline uses a templated configuration framework where configuration files can inherit parameters from other configuration files, through the parameter "`template_config`". Parameters defined after "`template_config`" override parameters from the inherited templates.

Default parameters depend on the configuration templates.

4.2 Configuration parameters

Configuration files ultimately derive from the `default-config.txt` configuration file, which enumerates all available parameters. The following sections enumerate all available configuration parameters.

Template

- `template_config`: An optional file name containing a template configuration file. To locate it, the software first looks in the directory defined by the environment variable `HTSEQGENIE_CONFIG` and, if not found, in the installed package HTSeqGenie. No default value.

Input

- `input_file`: The path of a FASTQ or gzipped FASTQ file containing the input reads. No default value.
- `input_file2`: An optional path of a FASTQ or gzipped FASTQ file containing the input paired reads in paired-end sequencing. If present, the parameter `paired_ends` must be set to `TRUE`. No default value.
- `paired_ends`: A logical indicating whether the reads are coming from paired-end sequencing. Default is `TRUE`.
- `quality_encoding`: An optional string indicating which quality encoding is used. Possible values are "sanger", "solexa", "illumina1.3", "illumina1.5" and "illumina1.8". If absent, a non-robust strategy is tried to guess it. No default value.
- `subsample_nreads`: An optional integer. If present, the preprocess module randomly subsamples the indicated number of reads from the input files before analysis. This feature is disabled by default.
- `chunk_size`: An integer indicating how many reads are present in a chunk. This parameter should not be changed. Default is `1e6`.
- `max_nbchunks`: An optional integer. If present, it sets the maximum number of chunks to be processed. This parameter is mostly present for debug purposes. This feature is disabled by default.

Output

- `save_dir`: A directory path indicating where the output results should be stored. No default value.
- `overwrite_save_dir`: A string indicating how to overwrite/save data in the `save_dir` directory. Possible values are "never", "erase" and "overwrite". If "never", the pipeline will stop if the indicated `save_dir` is present. If "erase", the `save_dir` directory is removed before starting the analysis. Default is "never".
- `prepend_str`: A string that will be prepended to output filenames. No default value.

- `remove_processedfastq`: A logical indicating whether the preprocessed FASTQ files before alignment should be removed. Default is TRUE.
- `remove_chunkdir`: A logical indicating whether the `chunks/` subdirectory should be removed. Default is TRUE.
- `tmp_dir`: A temporary directory where intermediate files will be written to. If set, this should point to a non-replicated, non-snapshotted directory such as `/gne/research/data/bioinfo/ngs_analysis/CGP_3.0/cgptmp`. If not set, temporary files will be written into the `save_dir`.

System

- `num_cores`: An integer indicating how many cores should be used on this machine. This value can be overridden by the `ngs_pipeline` calling script. Default value is 1.

Path

- `path.genomic_features`: A path to the directory that contains the genomic features RData files used for counting. No default.
- `path.gsnap_genomes`: A path to the directory that contains Gsnap's genomic indices. No default.

Debug

- `debug.level`: A string indicating which maximal level of debug should be used in the log file. Possible values are "ERROR", "WARN", "INFO" and "DEBUG". Default is "DEBUG".
- `debug.tracemem`: A logical indicating whether periodic memory information should be included in the log file. Default is TRUE.

Trim reads

- `trimReads.do`: A logical indicating whether the reads should be end-trimmed. Default is FALSE.
- `trimReads.length`: An integer indicating to what size (in nucleotides) the reads should be trimmed to. No default value.

Filter quality

- `filterQuality.do`: A logical indicating whether the reads should be filtered out based on read nucleotide qualities. A read is kept if the fraction "`filterQuality.minFrac`" of nucleotides have a quality higher than "`filterQuality.minQuality`". Default is TRUE.
- `filterQuality.minQuality`: An integer. See `filterQuality.do`. Default is 23.
- `filterQuality.minFrac`: A numeric ranging from 0.0 to 1.0. See `filterQuality.do`. Default is 0.7.

Detect adapter contamination

- `detectAdapterContam.do`: A logical indicating whether reads that look contaminated by adapter sequences should be reported (but NOT filtered). Default is TRUE.
- `detectAdapterContam.force_paired_end_adapter`: A logical indicating whether paired end adapter chemistry was used for single-end sequencing. Default is FALSE.

Detect ribosomal RNA

- `detectRRNA.do`: A logical indicating whether ribosomal RNAs should be filtered out. Default is TRUE.
- `detectRRNA.rna_genome`: A string indicating which gsnap genome index to use (-d flag) to detect ribosomal RNAs. No default value.

ShortRead report

- `shortReadReport.do`: A logical indicating whether a ShortRead report should be generated from the preprocessed reads. Default is TRUE.
- `shortReadReport.subsample_nbreads`: An optional integer indicating how many reads should be subsampled from the preprocessed reads for generating the report. No value indicates to take all of them (can be memory and time consuming). Default is 20e6.

Aligner

- `alignReads.genome`: A string indicating which gsnap genome index to use (-d flag). No default value.
- `alignReads.max_mismatches`: An optional integer indicating how many maximal mismatches gsnap should use (-m flag). If absent, gsnap uses an automatic value. No default value.
- `alignReads.sam_id`: A string (-read-group-id flag). No default value.
- `alignReads.snp_index`: An optional string containing gsnap SNP database index (-v flag). No default value.
- `alignReads.splice_index`: An optional string containing gsnap splice index (-s flag). No default value.
- `alignReads.static_parameters`: An optional string containing extra gsnap parameters. No default value.
- `alignReads.nbthreads_perchunk`: An optional integer indicating how many threads should be used to process one chunk (-t flag). If unspecified, the value is set to $\min(\text{num_cores}, 4)$. This parameter is mostly given for debug purposes. No default value.

Count genomic features

- `countGenomicFeatures.do`: A logical indicating whether counts per genomic feature should be computed. Default is TRUE.
- `countGenomicFeatures.gfeatures`: A filename containing the RData file used by the counting features module. The full path is "path.genomic_features"/ "countGenomicFeatures.gfeatures". No default value.

Variants

- `analyzeVariants.do`: A logical indicating whether variant calling should be performed. Default is TRUE.
- `analyzeVariants.use_read_length`: FALSE
- `analyzeVariants.with_qual`: A logical indicating whether variant calling should take mapping quality into account. Default is TRUE.
- `analyzeVariants.bqual`: Default is 23.
- `analyzeVariants.bin_fraction`: Default is 0.1.

Coverage

- `coverage.extendReads`: A logical whether reads should be extended in the coverage vector. Default is `FALSE`
- `coverage.fragmentLength`: Amount by which single ended reads should be extended. Usually defined by the fragment length. No default value.

5 Session Information

```
> toLatex(sessionInfo())
```

- R version 3.0.0 (2013-04-03), x86_64-unknown-linux-gnu
- Locale: `LC_CTYPE=en_US.UTF-8`, `LC_NUMERIC=C`, `LC_TIME=en_US.UTF-8`, `LC_COLLATE=C`,
`LC_MONETARY=en_US.UTF-8`, `LC_MESSAGES=en_US.UTF-8`, `LC_PAPER=C`, `LC_NAME=C`, `LC_ADDRESS=C`,
`LC_TELEPHONE=C`, `LC_MEASUREMENT=en_US.UTF-8`, `LC_IDENTIFICATION=C`
- Base packages: base, datasets, grDevices, graphics, methods, parallel, stats, tools, utils
- Other packages: AnnotationDbi 1.22.0, Biobase 2.20.0, BiocGenerics 0.6.0, Biostrings 2.28.0,
Cairo 1.5-2, DBI 0.2-5, GenomicRanges 1.12.0, HTSeqGenie 3.10.0, IRanges 1.18.0,
LungCancerLines 0.0.7, RColorBrewer 1.0-5, RSQLite 0.11.2, Rsamtools 1.12.0, ShortRead 1.18.0,
gmapR 1.2.0, hwriter 1.3, lattice 0.20-15, latticeExtra 0.6-24, org.Hs.eg.db 2.9.0, rtracklayer 1.20.0
- Loaded via a namespace (and not attached): BSgenome 1.28.0, GenomicFeatures 1.12.0,
Matrix 1.0-12, RBGL 1.36.0, RCurl 1.95-4.1, VariantAnnotation 1.6.0, VariantTools 1.2.0,
XML 3.96-1.1, biomaRt 2.16.0, bitops 1.0-5, graph 1.38.0, grid 3.0.0, stats4 3.0.0, zlibbioc 1.6.0