

Process scRNA-Seq reads in *scruff*

Zhe Wang

2018-10-30

Package

scruff 1.0.0

Contents

1	Introduction	2
2	Quick Start	2
3	Stepwise Tutorial For CEL-Seq Samples	3
3.1	Load Example Dataset	3
3.2	Demultiplex and Assign Cell Specific Reads	4
3.3	Alignment	5
3.4	UMI correction and Generation of Count Matrix	5
3.5	Visualization of QC metrics	6
3.6	Visualization of Read mapping locations.	11
4	10X BAM File Quality Assessment.	12
5	Session Information	14

1 Introduction

scruff is a toolkit for processing single cell RNA-seq FASTQ reads generated by CEL-Seq and CEL-Seq2 protocols. It does demultiplexing, alignment, Unique Molecular Identifier (UMI) filtering, and transcript counting in an automated fashion and generates the gene count matrix, QC metrics and provides visualizations of data quality. This vignette provides a brief introduction to the *scruff* package by walking through the demultiplexing, alignment, and UMI-counting of a built-in publicly available example dataset ([van den Brink, et al. 2017](#)).

2 Quick Start

```
# Run scruff on example dataset
# NOTE: Requires Rsubread index and TxDb objects for the reference genome.
# For generation of these files, please refer to the Stepwise Tutorial.

library(scruff)

# Get the paths to example FASTQ, FASTA, and GTF files.
v1h1R1 <- system.file("extdata",
  "vandenBrink_1h1_L001_R1_001.fastq.gz",
  package = "scruff")
v1h1R2 <- system.file("extdata",
  "vandenBrink_1h1_L001_R2_001.fastq.gz",
  package = "scruff")
vb1R1 <- system.file("extdata",
  "vandenBrink_b1_L001_R1_001.fastq.gz",
  package = "scruff")
vb1R2 <- system.file("extdata",
  "vandenBrink_b1_L001_R2_001.fastq.gz",
  package = "scruff")
fasta <- system.file("extdata", "GRCm38_MT.fa", package = "scruff")
gtf <- system.file("extdata", "GRCm38_MT.gtf", package = "scruff")
```

Build Rsubread alignment index. This is for the alignment step. For test purpose, here we are aligning the example FASTQ files to the genes on mitochondrial chromosome only.

```
# NOTE: Rsubread package does not support Windows environment.
if (!requireNamespace("Rsubread", quietly = TRUE)) {
  message("Package \"Rsubread\" needed for \"alignRsubread\",
    " function to work. ",
    "Please install it if you are using Linux or macOS systems. ",
    "The function is not available in Windows environment.")
} else {
  # Create index files for GRCm38_MT.
  # For details, please refer to Rsubread user manual.
  # Specify the basename for Rsubread index
  indexBase <- "GRCm38_MT"
  Rsubread::buildindex(basename = indexBase,
```

Process scRNA-Seq reads in *scruff*

```
reference = fasta,  
indexSplit = FALSE)  
}
```

Now that everything is ready, we can run *scruff*. In experiment 1h1, cell barcodes 95 and 96 are empty well controls. In experiment b1, cell barcode 95 is bulk sample containing 300 cells. This information can be set by the `cellPerWell` argument. *scruff* makes use of the [SingleCellExperiment](#) package. The following command returns a `SingleCellExperiment` object containing UMI filtered count matrix as well as gene and sample annotations and QC metrics.

```
data(barcodeExample, package = "scruff")  
  
if (!requireNamespace("Rsubread", quietly = TRUE)) {  
  message("Package \"Rsubread\" needed. ",  
    "Please install it if you are using Linux or macOS systems. ",  
    "The function is not available in Windows environment.")  
} else {  
  sce <- scruff(project = "example",  
    experiment = c("1h1", "b1"),  
    lane = c("L001", "L001"),  
    read1Path = c(v1h1R1, vb1R1),  
    read2Path = c(v1h1R2, vb1R2),  
    bc = barcodeExample,  
    index = indexBase,  
    unique = FALSE,  
    nBestLocations = 1,  
    reference = gtf,  
    bcStart = 1,  
    bcStop = 8,  
    umiStart = 9,  
    umiStop = 12,  
    keep = 75,  
    cellPerWell = c(rep(1, 94), 0, 0, rep(1, 94), 300, 1),  
    cores = 2,  
    verbose = TRUE)  
}
```

Visualize data quality.

```
data(sceExample, package = "scruff")  
qc <- qcplots(sceExample)
```

3 Stepwise Tutorial For CEL-Seq Samples

3.1 Load Example Dataset

The *scruff* package contains 4 single cell RNA-seq FASTQ example files. Each file has 10,000 sequenced reads.

Process scRNA-Seq reads in *scruff*

```
library(scruff)
v1h1R1 <- system.file("extdata",
  "vandenBrink_1h1_L001_R1_001.fastq.gz",
  package = "scruff")
v1h1R2 <- system.file("extdata",
  "vandenBrink_1h1_L001_R2_001.fastq.gz",
  package = "scruff")
vb1R1 <- system.file("extdata",
  "vandenBrink_b1_L001_R1_001.fastq.gz",
  package = "scruff")
vb1R2 <- system.file("extdata",
  "vandenBrink_b1_L001_R2_001.fastq.gz",
  package = "scruff")
```

3.2 Demultiplex and Assign Cell Specific Reads

Now the FASTQ files are ready to be demultiplexed. *scruff* package provides built-in predefined cell barcodes `barcodeExample` for demultiplexing the example dataset. In the example FASTQ files, read 1 contains cell barcode and UMI sequence information. Read 2 contains transcript sequences. The barcode sequence of each read starts at base 1 and ends at base 8. The UMI sequence starts at base 9 and ends at base 12. They can be set via `bcStart`, `bcStop`, and `umiStart`, `umiStop` arguments. By default, reads with any nucleotide in the barcode and UMI sequences with sequencing quality lower than 10 (Phred score) will be excluded. The following command demultiplexes the example FASTQ reads and trims reads longer than 75 nucleotides. The command returns a `SingleCellExperiment` object whose `colData` contains the cell index, barcode, reads, percentage of reads assigned, sample, and FASTQ file path information for each cell. By default, the cell specific demultiplexed fastq.gz files are stored in `./Demultiplex` folder.

```
data(barcodeExample, package = "scruff")
de <- demultiplex(project = "example",
  experiment = c("1h1", "b1"),
  lane = c("L001", "L001"),
  read1Path = c(v1h1R1, vb1R1),
  read2Path = c(v1h1R2, vb1R2),
  barcodeExample,
  bcStart = 1,
  bcStop = 8,
  bcEdit = 0,
  umiStart = 9,
  umiStop = 12,
  keep = 75,
  minQual = 10,
  yieldReads = 1e+06,
  verbose = TRUE,
  overwrite = TRUE,
  cores = 2)
```

3.3 Alignment

scruff provides an alignment function `alignRsubread` which is a wrapper function to `align` in `Rsubread` package. It aligns the reads to reference sequence index and outputs sequence alignment map files in “BAM” or “SAM” format. For demonstration purpose, the built-in mitochondrial DNA sequence from GRCm38 reference assembly `GRCm38MitochondrialFasta` will be used to map the reads. First, a `Rsubread` index for the reference sequence needs to be generated.

```
# Create index files for GRCm38_MT. For details, please refer to Rsubread
# user manual.
fasta <- system.file("extdata", "GRCm38_MT.fa", package = "scruff")
# NOTE: Rsubread package does not support Windows environment.
if (!requireNamespace("Rsubread", quietly = TRUE)) {
  message("Package \"Rsubread\" needed.",
    " Please install it if you are using Linux or macOS systems.",
    " The function is not available in Windows environment.\n")
} else {
  # Create index files for GRCm38_MT.
  # For details, please refer to Rsubread user manual.
  # Specify the basename for Rsubread index
  indexBase <- "GRCm38_MT"
  Rsubread::buildindex(basename = indexBase,
    reference = fasta,
    indexSplit = FALSE)
}
```

The following command maps the FASTQ files to GRCm38 mitochondrial reference sequence `GRCm38_MT.fa` and returns a `SingleCellExperiment` object. By default, the files are stored in BAM format in `./Alignment` folder.

```
# Align the reads using Rsubread
if (requireNamespace("Rsubread", quietly = TRUE)) {
  al <- alignRsubread(de,
    indexBase,
    unique = FALSE,
    nBestLocations = 1,
    format = "BAM",
    overwrite = TRUE,
    verbose = TRUE,
    cores = 2)
}
```

3.4 UMI correction and Generation of Count Matrix

Example GTF file `GRCm38_MT.gtf` will be used for feature counting. Currently, *scruff* applies the union counting mode of the HTSeq Python package. The following command generates the UMI filtered count matrix for the example dataset.

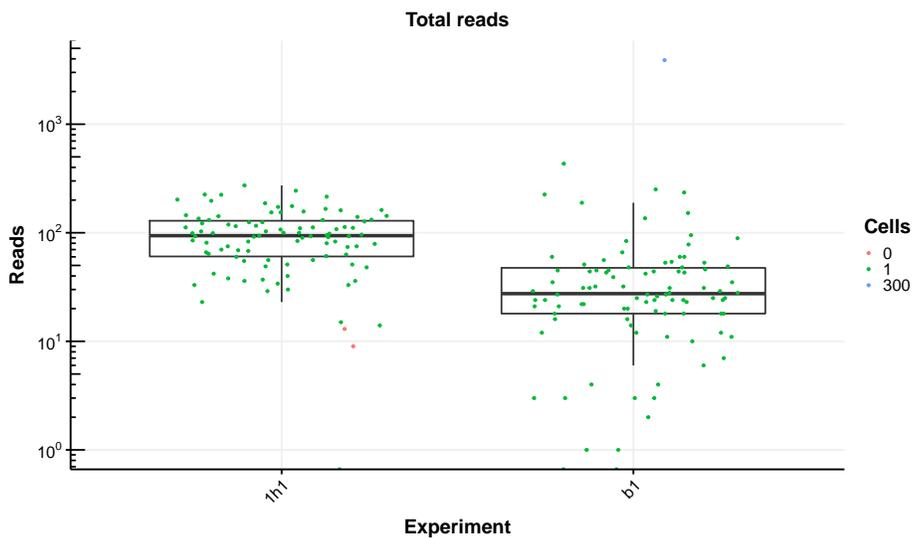
Process scRNA-Seq reads in *scruff*

```
gtf <- system.file("extdata", "GRCm38_MT.gtf", package = "scruff")
# get the molecular counts of transcripts for each cell
# In experiment 1h1, cell barcodes 95 and 96 are empty well controls.
# In experiment b1, cell barcode 95 is bulk sample containing 300 cells.
if (requireNamespace("Rsubread", quietly = TRUE)) {
  sce = countUMI(al,
    gtf,
    format = "BAM",
    cellPerWell = c(rep(1, 94), 0, 0, rep(1, 94), 300, 1),
    verbose = TRUE,
    cores = 2)
}
```

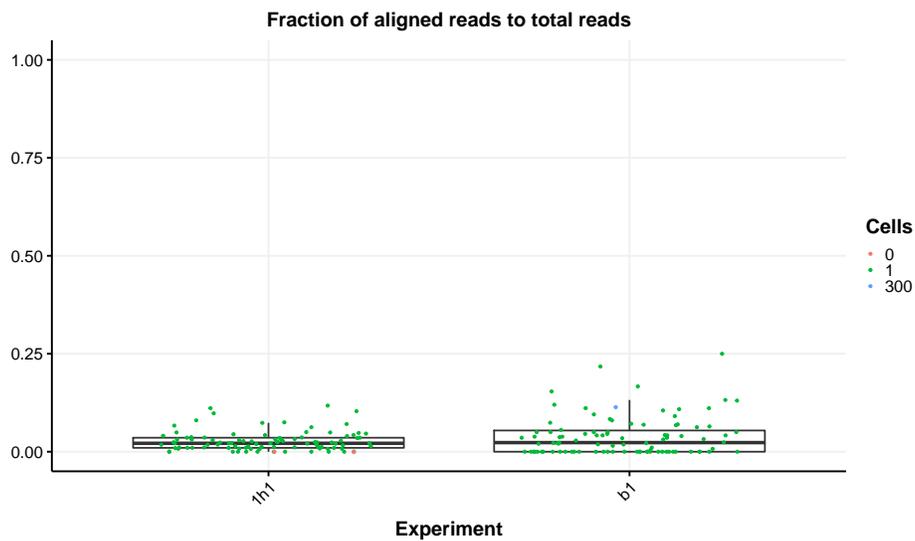
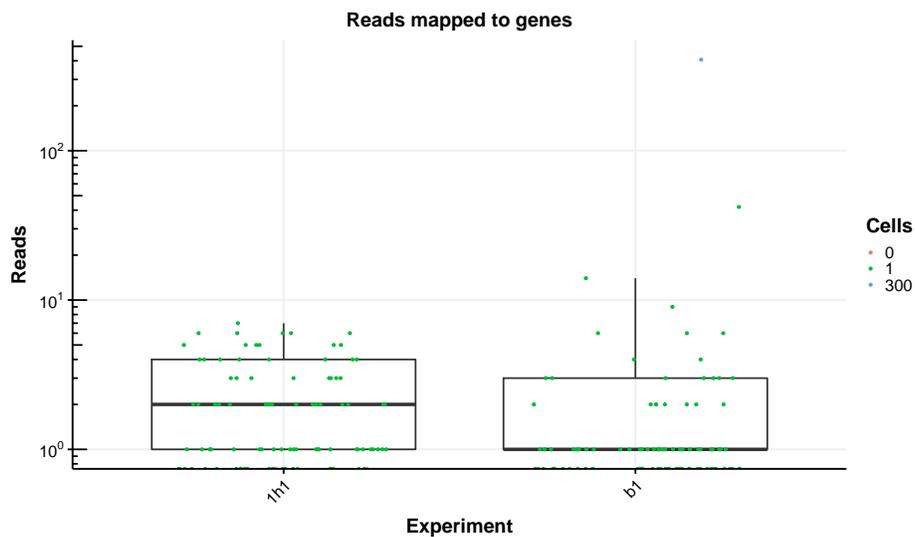
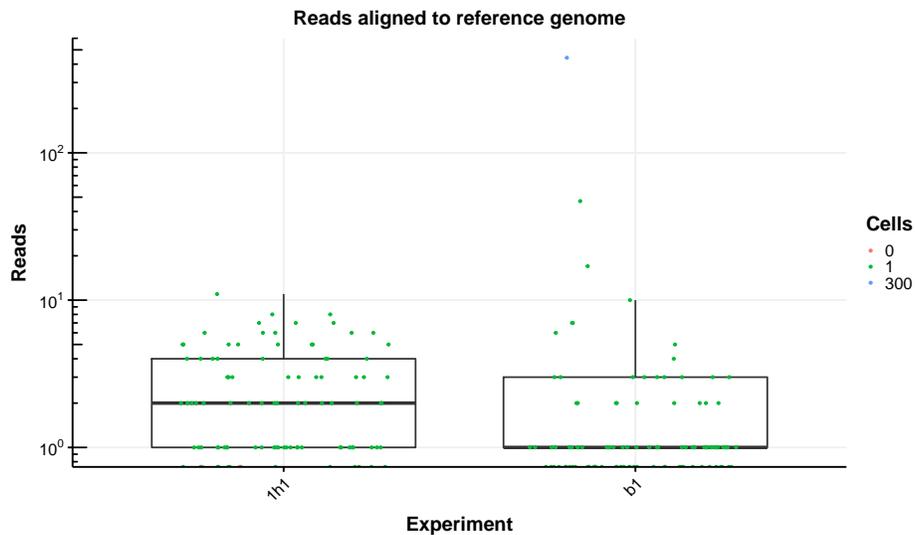
3.5 Visualization of QC metrics

The data quality diagnostic information are contained in the `colData` of the returned `SingleCellExperiment` object `sce`. They can be visualized using the `qcplots` function.

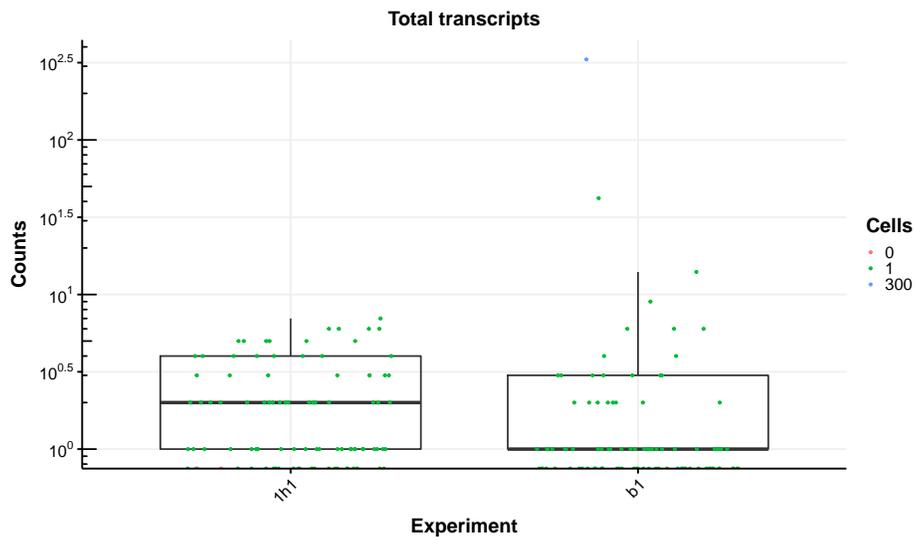
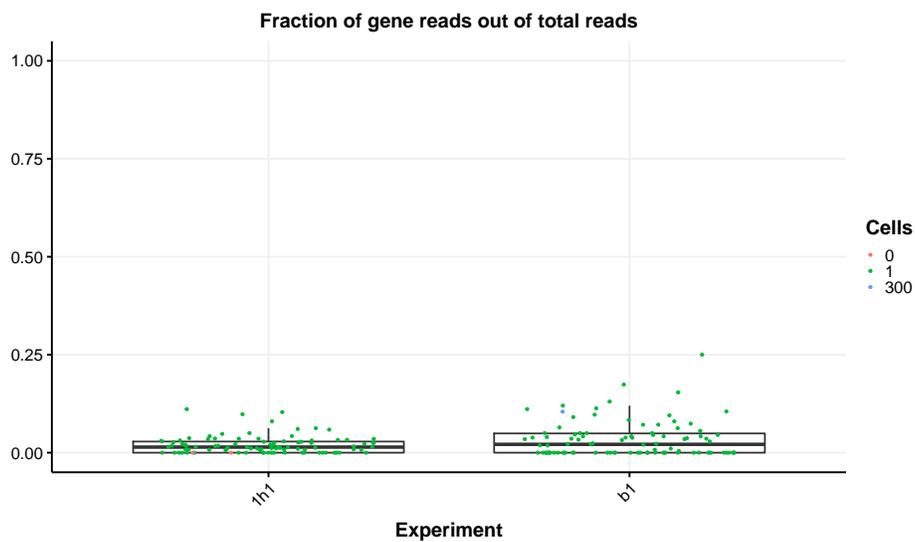
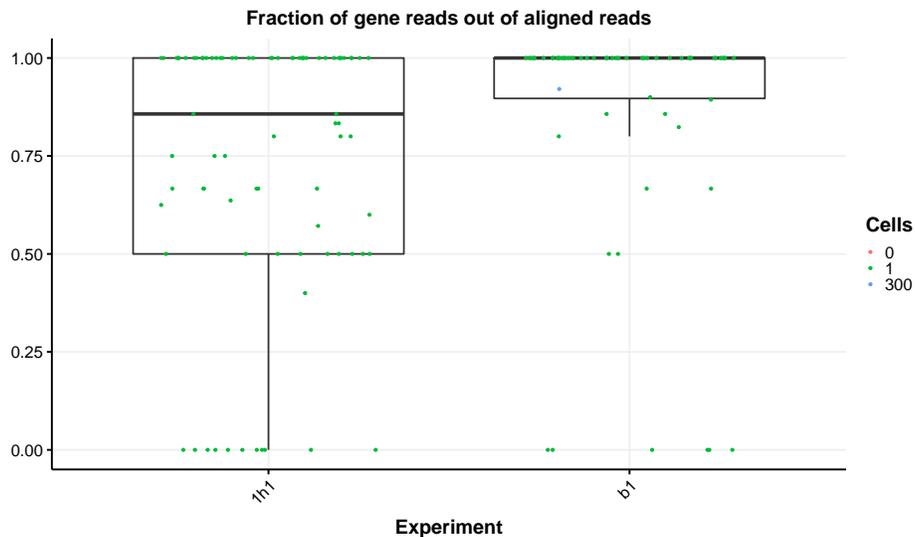
```
data(sceExample, package = "scruff")
qc <- qcplots(sceExample)
qc
```



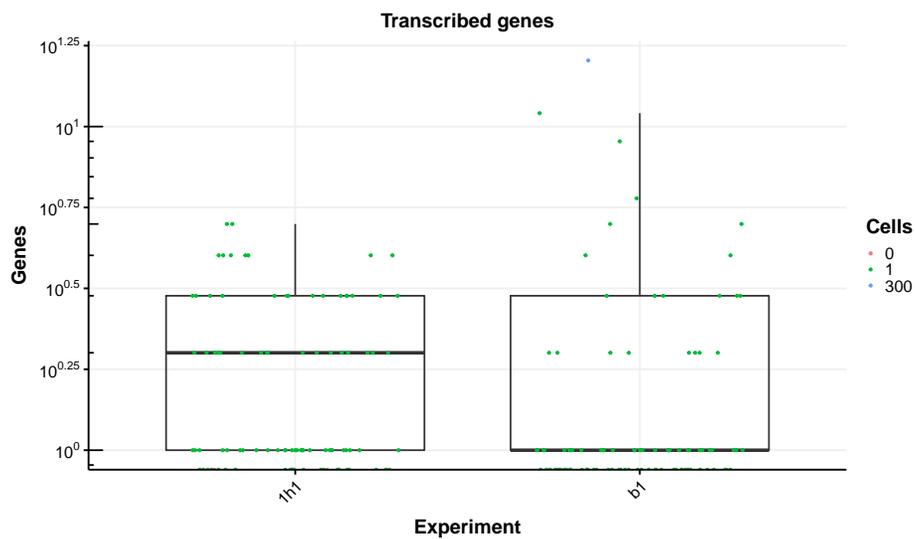
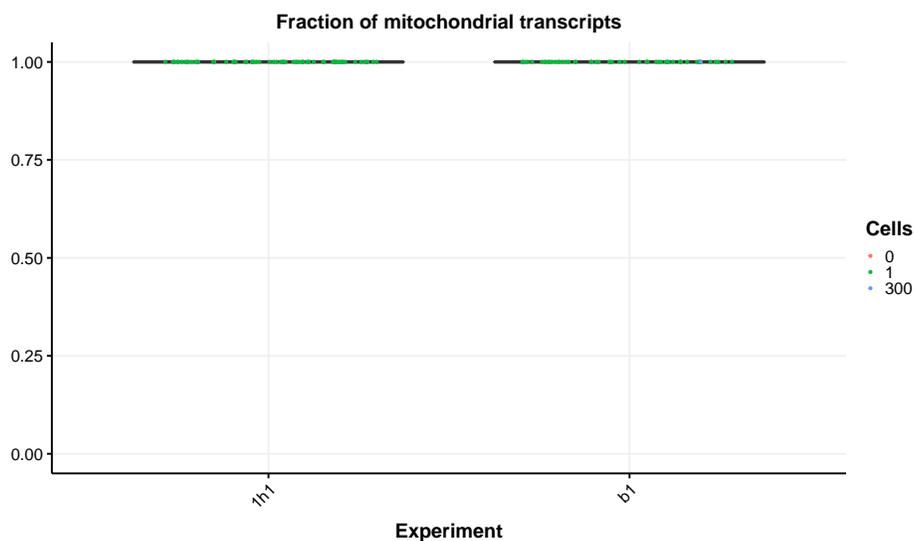
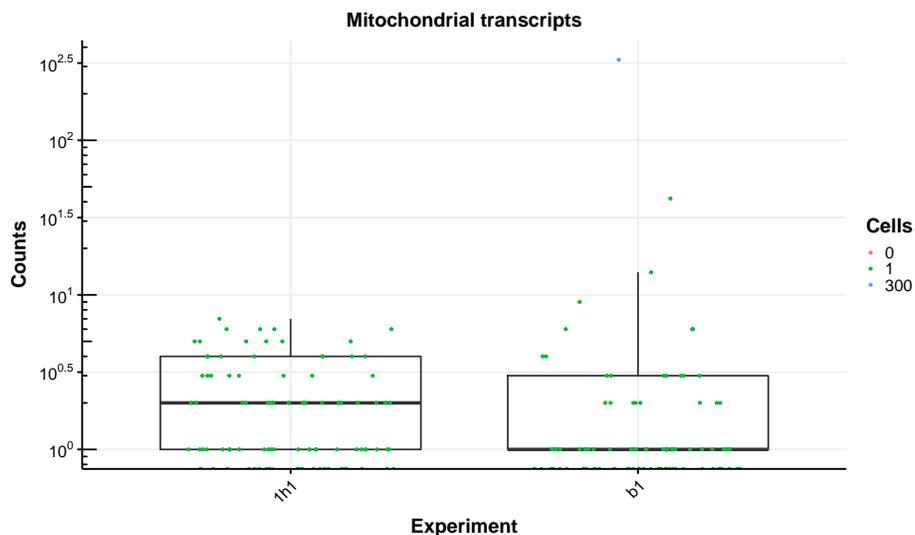
Process scRNA-Seq reads in *scruff*



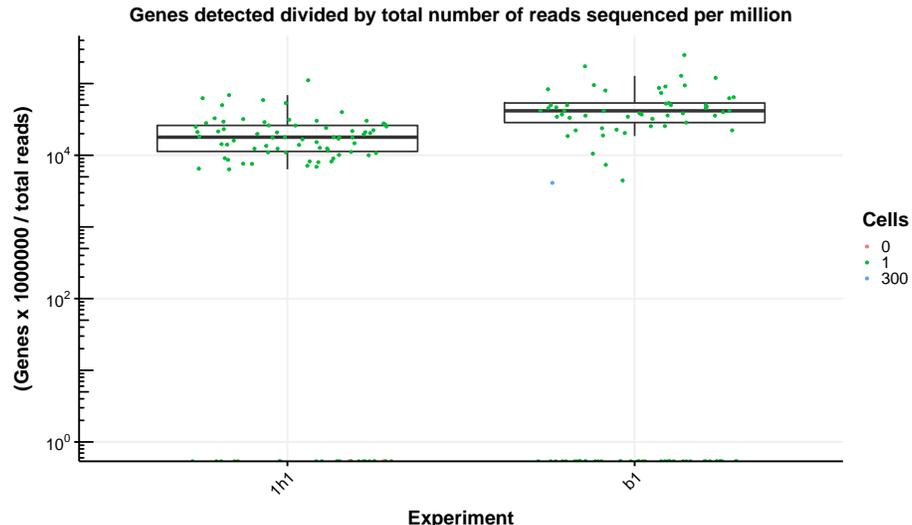
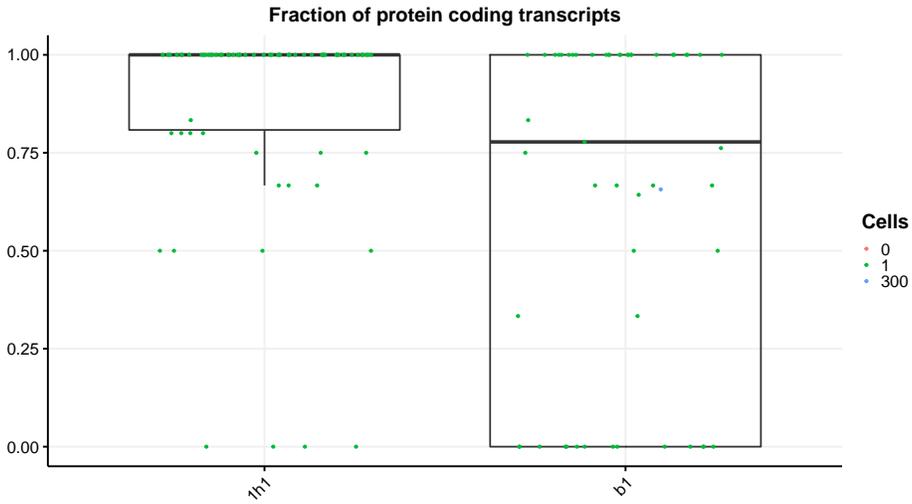
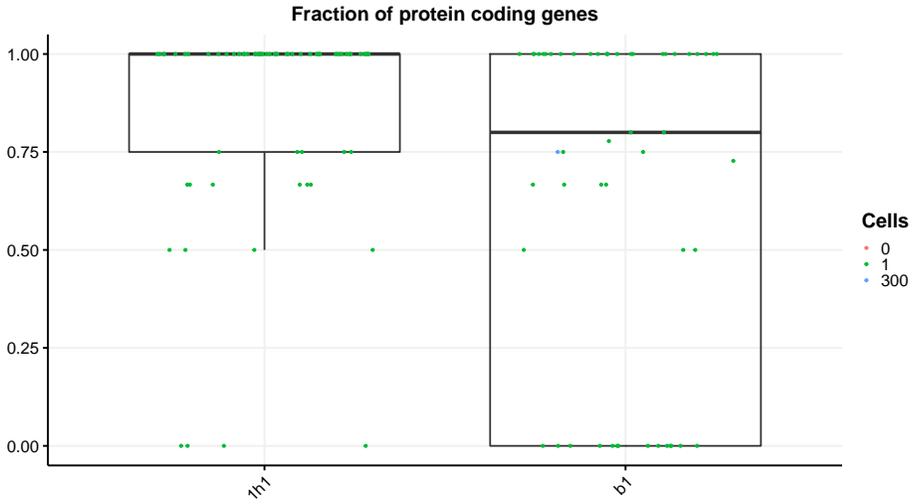
Process scRNA-Seq reads in *scruff*



Process scRNA-Seq reads in *scruff*



Process scRNA-Seq reads in *scruff*



Process scRNA-Seq reads in *scruff*

Each of these boxplots shows the distribution of a quality metric within and across experiments. Y axis shows the number or fraction of these metrics. Each point represents a unique cell barcode associated with the cells or bulk samples in the wells of the plate. These points are colored by the `cellPerWell` parameter which is the number of sorted cells in each well. For example, in experiment 1h1, cell barcodes 95 and 96 are empty well controls. In experiment b1, cell barcode 95 is a bulk sample containing 300 cells.

These quality metrics include total reads, aligned reads, reads mapped to genes, fraction of aligned reads, fraction of gene reads out of aligned reads, fraction of gene reads out of total reads, total transcripts, mitochondrial transcripts, fraction of mitochondrial transcripts, transcribed genes, fraction of protein coding genes and transcripts, and gene detection rate.

3.6 Visualization of Read mapping locations

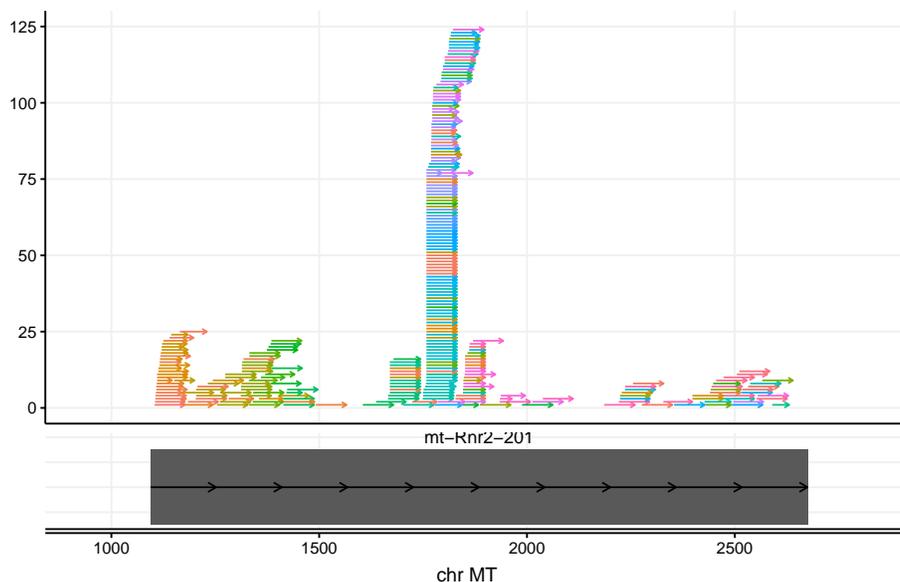
scruff package provides function to visualize read alignment locations at specified genomic coordinates on the reference genome. The following command visualize the reads mapped to gene `mt-Rnr2` for the bulk sample `vandenBrink_b1_cell_0095`. Reads are colored by their UMI tags.

```
# Visualize the reads mapped to gene "mt-Rnr2" in  
# cell "vandenBrink_b1_cell_0095".  
data(bamExample, package = "scruff")
```

```
gtfEG = refGenome::ensemblGenome(dirname(gtf))  
refGenome::read.gtf(gtfEG, filename = basename(gtf))
```

```
# gene mt-Rnr2 starts at 1094 and ends at 2675  
start <- 1094  
end <- 2675  
  
g1 <- rview(bamExample, chr = "MT", start = start, end = end)  
g2 <- gview(gtfEG, chr = "MT", start = start, end = end)  
g <- ggbio::tracks(g1, g2, heights = c(4,1), xlab = "chr MT")  
g
```

Process scRNA-Seq reads in *scruff*



This plot shows the read alignment information for gene mt-RNA2. The top panel shows the alignment and orientation of each read aligned to this gene. Y axis indicates the number of aligned reads at specific locations. Each arrow represents a read. The length of the arrow indicates the length of the read. The reads are colored by their UMI tags. In the bottom panel, each arrow represents an isoform of the gene. The isoforms are named by their transcript names. The grey boxes indicate the exonic regions of the transcript. X axis is the genomic coordinates on the corresponding chromosome (e.g. chr MT). It is consistent between the top and bottom panels if `legend` is set to `FALSE` for `rview` function.

4 10X BAM File Quality Assessment

The function `tenxBamqc` collects QC metrics from BAM files generated by 10X Cell Ranger pipeline. The collected QC metrics can be visualized by the function `qcplots`. The collected measurements include number of aligned reads and number of reads aligned to a gene for each valid cell barcode. Users can see the alignment quality for the filtered cells after plugging in the filtered barcode file from the Cell Ranger results. Here we show the visualization of an example BAM file from Cell Ranger output.

```
# The following example BAM file is the first 5000 BAM file records extracted
# from sample 01 of the 1.3 Million Brain Cells dataset from E18 Mice.
# (https://support.10xgenomics.com/single-cell-gene-expression/datasets/1.3.0/
# 1M_neurons)
# The BAM file for sample 01 is downloaded from here:
# http://sra-download.ncbi.nlm.nih.gov/srapub_files/
# SRR5167880_E18_20160930_Neurons_Sample_01.bam
# see details here:
# https://trace.ncbi.nlm.nih.gov/Traces/sra/?study=SRP096558
# and here:
# https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE93421
bamfile10x <- system.file("extdata",
  "SRR5167880_E18_20160930_Neurons_Sample_01_5000.bam",
```

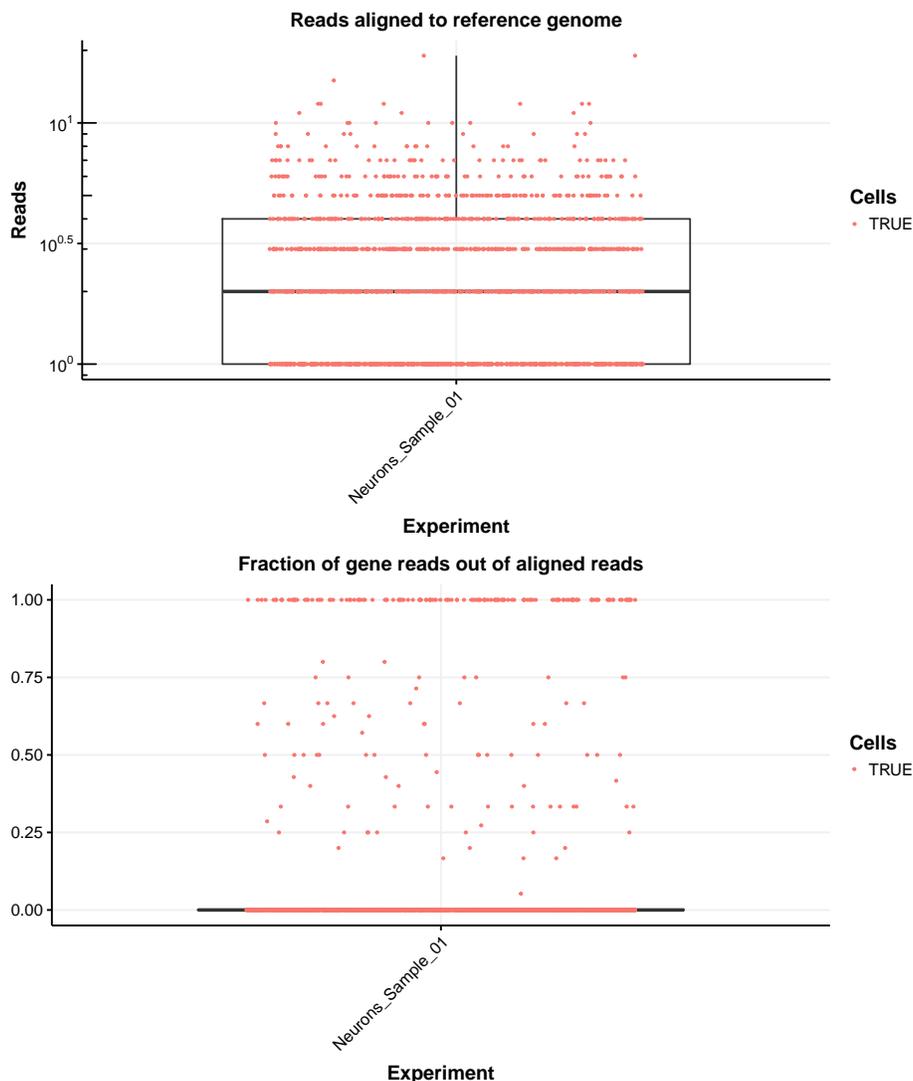
Process scRNA-Seq reads in *scruff*

```
package = "scruff")

# The filtered cell barcodes are generated using the following script:
# library(TENxBrainData)
# library(data.table)
# tenx <- TENxBrainData()
# # get filtered barcodes for sample 01
# filteredBcIndex <- tstrsplit(colData(tenx)[, "Barcode"], "-")[[2]] == 1
# filteredBc <- colData(tenx)[filteredBcIndex, ][["Barcode"]]
filteredBc <- system.file("extdata",
  "SRR5167880_E18_20160930_Neurons_Sample_01_filtered_barcode.tsv",
  package = "scruff")
# QC results are saved to current working directory
qcDt <- tenxBamqc(bam = bamfile10x,
  experiment = "Neurons_Sample_01",
  filter = filteredBc)
## tenxBamqc(bam = bamfile10x, experiment = "Neurons_Sample_01",
##   filter = filteredBc)
qcDt
## class: SingleCellExperiment
## dim: 1 1576
## metadata(0):
## assays(1): ''
## rownames: NULL
## rowData names(0):
## colnames(1576): AACCTGAGGCCGTT-Neurons_Sample_01
##   AACCTGAGTCCGGTC-Neurons_Sample_01 ...
##   TTTGTCATCGGTGTTA-Neurons_Sample_01
##   TTTGTCATCTGTATC-Neurons_Sample_01
## colData names(5): cell_barcode reads_mapped_to_genome
##   reads_mapped_to_genes experiment number_of_cells
## reducedDimNames(0):
## spikeNames(0):
```

```
g <- qcplots(qcDt)
g
```

Process scRNA-Seq reads in *scruff*



5 Session Information

```
## R version 3.5.1 Patched (2018-07-12 r74967)
## Platform: x86_64-pc-linux-gnu (64-bit)
## Running under: Ubuntu 16.04.5 LTS
##
## Matrix products: default
## BLAS: /home/biocbuild/bbs-3.8-bioc/R/lib/libRblas.so
## LAPACK: /home/biocbuild/bbs-3.8-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
```

Process scRNA-Seq reads in *scruff*

```
## [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] stats4      parallel  stats      graphics  grDevices  utils      datasets
## [8] methods    base
##
## other attached packages:
## [1] GenomicFeatures_1.34.0 AnnotationDbi_1.44.0 Biobase_2.42.0
## [4] GenomicRanges_1.34.0 GenomeInfoDb_1.18.0 IRanges_2.16.0
## [7] S4Vectors_0.20.0      BiocGenerics_0.28.0 scruff_1.0.0
## [10] BiocStyle_2.10.0
##
## loaded via a namespace (and not attached):
## [1] colorspace_1.3-2      hwriter_1.3.2
## [3] rprojroot_1.3-2      biovizBase_1.30.0
## [5] htmlTable_1.12        XVector_0.22.0
## [7] base64enc_0.1-3      dichromat_2.0-0
## [9] rstudioapi_0.8        bit64_0.9-7
## [11] splines_3.5.1         ggbio_1.30.0
## [13] doBy_4.6-2            knitr_1.20
## [15] Formula_1.2-3         Rsamtools_1.34.0
## [17] cluster_2.0.7-1      graph_1.60.0
## [19] BiocManager_1.30.3   compiler_3.5.1
## [21] httr_1.3.1            backports_1.1.2
## [23] assertthat_0.2.0     Matrix_1.2-14
## [25] lazyeval_0.2.1       acepack_1.4.1
## [27] htmltools_0.3.6      prettyunits_1.0.2
## [29] tools_3.5.1           bindr_0.2.2
## [31] gtable_0.2.0          glue_1.3.0
## [33] GenomeInfoDbData_1.2.0 reshape2_1.4.3
## [35] dplyr_0.7.7           ggthemes_4.0.1
## [37] ShortRead_1.40.0     Rcpp_0.12.19
## [39] Biostrings_2.50.0    rtracklayer_1.42.0
## [41] refGenome_1.7.3      xfun_0.4
## [43] stringr_1.3.1        ensemblDb_2.6.0
## [45] XML_3.98-1.16        zlibbioc_1.28.0
## [47] MASS_7.3-51          scales_1.0.0
## [49] BSgenome_1.50.0      VariantAnnotation_1.28.0
## [51] hms_0.4.2            ProtGenerics_1.14.0
## [53] SummarizedExperiment_1.12.0 RBGL_1.58.0
## [55] AnnotationFilter_1.6.0 RColorBrewer_1.1-2
## [57] SingleCellExperiment_1.4.0 yaml_2.2.0
## [59] curl_3.2              memoise_1.1.0
## [61] gridExtra_2.3         ggplot2_3.1.0
## [63] biomaRt_2.38.0       rpart_4.1-13
## [65] reshape_0.8.8        latticeExtra_0.6-28
## [67] stringi_1.2.4         RSQLite_2.1.1
## [69] checkmate_1.8.5      BiocParallel_1.16.0
## [71] Rsubread_1.32.0      rlang_0.3.0.1
```

Process scRNA-Seq reads in *scruff*

```
## [73] pkgconfig_2.0.2          matrixStats_0.54.0
## [75] bitops_1.0-6                evaluate_0.12
## [77] lattice_0.20-35             purrr_0.2.5
## [79] bindr_0.1.1                 labeling_0.3
## [81] GenomicAlignments_1.18.0    htmlwidgets_1.3
## [83] bit_1.1-14                  tidyselect_0.2.5
## [85] GGally_1.4.0                plyr_1.8.4
## [87] magrittr_1.5                 bookdown_0.7
## [89] R6_2.3.0                    Hmisc_4.1-1
## [91] DelayedArray_0.8.0          DBI_1.0.0
## [93] pillar_1.3.0                foreign_0.8-71
## [95] withr_2.1.2                 survival_2.43-1
## [97] RCurl_1.95-4.11             nnet_7.3-12
## [99] tibble_1.4.2                crayon_1.3.4
## [101] OrganismDbi_1.24.0          rmarkdown_1.10
## [103] progress_1.2.0              grid_3.5.1
## [105] data.table_1.11.8           blob_1.1.1
## [107] digest_0.6.18               munsell_0.5.0
```