

Package ‘scraper’

May 2, 2026

Version 1.7.0

Date 2026-04-24

Title Bindings to C++ Libraries for Single-Cell Analysis

Description Implements R bindings to C++ code for analyzing single-cell (expression) data, mostly from various libscrans libraries. Each function performs an individual step in the single-cell analysis workflow, ranging from quality control to clustering and marker detection. Additional wrappers are provided for easy construction of end-to-end workflows involving Bioconductor objects like SingleCellExperiments.

License MIT + file LICENSE

Imports methods, Rcpp, beachmat (>= 2.25.1), S4Vectors, SparseArray, DelayedArray, BiocNeighbors (>= 1.99.0), parallel

Suggests testthat, knitr, rmarkdown, BiocStyle, Matrix, IRanges, SummarizedExperiment, SingleCellExperiment, scRNAseq, org.Mm.eg.db, scater, igraph

LinkingTo Rcpp, assorthead (>= 1.5.16), beachmat, BiocNeighbors, Rigraphlib

biocViews Normalization, RNASeq, Software, GeneExpression, Transcriptomics, SingleCell, BatchEffect, QualityControl, DifferentialExpression, FeatureExtraction, PrincipalComponent, Clustering

SystemRequirements C++17, GNU make

URL <https://github.com/libscrans/scraper>

BugReports <https://github.com/libscrans/scraper/issues>

VignetteBuilder knitr

Encoding UTF-8

RoxygenNote 7.3.3

git_url <https://git.bioconductor.org/packages/scraper>

git_branch devel

git_last_commit 8005249

git_last_commit_date 2026-04-28

Repository Bioconductor 3.24

Date/Publication 2026-05-01

Author Aaron Lun [cre, aut]

Maintainer Aaron Lun <infinite.monkeys.with.keyboards@gmail.com>

Contents

scrapper-package	3
adt_quality_control	3
aggregateAcrossCells	6
aggregateAcrossCells.se	7
aggregateAcrossGenes	9
aggregateAcrossGenes.se	11
analyze	12
analyze.se	17
buildSnnGraph	22
centerSizeFactors	24
chooseHighlyVariableGenes	25
choosePseudoCount	26
chooseRnaHvgs.se	27
clusterGraph	29
clusterGraph.se	31
clusterKmeans	32
clusterKmeans.se	35
combineFactors	36
computeBlockWeights	37
computeClrm1Factors	38
convertAnalyzeResults	39
correctMnn	40
correctMnn.se	42
countGroupsByBlock	43
crispr_quality_control	44
fitVarianceTrend	46
getTestData.se	48
LogNormalizedMatrix	49
modelGeneVariances	50
normalizeAdtCounts.se	53
normalizeCounts	54
normalizeCrisprCounts.se	56
normalizeRnaCounts.se	57
quickAdtQc.se	58
quickCrisprQc.se	60
quickRnaQc.se	61
reportGroupMarkerStatistics	63
rna_quality_control	64
runAllNeighborSteps	67
runAllNeighborSteps.se	68
runPca	70
runPca.se	72
runTsne	74
runTsne.se	76
runUmap	77
runUmap.se	81
sanitizeSizeFactors	82
scaleByNeighbors	83
scaleByNeighbors.se	85
scoreGeneSet	86

scoreGeneSet.se	88
scoreMarkers	89
scoreMarkers.se	94
subsampleByNeighbors	96
summarizeEffects	97
testEnrichment	99

Index**101**

scrapper-package	<i>scrapper: Bindings to C++ Libraries for Single-Cell Analysis</i>
------------------	---

Description

Implements R bindings to C++ code for analyzing single-cell (expression) data, mostly from various libscrans libraries. Each function performs an individual step in the single-cell analysis workflow, ranging from quality control to clustering and marker detection. Additional wrappers are provided for easy construction of end-to-end workflows involving Bioconductor objects like SingleCellExperiments.

Author(s)

Maintainer: Aaron Lun <infinite.monkeys.with.keyboards@gmail.com>

See Also

Useful links:

- <https://github.com/libscrans/scrapper>
- Report bugs at <https://github.com/libscrans/scrapper/issues>

adt_quality_control	<i>Quality control for ADT count data</i>
---------------------	---

Description

Compute per-cell QC metrics from an initialized matrix of ADT counts, and use the metrics to suggest filter thresholds to retain high-quality cells.

Usage

```
computeAdtQcMetrics(x, subsets, num.threads = 1)

suggestAdtQcThresholds(
  metrics,
  block = NULL,
  min.detected.drop = 0.1,
  num.mads = 3
)

filterAdtQcMetrics(thresholds, metrics, block = NULL)
```

Arguments

x	A matrix-like object where rows are ADTs and columns are cells. Values are expected to be counts.
subsets	Named list of vectors specifying tag subsets of interest, typically control tags like IgGs. Each vector may be logical (whether to keep each row), integer (row indices) or character (row names). For character vectors, strings not present in <code>rownames(x)</code> are ignored.
num.threads	Integer scalar specifying the number of threads to use.
metrics	DataFrame of per-cell QC metrics. This should have the same structure as the return value of <code>computeAdtQcMetrics</code> .
block	Factor specifying the block of origin (e.g., batch, sample) for each cell in <code>metrics</code> . Alternatively NULL if all cells are from the same block. For <code>filterAdtQcMetrics</code> , a blocking factor should be provided if <code>block</code> was used to construct thresholds.
min.detected.drop	Minimum drop in the number of detected features from the median, in order to consider a cell to be of low quality.
num.mads	Number of median from the median, to define the threshold for outliers in each metric.
thresholds	List with the same structure as produced by <code>suggestAdtQcThresholds</code> .

Value

For `computeAdtQcMetrics`, a [DataFrame](#) is returned with one row per cell in `x`. This contains the following columns:

- `sum`, a numeric vector containing the total ADT count for each cell. In theory, this represents the efficiency of library preparation and sequencing. Compared to RNA, the sum is less useful as a QC metric for ADT data as it is strongly influenced by biological variation in the abundance of the targeted features. Nonetheless, we compute it for diagnostic purposes.
- `detected`, an integer vector containing the number of detected tags per cell. Even though ADTs are typically used in situations where few features are highly abundant (e.g., cell type-specific markers), we still expect detectable coverage of most features due to ambient contamination, non-specific binding or some background expression. Low numbers of detected tags indicates that library preparation or sequencing depth was suboptimal.
- `subsets`, a nested [DataFrame](#) where each column corresponds to a control subset and is a numeric vector containing the total count in that subset. The exact interpretation depends on the nature of the feature subset but the most common use case involves isotype control (IgG) features. IgG antibodies should not bind to anything so a high subset sum suggests that non-specific binding is a problem, e.g., due to antibody conjugates. (Unlike RNA quality control, we do not use proportions here as it is entirely possible for a cell to have low counts for other tags due to the absence of their targeted features; this would result in a high proportion even if the cell has a "normal" level of non-specific binding.)

Each vector is of length equal to the number of cells.

For `suggestAdtQcThresholds`, a named list is returned:

- If `block=NULL`, the list contains:

- detected, a numeric scalar containing the lower bound on the number of detected tags. This is defined as the lower of (i) num.mads MADs below the median for the log-transformed values across all cells, and (ii) the product of 1 - min.detected.drop and the median across all cells. The latter avoids overly aggressive filtering when the MAD is zero.
 - subsets, a numeric vector containing the upper bound on the sum of counts in each control subset. This is defined as num.mads MADs above the median of the log-transformed metrics across all cells.
- Otherwise, if block is supplied, the list contains:
 - detected, a numeric vector containing the lower bound on the number of detected tags for each blocking level. Here, the threshold is computed independently for each block, using the same method as the unblocked case.
 - subsets, a list of numeric vectors containing the upper bound on the sum of counts in each control subset for each blocking level. Here, the threshold is computed independently for each block, using the same method as the unblocked case.
 - block.ids, a vector containing the identities of the unique blocks.

Each vector is of length equal to the number of levels in block and is named accordingly.

For filterAdtQcMetrics, a logical vector of length ncol(x) is returned indicating which cells are of high quality. High-quality cells are defined as those with numbers of detected tags above the detected threshold and control subset sums below the subsets threshold.

Author(s)

Aaron Lun

See Also

The compute_adt_qc_metrics, compute_adt_qc_filters and compute_adt_qc_filters_blocked functions in https://libscrان.github.io/scrان_qc/.

[quickAdtQc.se](#), to run all of the ADT-related QC functions on a [SummarizedExperiment](#).

Examples

```
# Mocking a matrix:
library(Matrix)
x <- round(abs(rsparsematrix(1000, 100, 0.1) * 100))

# Mocking up a control set.
sub <- list(IgG=rbinom(nrow(x), 1, 0.1) > 0)

qc <- computeAdtQcMetrics(x, sub)
qc

filt <- suggestAdtQcThresholds(qc)
str(filt)

keep <- filterAdtQcMetrics(filt, qc)
summary(keep)
```

aggregateAcrossCells *Aggregate expression across cells*

Description

Aggregate expression values across cells based on one or more grouping factors. This is usually applied to a count matrix to create pseudo-bulk profiles for each cluster/sample combination. These profiles can then be used as if they were counts from bulk data, e.g., for differential analyses with **edgeR**.

Usage

```
aggregateAcrossCells(
  x,
  factors,
  compute.sum = TRUE,
  compute.detected = TRUE,
  compute.median = TRUE,
  num.threads = 1
)
```

Arguments

x	A matrix-like object where rows correspond to genes or genomic features and columns correspond to cells. Values are typically expected to be counts.
factors	A list or data frame (or their equivalents from S4Vectors) containing one or more grouping factors, see combineFactors .
compute.sum	Boolean indicating whether to compute the sum in each group.
compute.detected	Boolean indicating whether to compute the number of cells with detected expression in each group.
compute.median	Boolean indicating whether to compute the median in each group.
num.threads	Integer specifying the number of threads to be used for aggregation.

Value

A list containing:

- `sums`, a numeric matrix where each row corresponds to a gene and each column corresponds to a unique combination of levels from factors. Each entry contains the summed expression across all cells with that combination. Omitted if `compute.sum=FALSE`.
- `detected`, an integer matrix where each row corresponds to a gene and each column corresponds to a unique combination of levels from factors. Each entry contains the number of cells with detected (i.e., positive) expression in that combination. Omitted if `compute.detected=FALSE`.
- `medians`, a numeric matrix where each row corresponds to a gene and each column corresponds to a unique combination of levels from factors. Each entry contains the median expression across all cells with that combination. Omitted if `compute.median=FALSE`.
- `combinations`, a [DataFrame](#) containing the unique combination of levels from factors. Rows correspond to columns of `sums` and `detected`, while columns correspond to the factors in factors.

- counts, the number of cells associated with each combination. Each entry corresponds to a row of combinations.
- index, an integer vector of length equal to the number of cells in x. This specifies the combination in combinations to which each cell was assigned.

Author(s)

Aaron Lun

See Also

The `aggregate_across_cells` function in https://libscran.github.io/scran_aggregate/, `aggregateAcrossCells.se`, to perform aggregation on a `SummarizedExperiment`.
`aggregateAcrossGenes`, to aggregate expression values across gene sets.

Examples

```
# Mocking a matrix:
library(Matrix)
x <- round(abs(rsparsematrix(1000, 100, 0.1) * 100))

# Simple aggregation:
clusters <- sample(LETTERS, 100, replace=TRUE)
agg <- aggregateAcrossCells(x, list(cluster=clusters))
str(agg)

# Multi-factor aggregation
samples <- sample(1:5, 100, replace=TRUE)
agg2 <- aggregateAcrossCells(x, list(cluster=clusters, sample=samples))
str(agg2)
```

aggregateAcrossCells.se

Aggregate expression across cells in a SummarizedExperiment

Description

Aggregate expression values across groups of cells for each gene, by calling `aggregateAcrossCells` on an assay in a `SummarizedExperiment`.

Usage

```
aggregateAcrossCells.se(  
  x,  
  factors,  
  num.threads = 1,  
  more.aggr.args = list(),  
  assay.type = "counts",  
  output.prefix = "factor.",  
  counts.name = "counts",  
  meta.name = "aggregated",
```

```

include.coldata = TRUE,
more.coldata.args = list(),
altexps = NULL,
copy.altexps = FALSE
)

```

```
aggregateColData(coldata, index, number, only.atomic = TRUE, placeholder = NA)
```

Arguments

x	A SummarizedExperiment object or one of its subclasses. Rows correspond to genes and columns correspond to cells.
factors	List or data frame (or their equivalents from S4Vectors) containing grouping factors, see aggregateAcrossCells for more details. Alternatively, an atomic vector or factor representing a single variable.
num.threads	Number of threads, passed to aggregateAcrossCells .
more.aggr.args	Named list of additional arguments to pass to aggregateAcrossCells .
assay.type	Integer or string specifying the assay of x to be aggregated.
output.prefix	String specifying a prefix to add to the names of the <code>link[SummarizedExperiment]{colData}</code> columns storing the factor combinations. If NULL, no prefix is added.
counts.name	String specifying the name of the <code>colData</code> column in which to store the cell count for each factor combination. If NULL, the cell counts are not reported.
meta.name	String specifying the name of the <code>metadata</code> entry in which to store additional outputs like the combination indices. If NULL, additional outputs are not reported.
include.coldata	Logical scalar indicating whether to add the aggregated <code>colData</code> from x to the output.
more.coldata.args	Named list of additional arguments to pass to <code>aggregateColData</code> . Only relevant if <code>include.coldata=TRUE</code> .
altexps	Unnamed integer or character vector specifying the indices/names of alternative experiments of x to aggregate. The aggregated assay from each alternative experiment is determined by <code>assay.type</code> . Alternatively, this may be a named integer or character vector. Each name specifies an alternative experiment while each value is the index/name of the assay to aggregate from that experiment. Only relevant if x is a SingleCellExperiment .
copy.altexps	Logical scalar indicating whether to copy the <code>colData</code> and <code>metadata</code> of the output <code>SingleCellExperiment</code> into each of its alternative experiments.
coldata	DataFrame of column data, containing one row for each cell.
index	Integer vector containing the index of the factor combination to which each cell in <code>coldata</code> was assigned.
number	Integer specifying the total number of unique factor combinations. All elements of <code>index</code> should be less than or equal to <code>number</code> .
only.atomic	Logical scalar specifying whether to skip non-atomic, non-factor columns.
placeholder	Placeholder value to store in the output column when a factor combination does not have a single unique value.

Value

For `aggregateAcrossCells.se`, a `SummarizedExperiment` is returned where each column corresponds to a factor combination. Each row corresponds to a gene in `x` and the `rowData` is taken from `x`. The assays contain the sum of counts ("sums") and the number of detected cells ("detected") in each combination for each gene. The `colData` contains:

- The factor combinations, with column names prefixed by `output.prefix`.
- The cell count for each combination, named by `counts.name`.
- Additional `colData` from `x` if `include.coldata=TRUE`. This is aggregated with `aggregateColData` on the combination indices.

The metadata contains a list named as `meta.name`, containing a `index` integer vector of length equal to the number of cells in `x`. Each entry of this vector is an index of the factor combination (i.e., column of the output object) to which the corresponding cell was assigned.

If `altexps` is specified, a `SingleCellExperiment` is returned instead. The same aggregation for the main experiment is applied to each alternative experiment. If `copy.altexps=TRUE`, the `colData` for each alternative experiment will contain a copy of the factor combinations and cell counts, and the metadata will contain a copy of the index vector.

For `aggregateColData`, a `DataFrame` is returned with number of rows equal to `number`. Each atomic or factor column in `coldata` is represented by a column in the output `DataFrame`. In each column, the `j`-th entry is equal to the unique value of all rows where `index == j`, or placeholder if there is not exactly one unique value. If `only.atomic=FALSE`, any non-atomic/non-factor columns of `coldata` are represented in the output `DataFrame` by a vector of placeholder values. If `only.atomic=TRUE`, any non-atomic/non-factor columns of `coldata` are skipped.

Author(s)

Aaron Lun

Examples

```
library(SingleCellExperiment)
sce <- getTestRnaData.se("start")
aggr <- aggregateAcrossCells.se(sce, sce$level1class)
head(assay(aggr))
colData(aggr)

# We can also aggregate within alternative experiments.
aggr2 <- aggregateAcrossCells.se(sce, sce$level1class, altexps="ERCC")
head(assay(altExp(aggr2, "ERCC")))
```

`aggregateAcrossGenes` *Aggregate expression across genes*

Description

Aggregate expression values across genes, potentially with weights. This is typically used to summarize expression values for gene sets into a single per-cell score.

Usage

```
aggregateAcrossGenes(x, sets, average = FALSE, convert = TRUE, num.threads = 1)
```

Arguments

<code>x</code>	A matrix-like object where rows correspond to genes or genomic features and columns correspond to cells. Values are usually normalized expression values, possibly log-transformed depending on the application.
<code>sets</code>	List of vectors where each entry corresponds to a gene set. Each entry may be an integer vector of row indices, a logical vector of length equal to the number of rows, or a character vector of row names. For integer and character vectors, duplicate elements are ignored. For character vectors, any strings not present in <code>rownames(x)</code> are ignored. Alternatively, each entry may be a list of two vectors. The first vector should be either integer (row indices) or character (row names), specifying the genes in the set. The second vector should be numeric and of the same length as the first vector, specifying the weight associated with each gene. If duplicate genes are present, only the first occurrence is used. If the first vector contains gene names not present in <code>x</code> , those genes are ignored.
<code>average</code>	Logical scalar indicating whether to compute the average rather than the sum.
<code>convert</code>	Logical scalar indicating whether to convert gene identities to non-duplicate row indices in each entry of <code>sets</code> . Can be set to <code>FALSE</code> for greater efficiency if the <code>sets</code> already contains non-duplicated integer vectors.
<code>num.threads</code>	Integer specifying the number of threads to be used for aggregation.

Value

A list of length equal to that of `sets`. Each entry is a numeric vector of length equal to the number of columns in `x`, containing the (weighted) sum/mean of expression values for the corresponding set across all cells.

Author(s)

Aaron Lun

See Also

The `aggregate_across_genes` function in https://libscran.github.io/scran_aggregate/. [aggregateAcrossGenes](#), to perform aggregation on a [SummarizedExperiment](#). [aggregateAcrossCells](#), to aggregate expression values across groups of cells.

Examples

```
# Mocking a matrix:
library(Matrix)
x <- round(abs(rsparsmatrix(1000, 100, 0.1) * 100))

# Unweighted aggregation:
sets <- list(
  foo = sample(nrow(x), 20),
  bar = sample(nrow(x), 10)
)
```

```

agg <- aggregateAcrossGenes(x, sets)
str(agg)

# Weighted aggregation:
sets <- list(
  foo = list(sample(nrow(x), 20), runif(20)),
  bar = list(sample(nrow(x), 10), runif(10))
)
agg2 <- aggregateAcrossGenes(x, sets, average = TRUE)
str(agg2)

```

```
aggregateAcrossGenes.se
```

Aggregate expression across gene sets in a SummarizedExperiment

Description

Aggregate expression values across sets of genes for each cell, by calling `aggregateAcrossGenes` on an assay in a `SummarizedExperiment`.

Usage

```

aggregateAcrossGenes.se(
  x,
  sets,
  num.threads = 1,
  more.aggr.args = list(),
  assay.type = "logcounts",
  output.name = NULL
)

```

Arguments

<code>x</code>	A <code>SummarizedExperiment</code> object or one of its subclasses. Rows correspond to genes and columns correspond to cells.
<code>sets</code>	List of gene sets, see <code>aggregateAcrossGenes</code> for more details. Alternatively, <code>sets</code> may be a <code>List</code> subclass, in which case the <code>mcols</code> are used as the <code>rowData</code> of the output object. Weighted gene sets can be represented by a list of <code>DataFrames</code> or a <code>DataFrameList</code> , where each <code>DataFrame</code> contains two columns, i.e., the gene identities and the associated weights.
<code>num.threads</code>	Number of threads, passed to <code>aggregateAcrossGenes</code> .
<code>more.aggr.args</code>	Named list of additional arguments to pass to <code>aggregateAcrossGenes</code> .
<code>assay.type</code>	Integer or string specifying the assay of <code>x</code> to be aggregated.
<code>output.name</code>	String specifying the assay name in the output object. Defaults to <code>assay.type</code> if it is a string, otherwise "aggregated".

Value

A `SummarizedExperiment` with number of rows equal to the number of gene sets. The lone assay contains the aggregated values for each gene set for all cells. The `colData` is the same as that of `x`.

Author(s)

Aaron Lun

Examples

```
library(SingleCellExperiment)
sce <- getTestRnaData.se("norm")

library(org.Mm.eg.db)
set.df <- select(
  org.Mm.eg.db,
  keytype="GO",
  keys=c(
    "GO:0048709", # oligodendrocyte differentiation
    "GO:0048699", # neuron development
    "GO:0048143" # astrocyte activation
  ),
  columns="SYMBOL"
)
sets <- splitAsList(set.df$SYMBOL, set.df$GO)

aggregated <- aggregateAcrossGenes.se(sce, sets)
aggregated
assay(aggregated)[,1:10]
```

analyze

Analyze single-cell data

Description

Execute a simple single-cell analysis pipeline, starting from a count matrix and ending with clusters, visualizations and markers. This also supports integration of multiple modalities and correction of batch effects.

Usage

```
analyze(
  rna.x,
  adt.x = NULL,
  crispr.x = NULL,
  block = NULL,
  rna.subsets = list(),
  adt.subsets = list(),
  suggestRnaQcThresholds.args = list(),
  suggestAdtQcThresholds.args = list(),
  suggestCrisprQcThresholds.args = list(),
  filter.cells = TRUE,
  centerSizeFactors.args = list(),
  computeClrm1Factors.args = list(),
  normalizeCounts.args = list(),
  modelGeneVariances.args = list(),
```

```

chooseHighlyVariableGenes.args = list(),
runPca.args = list(),
use.rna.pcs = TRUE,
use.adt.pcs = TRUE,
use.crispr.pcs = TRUE,
scaleByNeighbors.args = list(),
correctMnn.args = list(),
runUmap.args = list(),
runTsne.args = list(),
buildSnnGraph.args = list(),
clusterGraph.args = list(),
runAllNeighborSteps.args = list(),
kmeans.clusters = NULL,
clusterKmeans.args = list(),
clusters.for.markers = c("graph", "kmeans"),
scoreMarkers.args = list(),
BNPARAM = AnnoyParam(),
rna.assay = 1L,
adt.assay = 1L,
crispr.assay = 1L,
num.threads = 3L
)

```

Arguments

<code>rna.x</code>	Matrix-like object containing RNA counts. This should have the same number of columns as the other <code>*.x</code> arguments. Alternatively, a SummarizedExperiment instance containing such a matrix in its <code>rna.assay</code> . Alternatively NULL, if no RNA counts are available.
<code>adt.x</code>	Matrix-like object containing ADT counts. This should have the same number of columns as the other <code>*.x</code> arguments. Alternatively, a SummarizedExperiment instance containing such a matrix in its <code>adt.assay</code> . Alternatively NULL, if no ADT counts are available.
<code>crispr.x</code>	Matrix-like object containing ADT counts. This should have the same number of columns as the other <code>*.x</code> arguments. Alternatively, a SummarizedExperiment instance containing such a matrix in its <code>crispr.assay</code> . Alternatively NULL, if no ADT counts are available.
<code>block</code>	Factor specifying the block of origin (e.g., batch, sample) for each cell in the <code>*_x</code> matrices. Alternatively NULL, if all cells are from the same block.
<code>rna.subsets</code>	Gene subsets for quality control, typically used for mitochondrial genes. See the subsets arguments in computeRnaQcMetrics for details.
<code>adt.subsets</code>	ADT subsets for quality control, typically used for IgG controls. See the subsets arguments in computeAdtQcMetrics for details.
<code>suggestRnaQcThresholds.args</code>	Named list of arguments to pass to suggestRnaQcThresholds .
<code>suggestAdtQcThresholds.args</code>	Named list of arguments to pass to suggestAdtQcThresholds .

<code>suggestCrisprQcThresholds.args</code>	Named list of arguments to pass to suggestCrisprQcThresholds .
<code>filter.cells</code>	Logical scalar indicating whether to filter the count matrices to only retain high-quality cells in all modalities. If FALSE, QC metrics and thresholds are still computed but are not used to filter the count matrices.
<code>centerSizeFactors.args</code>	Named list of arguments to pass to centerSizeFactors .
<code>computeClrm1Factors.args</code>	Named list of arguments to pass to computeClrm1Factors . Only used if <code>adt.x</code> is provided.
<code>normalizeCounts.args</code>	Named list of arguments to pass to normalizeCounts .
<code>modelGeneVariances.args</code>	Named list of arguments to pass to modelGeneVariances . Only used if <code>rna.x</code> is provided.
<code>chooseHighlyVariableGenes.args</code>	Named list of arguments to pass to chooseHighlyVariableGenes . Only used if <code>rna.x</code> is provided.
<code>runPca.args</code>	Named list of arguments to pass to runPca .
<code>use.rna.pcs</code>	Logical scalar indicating whether to use the RNA-derived PCs for downstream steps (i.e., clustering, visualization). Only used if <code>rna.x</code> is provided.
<code>use.adt.pcs</code>	Logical scalar indicating whether to use the ADT-derived PCs for downstream steps (i.e., clustering, visualization). Only used if <code>adt.x</code> is provided.
<code>use.crispr.pcs</code>	Logical scalar indicating whether to use the CRISPR-derived PCs for downstream steps (i.e., clustering, visualization). Only used if <code>crispr.x</code> is provided.
<code>scaleByNeighbors.args</code>	Named list of arguments to pass to scaleByNeighbors . Only used if multiple modalities are available and their corresponding <code>use.*.pcs</code> arguments are TRUE.
<code>correctMnn.args</code>	Named list of arguments to pass to correctMnn . Only used if <code>block</code> is supplied.
<code>runUmap.args</code>	Named list of arguments to pass to runUmap . If NULL, UMAP is not performed.
<code>runTsne.args</code>	Named list of arguments to pass to runTsne . If NULL, t-SNE is not performed.
<code>buildSnnGraph.args</code>	Named list of arguments to pass to buildSnnGraph . Ignored if <code>clusterGraph.args</code> = NULL.
<code>clusterGraph.args</code>	Named list of arguments to pass to clusterGraph . If NULL, graph-based clustering is not performed.
<code>runAllNeighborSteps.args</code>	Named list of arguments to pass to runAllNeighborSteps .
<code>kmeans.clusters</code>	Integer scalar specifying the number of clusters to use in k-means clustering. If NULL, k-means clustering is not performed.
<code>clusterKmeans.args</code>	Named list of arguments to pass to clusterKmeans . Ignored if <code>kmeans.clusters</code> = NULL.
<code>clusters.for.markers</code>	Character vector of clustering algorithms (either "graph" or "kmeans", specifying the clustering to be used for marker detection. The first available clustering will be chosen.

<code>scoreMarkers.args</code>	Named list of arguments to pass to <code>scoreMarkers</code> . Ignored if no suitable clusterings are available.
<code>BNPARAM</code>	A <code>BiocNeighborParam</code> instance specifying the nearest-neighbor search algorithm to use.
<code>rna.assay</code>	Integer scalar or string specifying the assay to use if <code>rna.x</code> is a <code>SummarizedExperiment</code> .
<code>adt.assay</code>	Integer scalar or string specifying the assay to use if <code>adt.x</code> is a <code>SummarizedExperiment</code> .
<code>crispr.assay</code>	Integer scalar or string specifying the assay to use if <code>crispr.x</code> is a <code>SummarizedExperiment</code> .
<code>num.threads</code>	Integer scalar specifying the number of threads to use in each step.

Value

List containing the results of the entire analysis:

- `rna.qc.metrics`: Results of `computeRnaQcMetrics`. If RNA data is not available, this is set to NULL instead.
- `rna.qc.thresholds`: Results of `suggestRnaQcThresholds`. If RNA data is not available, this is set to NULL instead.
- `rna.qc.filter`: Results of `filterRnaQcMetrics`. If RNA data is not available, this is set to NULL instead.
- `adt.qc.metrics`: Results of `computeAdtQcMetrics`. If ADT data is not available, this is set to NULL instead.
- `adt.qc.thresholds`: Results of `suggestAdtQcThresholds`. If ADT data is not available, this is set to NULL instead.
- `adt.qc.filter`: Results of `filterAdtQcMetrics`. If ADT data is not available, this is set to NULL instead.
- `crispr.qc.metrics`: Results of `computeCrisprQcMetrics`. If CRISPR data is not available, this is set to NULL instead.
- `crispr.qc.thresholds`: Results of `suggestCrisprQcThresholds`. If CRISPR data is not available, this is set to NULL instead.
- `crispr.qc.filter`: Results of `filterCrisprQcMetrics`. If CRISPR data is not available, this is set to NULL instead.
- `combined.qc.filter`: Logical vector indicating which cells are of high quality and should be retained for downstream analyses.
- `rna.filtered`: Matrix of RNA counts that has been filtered to only contain the high-quality cells in `combined.qc.filter`. If RNA data is not available, this is set to NULL instead.
- `adt.filtered`: Matrix of ADT counts that has been filtered to only contain the high-quality cells in `combined.qc.filter`. If ADT data is not available, this is set to NULL instead.
- `crispr.filtered`: Matrix of CRISPR counts that has been filtered to only contain the high-quality cells in `combined.qc.filter`. If CRISPR data is not available, this is set to NULL instead.
- `rna.size.factors`: Size factors for the RNA count matrix, derived from the sum of counts for each cell and centered with `centerSizeFactors`. If RNA data is not available, this is set to NULL instead.

- `rna.normalized`: Matrix of (log-)normalized expression values derived from RNA counts, as computed by `normalizeCounts` using `rna.size.factors`. If RNA data is not available, this is set to NULL instead.
- `adt.size.factors`: Size factors for the ADT count matrix, computed by `computeClrm1Factors` and centered with `centerSizeFactors`. If ADT data is not available, this is set to NULL instead.
- `adt.normalized`: Matrix of (log-)normalized expression values derived from ADT counts, as computed by `normalizeCounts` using `adt.size.factors`. If ADT data is not available, this is set to NULL instead.
- `crispr.size.factors`: Size factors for the CRISPR count matrix, derived from the sum of counts for each cell and centered with `centerSizeFactors`. If CRISPR data is not available, this is set to NULL instead.
- `crispr.normalized`: Matrix of (log-)normalized expression values derived from CRISPR counts, as computed by `normalizeCounts` using `crispr.size.factors`. If CRISPR data is not available, this is set to NULL instead.
- `rna.gene.variances`: Results of `modelGeneVariances`. If RNA data is not available, this is set to NULL instead.
- `rna.highly.variable.genes`: Results of `chooseHighlyVariableGenes`. If RNA data is not available, this is set to NULL instead.
- `rna.pca`: Results of calling `runPca` on `rna.normalized` with the `rna.highly.variable.genes` subset. If RNA data is not available, this is set to NULL instead.
- `adt.pca`: Results of calling `runPca` on `adt.normalized`. If ADT data is not available, this is set to NULL instead.
- `crispr.pca`: Results of calling `runPca` on `crispr.normalized`. If CRISPR data is not available, this is set to NULL instead.
- `combined.pca`: If only one modality is used for the downstream analysis, this is a string specifying the list element containing the components to be used, e.g., "rna.pca". If multiple modalities are to be combined for downstream analysis, this contains the results of `scaleByNeighbors` on the PCs of those modalities.
- `block`: Vector or factor containing the blocking factor for all cells (after filtering, if `filter.cells = TRUE`). This is set to NULL if no blocking factor was supplied.
- `mnn.corrected`: Results of `correctMnn` on the PCs in or referenced by `combined.pca`. If no blocking factor is supplied, this is set to NULL instead.
- `tsne`: Results of `runTsne`. This is NULL if t-SNE was not performed.
- `umap`: Results of `runUmap`. This is NULL if UMAP was not performed.
- `snn.graph`: Results of `buildSnnGraph`. This is NULL if graph-based clustering was not performed, or if `return.graph=FALSE` in `runAllNeighborSteps`.
- `graph.clusters`: Results of `clusterGraph`. This is NULL if graph-based clustering was not performed.
- `kmeans.clusters`: Results of `clusterKmeans`. This is NULL if k-means clustering was not performed.
- `clusters`: Integer vector containing the cluster assignment for each cell (after filtering, if `filter.cells = TRUE`). This may be derived from `graph.clusters` or `kmeans.clusters` depending on the choice of `clusters.for.markers`. If no suitable clusterings are available, this is set to NULL.
- `rna.markers`: Results of calling `scoreMarkers` on `rna.normalized`. This is NULL if RNA data is not available or no suitable clusterings are available.

adt.markers: Results of calling [scoreMarkers](#) on adt.normalized. This is NULL if ADT data is not available or no suitable clusterings are available.

crispr.markers: Results of calling [scoreMarkers](#) on crispr.normalized. This is NULL if CRISPR data is not available or no suitable clusterings are available.

Author(s)

Aaron Lun

See Also

[convertAnalyzeResults](#), to convert the results into a [SingleCellExperiment](#).

Examples

```
sce <- getTestRnaData.se()
sce <- sce[,1:500] # smaller dataset for a faster runtime for R CMD check.
res <- analyze(
  sce,
  rna.subsets=list(mito=grep("^mt-", rownames(sce))),
  num.threads=2 # keep R CMD check happy
)
str(res)
convertAnalyzeResults(res)
```

analyze.se

Analyze single-cell data from a SummarizedExperiment

Description

Execute a simple single-cell analysis pipeline, starting from a count matrix and ending with clusters, visualizations and markers. This also supports integration of multiple modalities and correction of batch effects.

Usage

```
analyze.se(
  x,
  rna.altexp = NA,
  adt.altexp = NULL,
  crispr.altexp = NULL,
  rna.assay.type = "counts",
  adt.assay.type = "counts",
  crispr.assay.type = "counts",
  block = NULL,
  block.name = "block",
  rna.qc.subsets = list(),
  rna.qc.output.prefix = NULL,
  more.rna.qc.args = list(),
  adt.qc.subsets = list(),
  adt.qc.output.prefix = NULL,
```

```

more.adt.qc.args = list(),
crispr.qc.output.prefix = NULL,
more.crispr.qc.args = list(),
filter.cells = TRUE,
rna.norm.output.name = "logcounts",
more.rna.norm.args = list(),
adt.norm.output.name = "logcounts",
more.adt.norm.args = list(),
crispr.norm.output.name = "logcounts",
more.crispr.norm.args = list(),
rna.hvg.output.prefix = NULL,
more.rna.hvg.args = list(),
rna.pca.output.name = "PCA",
more.rna.pca.args = list(),
adt.pca.output.name = "PCA",
more.adt.pca.args = list(),
use.rna.pcs = TRUE,
use.adt.pcs = TRUE,
scale.output.name = "combined",
more.scale.args = list(),
mnn.output.name = "MNN",
more.mnn.args = list(),
more.umap.args = list(),
more.tsne.args = list(),
cluster.graph.output.name = "graph.cluster",
more.build.graph.args = list(),
more.cluster.graph.args = list(),
more.neighbor.args = list(),
kmeans.clusters = NULL,
kmeans.clusters.output.name = "kmeans.cluster",
more.kmeans.args = list(),
clusters.for.markers = c("graph", "kmeans"),
more.rna.marker.args = list(),
more.adt.marker.args = list(),
more.crispr.marker.args = list(),
BNPARAM = AnnoyParam(),
num.threads = 3L
)

```

Arguments

x	A SummarizedExperiment object or one of its subclasses. Rows correspond to genomic features (genes, ADTs or CRISPR guides) and columns correspond to cells.
rna.altexp	String or integer specifying the alternative experiment of x containing the RNA data. If NA, the main experiment is assumed to contain the RNA data. If NULL, it is assumed that no RNA data is available.
adt.altexp	String or integer specifying the alternative experiment of x containing the ADT data. If NA, the main experiment is assumed to contain the ADT data. If NULL, it is assumed that no ADT data is available.
crispr.altexp	String or integer specifying the alternative experiment of x containing the CRISPR data. If NA, the main experiment is assumed to contain the CRISPR data. If NULL,

it is assumed that no CRISPR data is available.

rna.assay.type	String or integer specifying the assay containing the RNA count data. Only used if rna.altexp is not NULL.
adt.assay.type	String or integer specifying the assay containing the ADT count data. Only used if adt.altexp is not NULL.
crispr.assay.type	String or integer specifying the assay containing the CRISPR count data. Only used if crispr.altexp is not NULL.
block	Vector or factor specifying the block of origin (e.g., batch, sample) for each cell in x. Alternatively NULL, if all cells are from the same block.
block.name	String specifying the name of the colData column in which to store the blocking factor. Only used if block is not NULL. If NULL, the blocking factor is not stored in the colData.
rna.qc.subsets	Passed to quickRnaQc.se as the subsets argument. Only used if rna.altexp is not NULL.
rna.qc.output.prefix	Passed to quickRnaQc.se as the output.prefix argument. Only used if rna.altexp is not NULL.
more.rna.qc.args	Named list of additional arguments to pass to quickRnaQc.se . Only used if rna.altexp is not NULL.
adt.qc.subsets	Passed to quickAdtQc.se as the subsets argument. Only used if adt.altexp is not NULL.
adt.qc.output.prefix	Passed to quickAdtQc.se as the output.prefix argument. Only used if adt.altexp is not NULL.
more.adt.qc.args	Named list of additional arguments to pass to quickAdtQc.se . Only used if adt.altexp is not NULL.
crispr.qc.output.prefix	Passed to quickCrisprQc.se as the output.prefix argument. Only used if crispr.altexp is not NULL.
more.crispr.qc.args	Named list of additional arguments to pass to quickCrisprQc.se . Only used if crispr.altexp is not NULL.
filter.cells	Logical scalar indicating whether to filter x to only retain high-quality cells in all modalities. If FALSE, QC metrics and thresholds are still computed but are not used to filter the count matrices.
rna.norm.output.name	Passed to normalizeRnaCounts.se as the output.name argument. Only used if rna.altexp is not NULL.
more.rna.norm.args	Named list of arguments to pass to normalizeRnaCounts.se . Only used if rna.altexp is not NULL.
adt.norm.output.name	Passed to normalizeAdtCounts.se as the output.name argument. Only used if adt.altexp is not NULL.
more.adt.norm.args	Named list of arguments to pass to normalizeAdtCounts.se . Only used if adt.altexp is not NULL.

`crispr.norm.output.name`
Passed to [normalizeCrisprCounts.se](#) as the `output.name` argument. Only used if `crispr.altexp` is not NULL.

`more.crispr.norm.args`
Named list of arguments to pass to [normalizeCrisprCounts.se](#). Only used if `crispr.altexp` is not NULL.

`rna.hvg.output.prefix`
Passed to [chooseRnaHvgs.se](#) as the `output.prefix` argument. Only used if `rna.altexp` is provided.

`more.rna.hvg.args`
Named list of arguments to pass to [chooseRnaHvgs.se](#). Only used if `rna.altexp` is provided.

`rna.pca.output.name`
Passed to [runPca.se](#) as the `output.name` argument. Only used if `rna.altexp` is provided.

`more.rna.pca.args`
Named list of arguments to pass to [runPca.se](#). Only used if `rna.altexp` is provided.

`adt.pca.output.name`
Passed to [runPca.se](#) as the `output.name` argument. Only used if `adt.altexp` is provided.

`more.adt.pca.args`
Named list of arguments to pass to [runPca.se](#). Only used if `adt.altexp` is provided.

`use.rna.pcs`
Logical scalar indicating whether to use the RNA-derived PCs for downstream steps (i.e., clustering, visualization). Only used if `rna.altexp` is provided.

`use.adt.pcs`
Logical scalar indicating whether to use the ADT-derived PCs for downstream steps (i.e., clustering, visualization). Only used if `adt.altexp` is provided.

`scale.output.name`
Passed to [scaleByNeighbors.se](#) as the `output.name` argument. Only used if multiple modalities are available and their corresponding `use.*.pcs` arguments are TRUE.

`more.scale.args`
Named list of arguments to pass to [scaleByNeighbors.se](#). Only used if multiple modalities are available and their corresponding `use.*.pcs` arguments are TRUE.

`mnn.output.name`
Passed to [correctMnn.se](#) as the `output.name` argument. Only used if `block` is supplied.

`more.mnn.args`
Named list of arguments to pass to [correctMnn.se](#). Only used if `block` is supplied.

`more.umap.args`
Passed to [runAllNeighborSteps.se](#).

`more.tsne.args`
Passed to [runAllNeighborSteps.se](#).

`cluster.graph.output.name`
Passed to [runAllNeighborSteps.se](#) as `cluster.output.name`.

`more.build.graph.args`
Passed to [runAllNeighborSteps.se](#).

`more.cluster.graph.args`
Passed to [runAllNeighborSteps.se](#).

<code>more.neighbor.args</code>	Passed to runAllNeighborSteps.se .
<code>kmeans.clusters</code>	Passed to clusterKmeans.se as the <code>k</code> argument. If <code>NULL</code> , k-means clustering is not performed.
<code>kmeans.clusters.output.name</code>	Passed to clusterKmeans.se as the <code>output.name</code> argument. Ignored if <code>kmeans.clusters = NULL</code> .
<code>more.kmeans.args</code>	Named list of arguments to pass to clusterKmeans.se . Ignored if <code>kmeans.clusters = NULL</code> .
<code>clusters.for.markers</code>	Character vector of clustering algorithms (either "graph" or "kmeans"), specifying the clustering to be used for marker detection. The first available clustering will be chosen. If no clustering is available from the list, markers will not be computed.
<code>more.rna.marker.args</code>	Named list of arguments to pass to scoreMarkers.se for the RNA data. Ignored if no suitable clusterings are available or if <code>rna.altexp=NULL</code> .
<code>more.adt.marker.args</code>	Named list of arguments to pass to scoreMarkers.se for the ADT data. Ignored if no suitable clusterings are available or if <code>adt.altexp=NULL</code> .
<code>more.crispr.marker.args</code>	Named list of arguments to pass to scoreMarkers.se for the CRISPR data. Ignored if no suitable clusterings are available or if <code>crispr.altexp=NULL</code> .
<code>BNPARAM</code>	A BiocNeighborParam instance specifying the nearest-neighbor search algorithm to use.
<code>num.threads</code>	Integer scalar specifying the number of threads to use in each step.

Details

This function is equivalent to:

- Running [quickRnaQc.se](#), [quickAdtQc.se](#) and/or [quickCrisprQc.se](#), for quality control.
- Subsetting `x` to only retain the high-quality cells in all modalities, based on `filter.cells`.
- Running [normalizeRnaCounts.se](#), [normalizeAdtCounts.se](#) and/or [normalizeCrisprCounts.se](#), for normalization.
- Running [chooseRnaHvgs.se](#) to identify highly variable genes.
- Running [runPca.se](#) on the RNA and/or ADT data.
- Running [scaleByNeighbors.se](#) if multiple modalities are present.
- Running [correctMnn.se](#) if multiple batches are present in `block`.
- Running [runAllNeighborSteps.se](#) to obtain t-SNE and UMAP coordinates, and to perform graph-based clustering.
- Running [clusterKmeans.se](#) to perform k-means clustering, depending on `kmeans.clusters`.
- Running [scoreMarkers.se](#) to compute markers for each modality based on one of the clusterings.

Value

List containing:

- `x`, a [SingleCellExperiment](#) that is a copy of the input `x`. It is also decorated with the results of each analysis step - see [Details](#).
- `markers`, a list of list of [DataFrames](#) containing the marker statistics for each modality. Each inner list corresponds to a modality (RNA, ADT, etc.) while each `DataFrame` corresponds to a cluster. If no clusterings were generated, this is set to `NULL`.

Author(s)

Aaron Lun

Examples

```
library(SingleCellExperiment)
sce <- getTestRnaData.se("start")
res <- analyze.se(
  sce,
  rna.qc.subsets=list(mito=grep("^mt-", rownames(sce))),
  num.threads=2 # keep R CMD check happy
)
assayNames(res$x)
reducedDimNames(res$x)
colData(res$x)
previewMarkers(res$markers$rna[[1]], "cohens.d.mean")
```

buildSnnGraph

Build a shared nearest neighbor graph

Description

Build a shared nearest neighbor (SNN) graph where each node is a cell. Edges are formed between cells that share one or more nearest neighbors, weighted by the number or ranking of those shared neighbors. If two cells are close together but have distinct sets of neighbors, the corresponding edge is downweighted as the two cells are unlikely to be part of the same neighborhood. In this manner, strongly weighted edges will only form within highly interconnected neighborhoods where many cells share the same neighbors. This provides a more sophisticated definition of similarity between cells compared to a simpler (unweighted) nearest neighbor graph that just focuses on immediate proximity.

Usage

```
buildSnnGraph(
  x,
  num.neighbors = 10,
  weight.scheme = "ranked",
  num.threads = 1,
  BNPARAM = AnnoyParam(),
  as.pointer = FALSE
)
```

Arguments

<code>x</code>	For <code>buildSnnGraph</code> , a numeric matrix where rows are dimensions and columns are cells, typically containing a low-dimensional representation from, e.g., <code>runPca</code> . Alternatively, a named list of nearest-neighbor search results. This should contain <code>index</code> , an integer matrix where rows are neighbors and columns are cells. Each column contains 1-based indices for the nearest neighbors of the corresponding cell, ordered by increasing distance. The number of neighbors for each cell should be equal to <code>num.neighbors</code> , otherwise a warning is raised. Alternatively, an index constructed by <code>buildIndex</code> .
<code>num.neighbors</code>	Integer scalar specifying the number of neighbors to use to construct the graph. Larger values increase the connectivity of the graph and reduce the granularity of subsequent community detection steps, at the cost of speed. Ignored if <code>x</code> contains pre-computed neighbor search results.
<code>weight.scheme</code>	String specifying the weighting scheme to use for constructing the SNN graph. This can be one of: <ul style="list-style-type: none"> • "ranked", where the weight of the edge is defined by the smallest sum of ranks across all shared neighbors. More shared neighbors, or shared neighbors that are close to both observations, will generally yield larger weights. • "number", where the weight of the edge is the number of shared nearest neighbors between them. This is a simpler scheme that is also slightly faster but does not account for the ranking of neighbors within each set. • "jaccard", where the weight of the edge is the Jaccard index of their neighbor sets. This is a monotonic transformation of the weight used in "number".
<code>num.threads</code>	Integer scalar specifying the number of threads to use. Only used if <code>x</code> is not a list of existing nearest-neighbor search results.
<code>BNPARAM</code>	A <code>BiocNeighborParam</code> object specifying the algorithm to use. Only used if <code>x</code> is not a list of existing nearest-neighbor search results.
<code>as.pointer</code>	Logical scalar indicating whether to return an external pointer for direct use in <code>clusterGraph</code> . This avoids the extra memory usage caused by conversion to/from an R list.

Value

If `as.pointer=FALSE`, a list is returned containing:

- `vertices`, an integer scalar specifying the number of vertices in the graph (i.e., cells in `x`).
- `edges`, an integer vector of 1-based indices for graph edges. Pairs of values represent the endpoints of an (undirected) edge, i.e., `edges[1:2]` form the first edge, `edges[3:4]` form the second edge and so on.
- `weights`, a numeric vector of weights for each edge. This has length equal to half the length of edges.

If `as.pointer=TRUE`, an external pointer to the graph is returned that can be directly used in `clusterGraph`.

Author(s)

Aaron Lun

See Also

The `build_snn_graph` function in https://libscran.github.io/scran_graph_cluster/. `clusterGraph`, to define clusters (i.e., communities) from the graph. `clusterGraph.se`, which builds an SNN graph from a `SingleCellExperiment`.

Examples

```
data <- matrix(rnorm(10000), ncol=1000)
out <- buildSnnGraph(data)
str(out)

# We can use this to make an igraph::graph.
g <- igraph::make_undirected_graph(out$edges, n = out$vertices)
igraph::E(g)$weight <- out$weight
```

centerSizeFactors *Center size factors*

Description

Scale the size factors so they are centered at unity. This ensures that the original scale of the counts is preserved in the normalized values from `normalizeCounts`, which simplifies interpretation and ensures that any pseudo-count added prior to log-transformation has a predictable shrinkage effect.

Usage

```
centerSizeFactors(size.factors, block = NULL, mode = c("lowest", "per-block"))
```

Arguments

<code>size.factors</code>	Numeric vector of size factors across cells. Invalid size factors (e.g., non-positive, non-finite) will be ignored.
<code>block</code>	Vector or factor of length equal to <code>size.factors</code> , specifying the block of origin for each cell. Alternatively <code>NULL</code> , in which case all cells are assumed to be in the same block.
<code>mode</code>	String specifying how to scale size factors across blocks. This can be either "lowest" or "per-block", see Details. Only used if <code>block</code> is provided.

Details

"lowest" will compute the average size factor in each block, identify the lowest average across all blocks, and then scale all size factors by that value. Here, our normalization strategy involves downscaling all blocks to match the coverage of the lowest-coverage block. This is useful for datasets with big differences in coverage between blocks as it avoids egregious upscaling of low-coverage blocks. Specifically, strong upscaling allows the log-transformation to ignore any shrinkage from the pseudo-count. This is problematic as it inflates differences between cells at log-values derived from low counts, increasing noise and overstating log-fold changes. Downscaling is safer as it allows the pseudo-count to shrink the log-differences between cells towards zero at low counts, effectively sacrificing some information in the higher-coverage batches so that they can be compared

to the low-coverage batches (which is preferable to exaggerating the informativeness of the latter for comparison to the former).

"per-block" will compute the average size factor in each block, and then scale each size factor by the average of block to which it belongs. The scaled size factors are identical to those obtained by separate invocations of 'center_size_factors()' on the size factors for each block. This can be desirable to ensure consistency with independent analyses of each block - otherwise, the centering would depend on the size factors in other blocks. However, any systematic differences in the size factors between blocks are lost, i.e., systematic changes in coverage between blocks will not be normalized.

Value

Numeric vector of length equal to size.factors, containing the centered size factors.

Author(s)

Aaron Lun

See Also

The center_size_factors and center_size_factors_blocked functions in https://libscran.github.io/scran_norm/.

[normalizeRnaCounts.se](#) and related functions, which center the size factors prior to normalization of a [SummarizedExperiment](#).

Examples

```
centerSizeFactors(runif(100))  
centerSizeFactors(runif(100), block=sample(3, 100, replace=TRUE))
```

chooseHighlyVariableGenes

Choose highly variable genes

Description

Choose highly variable genes (HVGs) based on a variance-related statistic. This is typically used to subset the gene-cell matrix prior to calling [runPca](#).

Usage

```
chooseHighlyVariableGenes(  
  stats,  
  top = 4000,  
  larger = TRUE,  
  keep.ties = TRUE,  
  bound = 0  
)
```

Arguments

stats	Numeric vector of variances (or a related statistic) across all genes. Typically, the residuals from <code>modelGeneVariances</code> are used here.
top	Integer specifying the number of top genes to retain. Note that the actual number of retained genes may not be equal to top, depending on the other options.
larger	Logical scalar indicating whether larger values of stats correspond to more variable genes. If TRUE, HVGs are defined as those with the largest values of stats. This is typically the case for variances or related statistics, e.g., residuals.
keep.ties	Logical scalar indicating whether to keep tied values of stats, even if top may be exceeded.
bound	Numeric scalar specifying the lower bound (if larger=TRUE) or upper bound (otherwise) to be applied to stats. Genes are not considered to be HVGs if they do not satisfy this bound, even if they are within the top genes. For example, residuals from the fitted trend should be positive, which can be enforced by setting bound to zero. Ignored if NULL.

Value

Integer vector containing the indices of genes in stats that are considered to be highly variable.

Author(s)

Aaron Lun

See Also

The `choose_highly_variable_genes` function in https://libscrn.github.io/scrn_variances/chooseRnaHvgs.se, which choose the HVGs from the residuals computed from a `SummarizedExperiment`.

Examples

```
resids <- rexp(10000)
str(chooseHighlyVariableGenes(resids))
```

choosePseudoCount *Choose a suitable pseudo-count*

Description

Choose a suitable pseudo-count to control the bias introduced by log-transformation of normalized counts from `normalizeCounts`. Larger pseudo-counts shrink log-expression values towards the zero-expression baseline, reducing the impact of the transformation bias at the cost of some sensitivity.

Usage

```
choosePseudoCount(size.factors, quantile = 0.05, max.bias = 1, min.value = 1)
```

Arguments

size.factors	Numeric vector of size factors for all cells. It is expected that these have already been centered, e.g., with centerSizeFactors . Invalid size factors (e.g., non-positive, non-finite) will be ignored.
quantile	Numeric scalar specifying the quantile to use for finding the smallest/largest size factors. Setting this to zero will use the observed minimum and maximum, though in practice, this is usually too sensitive to outliers. The default is to take the 5th and 95th percentile to obtain a range that captures most of the distribution.
max.bias	Numeric scalar specifying the maximum allowed bias. This is the maximum absolute value of any spurious log ₂ -fold change between the cells with the smallest and largest size factors.
min.value	Numeric scalar specifying the minimum value for the pseudo-count. Defaults to 1 to stabilize near-zero normalized expression values, otherwise these manifest as avoid large negative values.

Value

A choice of pseudo-count for [normalizeCounts](#).

Author(s)

Aaron Lun

References

Lun ATL (2018). Overcoming systematic errors caused by log-transformation of normalized single-cell RNA sequencing data. *bioRxiv* doi:10.1101/404962

See Also

choose_pseudo_count in https://libscrn.github.io/scrn_norm/.

Examples

```
sf <- centerSizeFactors(runif(100))
choosePseudoCount(sf)
choosePseudoCount(sf, quantile=0.01)
choosePseudoCount(sf, max.bias=0.5)
```

chooseRnaHvgs.se

Choose highly variable genes from a SummarizedExperiment

Description

Model the mean-variance relationship across genes and choose highly variable genes (HVGs) based on the residuals of the fitted trend. This calls [modelGeneVariances](#) on an assay of a [SummarizedExperiment](#), and then calls [chooseHighlyVariableGenes](#) on the residuals.

Usage

```

chooseRnaHvgs.se(
  x,
  block = NULL,
  num.threads = 1,
  more.var.args = list(),
  top = 4000,
  more.choose.args = list(),
  assay.type = "logcounts",
  output.prefix = NULL,
  include.per.block = FALSE
)

formatModelGeneVariancesResult(
  model.res,
  choose.res = NULL,
  include.per.block = FALSE
)

```

Arguments

x	A SummarizedExperiment object or one of its subclasses. Rows correspond to genes and columns correspond to cells.
block	Block assignment for each cell, to pass to modelGeneVariances .
num.threads	Number of threads, to pass to modelGeneVariances .
more.var.args	Named list of arguments to pass to modelGeneVariances .
top	Number of HVGs to choose, to pass to chooseHighlyVariableGenes .
more.choose.args	Named list of arguments to pass to chooseHighlyVariableGenes .
assay.type	Integer or string specifying the assay of x containing the log-normalized expression matrix for the RNA data.
output.prefix	String containing a prefix to add to the names of the <code>link[SummarizedExperiment]{rowData}</code> columns containing the output statistics.
include.per.block	Logical scalar indicating whether the per-block statistics should be stored in the output <code>rowData</code> . Only relevant if <code>block</code> is specified.
model.res	List returned by modelGeneVariances .
choose.res	Integer vector returned by chooseHighlyVariableGenes . This may be <code>NULL</code> , in which case the identities of the HVGs will not be stored.

Value

For `chooseRnaHvgs.se`, `x` is returned with the per-gene variance modelling statistics added to its `rowData`. The `hvg` column in the `rowData` indicates whether a gene was chosen as a HVG. If `include.per.block=TRUE` and `block` is specified, the per-block statistics are stored as a nested `DataFrame` in the `per.block` column.

For `formatModelGeneVariancesResult`, a `DataFrame` is returned with the per-gene variance modelling statistics. If `choose.res` is provided, a `hvg` column is also stored that indicates whether a gene was chosen as a HVG.

Author(s)

Aaron Lun

Examples

```
library(SingleCellExperiment)
sce <- getTestRnaData.se("norm")
sce <- chooseRnaHvgs.se(sce, more.var.args=list(use.min.width=TRUE))
summary(rowData(sce)$hvg)

plot(rowData(sce)$means, rowData(sce)$variances, col=factor(rowData(sce)$hvg))
curve(approxfun(rowData(sce)$means, rowData(sce)$fitted)(x), col="dodgerblue", add=TRUE)
```

clusterGraph

*Graph-based clustering of cells***Description**

Identify clusters by applying community detection algorithms to a graph. This assumes that the nodes on the graph represent cells and weighted edges are formed between related cells.

Usage

```
clusterGraph(
  x,
  method = c("multilevel", "leiden", "walktrap"),
  multilevel.resolution = 1,
  leiden.resolution = 1,
  leiden.objective = c("modularity", "cpm", "er"),
  walktrap.steps = 4,
  seed = 42
)
```

Arguments

x	List containing graph information or an external pointer to a graph, as returned by <code>buildSnnGraph</code> . Alternatively, an <code>igraph</code> object with edge weights.
method	String specifying the algorithm to use. <ul style="list-style-type: none"> "multilevel" uses multi-level modularity optimization, also known as the Louvain algorithm, see https://igraph.org/c/doc/igraph-Community.html#igraph_community_multilevel for details. "walktrap" uses the Walktrap community finding algorithm, see https://igraph.org/c/doc/igraph-Community.html#igraph_community_walktrap for details. "leiden" uses the Leiden algorithm, see https://igraph.org/c/doc/igraph-Community.html#igraph_community_leiden for details.
multilevel.resolution	Numeric scalar specifying the resolution when method="multilevel". Lower values favor fewer, larger communities; higher values favor more, smaller communities.

leiden.resolution	Numeric scalar specifying the resolution when method="leiden". Lower values favor fewer, larger communities; higher values favor more, smaller communities.
leiden.objective	String specifying the objective function when method="leiden". "modularity" uses the generalized modularity, "cpm" uses the Constant Potts Model, and "er" uses the Erdős-Rényi G(n, p) model. The CPM typically yields more fine-grained clusters than the modularity at the same leiden.resolution.
walktrap.steps	Integer scalar specifying the number of steps to use when method="walktrap". This determines the ability of the Walktrap algorithm to distinguish highly interconnected communities from the rest of the graph.
seed	Integer scalar specifying the random seed to use for method="multilevel" or "leiden".

Value

A list containing membership, a factor containing the cluster assignment for each cell. Additional fields may be present depending on the method:

- For method="multilevel", the levels list contains the clustering result at each level of the algorithm. A modularity numeric vector also contains the modularity at each level, the highest of which corresponds to the reported membership.
- For method="leiden", a quality numeric scalar containing the quality of the partitioning.
- For method="walktrap", a merges matrix specifies the pair of cells or clusters that were merged at each step of the algorithm. A modularity numeric scalar also contains the modularity of the final partitioning.

Author(s)

Aaron Lun

See Also

The various cluster_* functions in https://libscran.github.io/scran_graph_cluster/. [clusterGraph.se](#), which performs clustering on graph constructed from a [SingleCellExperiment](#).

Examples

```
data <- matrix(rnorm(10000), ncol=1000)
gout <- buildSnnGraph(data)
str(gout)

str(clusterGraph(gout))
str(clusterGraph(gout, method="leiden"))
str(clusterGraph(gout, method="walktrap"))
```

clusterGraph.se

*Graph-based clustering of cells in a SingleCellExperiment***Description**

Construct a shared-nearest neighbor (SNN) graph from an existing low-dimensional embedding by calling `buildSnnGraph` on a reduced dimension entry in a `SingleCellExperiment`. Then, apply community detection algorithms to obtain clusters of cells with `clusterGraph`.

Usage

```
clusterGraph.se(
  x,
  num.neighbors = 10,
  num.threads = 1,
  more.build.args = list(),
  method = "multilevel",
  resolution = NULL,
  more.cluster.args = list(),
  reddim.type = "PCA",
  output.name = "clusters",
  meta.name = NULL,
  graph.name = NULL
)
```

Arguments

<code>x</code>	A <code>SingleCellExperiment</code> object or one of its subclasses. Rows correspond to genomic features and columns correspond to cells.
<code>num.neighbors</code>	Number of neighbors for constructing the graph, passed to <code>buildSnnGraph</code> .
<code>num.threads</code>	Number of threads for graph construction, passed to <code>buildSnnGraph</code> .
<code>more.build.args</code>	Named list of further arguments to be passed to <code>buildSnnGraph</code> .
<code>method</code>	Clustering method to use, passed to <code>clusterGraph</code> .
<code>resolution</code>	Resolution for the community detection method in <code>clusterGraph</code> . This is either passed to <code>multilevel.resolution</code> or <code>leiden.resolution</code> depending on method.
<code>more.cluster.args</code>	Named list of further arguments to be passed to <code>clusterGraph</code> .
<code>reddim.type</code>	Integer or string specifying the existing embedding in the <code>reducedDim</code> of <code>x</code> . Alternatively, a named integer or character vector of length 1, where the name specifies an alternative experiment of <code>x</code> and the value is the name/index of a <code>reducedDim</code> entry in that alternative experiment.
<code>output.name</code>	String containing the name of the column of the <code>colData</code> in which to store the cluster assignments.
<code>meta.name</code>	String containing the name of the <code>metadata</code> entry in which to store extra clustering output. If <code>NULL</code> , no extra clustering output is stored.
<code>graph.name</code>	String containing the name of the <code>metadata</code> entry in which to store the SNN graph. If <code>NULL</code> , the SNN graph is not stored.

Value

x is returned with the cluster assignment for each cell stored in the colData. Additional clustering output is stored in the metadata.

Author(s)

Aaron Lun

Examples

```
sce <- getTestRnaData.se("pca")
sce <- clusterGraph.se(sce)
table(sce$clusters)
```

clusterKmeans

K-means clustering

Description

Perform k-means clustering with a variety of different initialization and refinement algorithms.

Usage

```
clusterKmeans(
  x,
  k,
  init.method = c("var-part", "kmeans++", "random"),
  refine.method = c("hartigan-wong", "lloyd"),
  var.part.optimize.partition = TRUE,
  var.part.size.adjustment = 1,
  lloyd.iterations = 100,
  hartigan.wong.iterations = 10,
  hartigan.wong.quick.transfer.iterations = 50,
  hartigan.wong.quit.quick.transfer.failure = FALSE,
  seed = 5489L,
  warn = TRUE,
  num.threads = 1
)
```

Arguments

- | | |
|-------------|--|
| x | Matrix-like object where rows are dimensions and columns are cells. This is typically a dense double-precision matrix containing a low-dimensional representation from, e.g., <code>runPca</code> . However, any matrix representation supported by <code>initializeCpp</code> can also be used. |
| k | Integer scalar specifying the number of clusters. |
| init.method | String specifying the initialization method for the centers: |

- "var-part" uses variance partitioning as described by Su and Dy (2007). The dataset is repeatedly split along the dimension of greatest variance until k partitions are formed, the centroids of which form the initial clusters. This approach is slower than the others but fully deterministic.
 - "kmeans++" uses the weighted sampling method described by Arthur and Vassilvitskii (2007). k points are sampled with probability based on the smallest distance to any previously sampled point. This improves the likelihood of choosing initial centroids that are far apart from each other.
 - "random" initialization involves choosing k random points as the initial centers. This is the simplest and fastest method but may not yield good starting points.
- `refine.method` String specifying the refinement method.
- "lloyd" uses Lloyd's algorithm, which performs a batch update in each iteration. This is simple and amenable to parallelization but may not converge.
 - "hartigan-wong" uses the Hartigan-Wong algorithm, which transfers points between clusters to optimize the drop in the within-cluster sum of squares. This is slower but has a greater chance of convergence.
- `var.part.optimize.partition` Logical scalar indicating whether each partition boundary should be optimized to reduce the sum of squares in the child partitions. This is slower but improves the quality of the partition. Only used if `init.method = "var.part"`.
- `var.part.size.adjustment` Numeric scalar between 0 and 1, specifying the adjustment to the cluster size when selecting the next cluster to partition. Setting this to 0 or 1 will select the cluster with the highest variance or sum of squares, respectively, for partitioning. In other words, a value of 0 will ignore the cluster size while setting a value of 1 will generally cause larger clusters to be selected. Only used if `init.method = "var.part"`.
- `lloyd.iterations` Integer scalar specifying the maximum number of iterations for the Lloyd algorithm. Larger values increase the chance of convergence at the cost of increasing compute time. Only used if `refine.method = "lloyd"`.
- `hartigan.wong.iterations` Integer scalar specifying the maximum number of iterations for the Hartigan-Wong algorithm. Larger values increase the chance of convergence at the cost of increasing compute time. Only used if `refine.method = "hartigan-wong"`.
- `hartigan.wong.quick.transfer.iterations` Integer scalar specifying the maximum number of quick transfer iterations for the Hartigan-Wong algorithm. Larger values increase the chance of convergence at the cost of increasing compute time. Only used if `refine.method = "hartigan-wong"`.
- `hartigan.wong.quit.quick.transfer.failure` Logical scalar indicating whether to quit the Hartigan-Wong algorithm upon convergence failure during quick transfer iterations. Setting this to FALSE gives the algorithm another chance to converge by attempting another optimal transfer iteration, at the cost of more compute time. If TRUE, the function follows the same behavior as R's `kmeans`. Only used if `refine.method = "hartigan-wong"`.
- `seed` Integer scalar specifying the seed for random number generation. Only used if `init.method = "random" or "kmeans++"`.

warn	Boolean specifying whether a warning should be emitted if the k-means algorithm failed to converge.
num.threads	Integer scalar specifying the number of threads to use.

Value

By default, a list is returned containing:

- `clusters`, a factor containing the cluster assignment for each cell. The number of levels is no greater than `k`, where each level is an integer that refer to a column of centers.
- `centers`, a numeric matrix with the coordinates of the cluster centroids (dimensions in rows, centers in columns). The number of columns is no greater than `k`. Empty clusters are automatically removed.
- `iterations`, an integer scalar specifying the number of refinement iterations that were performed.
- `status`, an integer scalar specifying the completion status of the algorithm. A value of zero indicates success while the meaning of any non-zero value depends on the choice of `refine.method`:
 - For Lloyd, a value of 2 indicates convergence failure.
 - For Hartigan-Wong, a value of 2 indicates convergence failure in the optimal transfer iterations. A value of 4 indicates convergence failure in the quick transfer iterations when `hartigan.wong.quit.quick.transfer.failure = TRUE`.

Author(s)

Aaron Lun

References

Hartigan JA. and Wong MA (1979). Algorithm AS 136: A K-means clustering algorithm. *Applied Statistics* 28, 100-108.

Arthur D and Vassilvitskii S (2007). k-means++: the advantages of careful seeding. *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms* 1027-1035.

Su T and Dy JG (2007). In Search of Deterministic Methods for Initializing K-Means and Gaussian Mixture Clustering. *Intelligent Data Analysis* 11, 319-338.

See Also

<https://libscan.github.io/kmeans/>, which describes the various initialization and refinement algorithms in more detail.

clusterkmeans.se, for k-means clustering on a [SingleCellExperiment](#).

Examples

```
x <- t(as.matrix(iris[,1:4]))
clustering <- clusterKmeans(x, k=3)
table(clustering$clusters, iris[,"Species"])
```

clusterKmeans.se	<i>k-means clustering of cells in a SingleCellExperiment</i>
------------------	--

Description

Perform k-means clustering on an existing low-dimensional embedding by calling `clusterKmeans` on a reduced dimension entry in a `SingleCellExperiment`.

Usage

```
clusterKmeans.se(
  x,
  k,
  num.threads = 1,
  more.kmeans.args = list(),
  reddim.type = "PCA",
  output.name = "clusters",
  meta.name = NULL
)
```

Arguments

<code>x</code>	A <code>SingleCellExperiment</code> object or one of its subclasses. Rows correspond to genomic features and columns correspond to cells.
<code>k</code>	Number of clusters, passed to <code>clusterKmeans</code> .
<code>num.threads</code>	Number of threads, passed to <code>clusterKmeans</code> .
<code>more.kmeans.args</code>	Named list of further arguments to be passed to <code>clusterKmeans</code> .
<code>reddim.type</code>	Integer or string specifying the existing embedding in the <code>reducedDim</code> of <code>x</code> . Alternatively, a named integer or character vector of length 1, where the name specifies an alternative experiment of <code>x</code> and the value is the name/index of a <code>reducedDim</code> entry in that alternative experiment.
<code>output.name</code>	String containing the name of the <code>colData</code> column in which to store the cluster assignments.
<code>meta.name</code>	String containing the name of the <code>metadata</code> entry in which to store extra clustering output. If <code>NULL</code> , no extra clustering output is stored.

Value

`x` is returned with the cluster assignment for each cell stored in the `colData`. Additional clustering output is stored in the `metadata`.

Author(s)

Aaron Lun

Examples

```
sce <- getTestRnaData.se("pca")
sce <- clusterKmeans.se(sce, k=10)
table(sce$clusters)
```

 combineFactors

Combine multiple factors

Description

Combine multiple factors into a single factor where each level of the latter is a unique combination of levels from the former.

Usage

```
combineFactors(factors, keep.unused = FALSE)
```

Arguments

factors An ordinary list or [List](#) of vectors or factors of the same length. Corresponding elements across all vectors/factors represent the combination of levels for a single observation. For factors, any existing levels are respected. For other vectors, the sorted and unique values are used as levels.

Alternatively, a data frame or [DataFrame](#) where each column is a vector or factor and each row corresponds to an observation.

keep.unused Logical scalar indicating whether to report unused combinations of levels.

Value

List containing:

- **levels**, a [DataFrame](#) containing the sorted and unique combinations of levels from factors. Each column corresponds to a factor in **factors** while each row corresponds to a unique combination.
- **index**, an integer vector specifying the index into **levels** for each observation.

For observation *i* and factor *j*, `levels[[j]][index[i]]` will recover `factors[[j]][i]`.

Author(s)

Aaron Lun

See Also

The `combine_to_factor` function in <https://libscran.github.io/factorize/>.

Examples

```

combineFactors(list(
  sample(LETTERS[1:5], 100, replace=TRUE),
  sample(3, 100, replace=TRUE)
))

combineFactors(list(
  factor(sample(LETTERS[1:5], 10, replace=TRUE), LETTERS[1:5]),
  factor(sample(5, 10, replace=TRUE), 1:5)
), keep.unused=TRUE)

```

computeBlockWeights *Compute block weights*

Description

Compute a weight for each block based on the number of cells in each block. This is typically used to aggregate statistics across blocks, e.g., with weighted sums/averages.

Usage

```

computeBlockWeights(
  sizes,
  block.weight.policy = c("variable", "equal", "size", "none"),
  variable.block.weight = c(0, 1000)
)

```

Arguments

sizes Numeric vector containing the size of (i.e., number of cells in) each block.

block.weight.policy String specifying the policy to use for weighting different blocks. This should be one of:

- "size": the contribution of each block is proportional to its size. "none" is also a deprecated alias for "size".
- "equal": blocks are equally weighted regardless of their size. The exception is that of empty blocks with no cells, which receive zero weight.
- "variable": blocks are equally weighted past a certain threshold size. Below that size, the contribution of each block is proportional to its size. This avoids oversized contributions from very large blocks.

variable.block.weight Numeric vector of length 2, specifying the parameters for variable block weighting. The first and second values are used as the lower and upper bounds, respectively, for the variable weight calculation. Only used if `block.weight.policy = "variable"`.

Value

Numeric vector containing the relative block weights.

Author(s)

Aaron Lun

See AlsoThe `compute_weights` function from https://libscran.github.io/scran_blocks/.**Examples**

```
computeBlockWeights(c(1, 10, 100, 1000, 10000))
computeBlockWeights(c(1, 10, 100, 1000, 10000), block.weight.policy="equal")
computeBlockWeights(c(1, 10, 100, 1000, 10000), variable.block.weight=c(50, 5000))
```

`computeClrm1Factors` *Compute size factors for ADT counts*

Description

Compute size factors from an ADT count matrix using the CLRm1 method. This is a variant of the centered log-ratio (CLR) method, where the size factors are defined from the geometric mean of counts within each cell.

Usage

```
computeClrm1Factors(x, num.threads = 1)
```

Arguments

<code>x</code>	A matrix-like object containing ADT count data. Rows correspond to tags and columns correspond to cells.
<code>num.threads</code>	Number of threads to use.

Value

Numeric vector containing the CLRm1 size factor for each cell. Note that these size factors are not centered and should be passed through, e.g., `centerSizeFactors` before normalization.

Author(s)

Aaron Lun

See Also

<https://github.com/libscran/clrm1>, for a description of the CLRm1 method.
[normalizeAdtCounts.se](#), which computes CLRm1 factors prior to normalization.

Examples

```
library(Matrix)
x <- abs(rsparsematrix(1000, 100, 0.1) * 10)
head(computeClrm1Factors(x))
```

convertAnalyzeResults *Convert analysis results into a SingleCellExperiment*

Description

Convert results from [analyze](#) into a [SingleCellExperiment](#) for further analysis with Bioconductor packages.

Usage

```
convertAnalyzeResults(  
  results,  
  main.modality = NULL,  
  flatten.qc.subsets = TRUE,  
  include.per.block.variances = FALSE  
)
```

Arguments

`results` List of results produced by [analyze](#).

`main.modality` String specifying the modality to use as the main experiment of a [SingleCellExperiment](#).

`flatten.qc.subsets` Logical scalar indicating whether QC metrics for subsets should be flattened in the column data. If FALSE, subset metrics are reported as a nested [DataFrame](#).

`include.per.block.variances` Logical scalar indicating whether the per-block variances should be reported as a nested [DataFrame](#) in the row data.

Value

A [SingleCellExperiment](#) containing most of the analysis results. Filtered and normalized matrices are stored in the assays. QC metrics, size factors and clusterings are stored in the column data. Gene variances are stored in the row data. PCA, t-SNE and UMAP results are stored in the reduced dimensions. Further modalities are stored as alternative experiments.

Author(s)

Aaron Lun

See Also

[analyze](#), to generate results.

correctMnn

*Batch correction with mutual nearest neighbors***Description**

Apply mutual nearest neighbor (MNN) correction to remove batch effects from a low-dimensional embedding.

Usage

```
correctMnn(
  x,
  block,
  num.neighbors = 15,
  num.steps = 1,
  merge.policy = c("rss", "size", "variance", "input"),
  num.mads = NULL,
  robust.iterations = NULL,
  robust.trim = NULL,
  mass.cap = NULL,
  order = NULL,
  reference.policy = NULL,
  BNPARAM = AnnoyParam(),
  num.threads = 1
)
```

Arguments

x	Numeric matrix where rows are dimensions and columns are cells, typically containing coordinates in a low-dimensional embedding (e.g., from runPca).
block	Factor specifying the block of origin (e.g., batch, sample) for each cell in x.
num.neighbors	Integer scalar specifying the number of neighbors in the various search steps. Larger values improve the stability of the correction by increasing the number of MNN pairs and including more observations in each center of mass. However, this comes at the cost of reduced resolution when matching subpopulations across batches.
num.steps	Integer scalar specifying the number of steps for the recursive neighbor search to compute the center of mass. Larger values mitigate the kissing effect but increase the risk of including inappropriately distant subpopulations into the center of mass.
merge.policy	String specifying the policy to use to choose the order of batches to merge. <ul style="list-style-type: none"> "input" will use the input order of the batches. Observations in the last batch are corrected first, and then the second-last batch, and so on. This allows users to control the merge order by simply changing the inputs. "size" will merge batches in order of increasing size (i.e., the number of observations). So, the smallest batch is corrected first while the largest batch is unchanged. The aim is to lower compute time by reducing the number of observations that need to be reprocessed in later merge steps.

- "variance" will merge batches in order of increasing variance between observations. So, the batch with the lowest variance is corrected first while the batch with the highest variance is unchanged. The aim is to lower compute time by encouraging more observations to be corrected to the most variable batch, thus avoid reprocessing in later merge steps.
- "rss" will merge batches in order of increasing residual sum of squares (RSS). This is effectively a compromise between "variance" and "size".

num.mads	Deprecated and ignored.
robust.iterations	Deprecated and ignored.
robust.trim	Deprecated and ignored.
mass.cap	Deprecated and ignored.
order	Deprecated and ignored, the merge order is now always automatically determined.
reference.policy	Deprecated, use merge.policy instead.
BNPARAM	A BiocNeighborParam object specifying the nearest-neighbor algorithm to use.
num.threads	Integer scalar specifying the number of threads to use.

Value

List containing corrected, a numeric matrix of the same dimensions as x, containing the corrected values.

Author(s)

Aaron Lun

References

Haghverdi L, Lun ATL, Morgan MD, Marioni JC (2018). Batch effects in single-cell RNA-sequencing data are corrected by matching mutual nearest neighbors. *Nat. Biotechnol.* 36(5):421-427

See Also

The compute function in <https://libscrn.github.io/mnncorrect/correctMnn.se>, to perform MNN correction on a [SingleCellExperiment](#).

Examples

```
# Mocking up a dataset with multiple batches.
x <- matrix(rnorm(10000), nrow=10)
b <- sample(3, ncol(x), replace=TRUE)
x[,b==2] <- x[,b==2] + 3
x[,b==3] <- x[,b==3] + 5
lapply(split(colMeans(x), b), mean) # indeed the means differ...

corrected <- correctMnn(x, b)
str(corrected)
lapply(split(colMeans(corrected$corrected), b), mean) # now merged.
```

correctMnn.se

*MNN correction on a SingleCellExperiment***Description**

Correct batch effects from an existing embedding with mutual nearest neighbors (MNNs), by calling `correctMnn` on a reduced dimension entry of a `SingleCellExperiment`.

Usage

```
correctMnn.se(
  x,
  block,
  BNPARAM = AnnoyParam(),
  num.threads = 1,
  more.mnn.args = list(),
  reddim.type = "PCA",
  output.name = "MNN",
  delayed.transpose = FALSE
)
```

Arguments

<code>x</code>	A <code>SingleCellExperiment</code> object or one of its subclasses. Rows correspond to genomic features and columns correspond to cells.
<code>block</code>	Block assignment for each cell, passed to <code>correctMnn</code> .
<code>BNPARAM</code>	Algorithm for the nearest neighbor search, passed to <code>correctMnn</code> .
<code>num.threads</code>	Number of threads, passed to <code>correctMnn</code> .
<code>more.mnn.args</code>	Named list of additional arguments to pass to <code>correctMnn</code> .
<code>reddim.type</code>	String or integer specifying the <code>reducedDim</code> entry on which to perform MNN correction. Alternatively, a named integer or character vector of length 1, where the name specifies an alternative experiment of <code>x</code> and the value is the name/index of a <code>reducedDim</code> entry in that alternative experiment.
<code>output.name</code>	String containing the name of the <code>reducedDim</code> entry in which to store the corrected embedding.
<code>delayed.transpose</code>	Logical scalar indicating whether to delay the transposition when storing coordinates in the <code>reducedDims</code> .

Value

`x` is returned with the corrected embedding stored as a `reducedDim` entry.

Author(s)

Aaron Lun

Examples

```
library(SingleCellExperiment)
sce <- getTestRnaData.se("pca")
# Treating the tissue of origin as the batch.
sce <- correctMnn.se(sce, sce$tissue)
reducedDimNames(sce)
```

countGroupsByBlock *Count cells in groups and blocks*

Description

Tabulate the frequency of cells in each combination of groups and blocks. This is typically used to examine the distribution of cells across batches for each cluster - the presence of a batch-specific cluster may be indicative of a batch effect.

Usage

```
countGroupsByBlock(
  groups,
  block,
  normalize.block = FALSE,
  normalize.groups = FALSE
)
```

Arguments

groups	Factor specifying the group to which each cell was assigned. This is typically used for clusters.
block	Factor specifying the block to which each cell was assigned. This is typically used for batches or samples.
normalize.block	Boolean indicating whether to normalize the number of cells across blocks. If TRUE, frequencies are divided by the column sums.
normalize.groups	Boolean indicating whether to normalize the number of cells across groups. If TRUE, frequencies are divided by the row sums. This is performed after normalization of the block counts if <code>normalize.block=TRUE</code> .

Value

Matrix of (normalized) frequencies. Each row corresponds to a group and each column corresponds to a block.

Author(s)

Aaron Lun

See Also

[table](#), which is used internally by this function.

Examples

```
groups <- sample(10, 100, replace=TRUE)
block <- sample(LETTERS[1:6], 100, replace=TRUE)

countGroupsByBlock(groups, block)
countGroupsByBlock(groups, block, normalize.block=TRUE)
countGroupsByBlock(groups, block, normalize.groups=TRUE)
countGroupsByBlock(groups, block, normalize.block=TRUE, normalize.groups=TRUE)
```

crispr_quality_control

Quality control for CRISPR count data

Description

Compute per-cell QC metrics from an initialized matrix of CRISPR counts, and use the metrics to suggest filter thresholds to retain high-quality cells.

Usage

```
computeCrisprQcMetrics(x, num.threads = 1)

suggestCrisprQcThresholds(metrics, block = NULL, num.mads = 3)

filterCrisprQcMetrics(thresholds, metrics, block = NULL)
```

Arguments

x	A matrix-like object where rows are CRISPRs and columns are cells. Values are expected to be counts.
num.threads	Integer scalar specifying the number of threads to use.
metrics	DataFrame of per-cell QC metrics. This should have the same structure as the return value of <code>computeCrisprQcMetrics</code> .
block	Factor specifying the block of origin (e.g., batch, sample) for each cell in <code>metrics</code> . Alternatively NULL if all cells are from the same block. For <code>filterCrisprQcMetrics</code> , a blocking factor should be provided if <code>block</code> was used to construct thresholds.
num.mads	Number of median from the median, to define the threshold for outliers in each metric.
thresholds	List with the same structure as produced by <code>suggestCrisprQcThresholds</code> .

Details

In CRISPR data, a cell is considered to be of low quality if it has a low count for its most abundant guide. However, directly defining a MAD-based outlier threshold on the maximum count is somewhat tricky as unsuccessful transfection can be common. This often results in a large subpopulation with low maximum counts, inflating the MAD and compromising the threshold calculation. Instead, we use the following approach:

- Compute the proportion of counts in the most abundant guide (i.e., the maximum proportion) in each cell. Cells that were successfully transfected should have high maximum proportions. In contrast, unsuccessfully transfected cells will be dominated by ambient contamination and have low proportions.
- Subset the dataset to only retain those cells with maximum proportions above the median. This assumes that at least 50 Thus, we remove all of the unsuccessful transfections and enrich for mostly-high-quality cells.
- Define a MAD-based threshold for low outliers on the log-transformed maximum count within the subset (see ‘choose_filter_thresholds()’ for details). This is now possible as we can assume that most of the remaining cells are of high quality.

Note that the maximum proportion is only used to define the subset for threshold calculation. Once the maximum count threshold is computed, it is applied to all cells regardless of their maximum proportions. This ensures that we correctly remove cells with low coverage, even if the proportion is high. It also allows us to retain cells transfected with multiple guides, as long as the maximum is high enough - such cells are not necessarily uninteresting, e.g., for examining interaction effects, so we will err on the side of caution and leave them in.

Value

For `computeCrisprQcMetrics`, a [DataFrame](#) is returned with one row per cell in `x`. This contains the following columns:

- `sum`, a numeric vector containing the total CRISPR count for each cell. Low counts indicate that the cell was not successfully transfected with a construct or that library preparation and sequencing failed.
- `detected`, an integer vector containing the number of detected guides per cell. In theory, this should be 1, as each cell should express no more than one guide construct. However, ambient contamination may introduce non-zero counts for multiple guides, without necessarily interfering with downstream analyses. As such, this metric is less useful for guide data, though we compute it anyway.
- `max.value`, a numeric vector containing the count for the most abundant guide in cell. Low values indicate that the cell was not successfully transfected or that library preparation and sequencing failed.
- `max.index`, an integer vector containing the row index of the most abundant guide in cell.

Each vector is of length equal to the number of cells.

For `suggestCrisprQcThresholds`, a named list is returned.

- If `block=NULL`, the list contains:
 - `max.value`, a numeric scalar containing the lower bound on the maximum count. This is defined as `num.mads` MADs below the median of the log-transformed metrics across cells with high maximum proportions (see [Details](#)).
- Otherwise, if `block` is supplied, the list contains:
 - `max.value`, a numeric vector containing the lower bound on the maximum counts for each blocking level. Here, the threshold is computed independently for each block, using the same method as the unblocked case.
 - `block.ids`, a vector containing the identities of the unique blocks.

Each vector is of length equal to the number of levels in `block` and is named accordingly.

For `filterCrisprQcMetrics`, a logical vector of length `ncol(x)` is returned indicating which cells are of high quality. High-quality cells are defined as those with maximum counts above the `max.value` threshold.

Author(s)

Aaron Lun

See Also

The `compute_crispr_qc_metrics`, `compute_crispr_qc_filters` and `compute_crispr_qc_filters_blocked` functions in https://libscran.github.io/scran_qc/.

[quickCrisprQc.se](#), to run all of the CRISPR-related QC functions on a [SummarizedExperiment](#).

Examples

```
# Mocking a matrix:
library(Matrix)
x <- round(abs(rsparsematrix(100, 100, 0.1) * 100))

qc <- computeCrisprQcMetrics(x)
qc

filt <- suggestCrisprQcThresholds(qc)
str(filt)

keep <- filterCrisprQcMetrics(filt, qc)
summary(keep)
```

`fitVarianceTrend`*Fit a mean-variance trend*

Description

Fit a trend to the per-gene variances with respect to their means, typically from normalized and log-transformed expression values.

Usage

```
fitVarianceTrend(
  means,
  variances,
  mean.filter = TRUE,
  min.mean = 0.1,
  transform = TRUE,
  span = 0.3,
  use.min.width = FALSE,
  min.width = 1,
  min.window.count = 200,
  num.threads = 1
)
```

Arguments

means	Numeric vector containing the mean (log-)expression for each gene.
variances	Numeric vector containing the variance in the (log-)expression for each gene.
mean.filter	Logical scalar indicating whether to filter on the means before trend fitting. The assumption is that there is a bulk of low-abundance genes that are uninteresting and should be removed to avoid skewing the windows of the LOWESS smoother.
min.mean	Numeric scalar specifying the minimum mean of genes to use in trend fitting. Genes with lower means do not participate in the LOWESS fit, to ensure that windows are not skewed towards the majority of low-abundance genes. Instead, the fitted values for these genes are defined by extrapolating the left edge of the fitted trend is extrapolated to the origin. The default value is chosen based on the typical distribution of means of log-expression values across genes. Only used if mean.filter=TRUE.
transform	Logical scalar indicating whether a quarter-root transformation should be applied before trend fitting. This transformation is copied from <code>limma::voom</code> and shrinks all values towards 1, flattening any sharp gradients in the trend for an easier fit. The default of TRUE assumes that the variances are computed from log-expression values, in which case there is typically a strong “hump” in the mean-variance relationship.
span	Numeric scalar specifying the span of the LOWESS smoother, as a proportion of the total number of points. Larger values improve stability at the cost of sensitivity to changes in low-density regions. Ignored if use.min.width=TRUE.
use.min.width	Logical scalar indicating whether a minimum width constraint should be applied to the LOWESS smoother. This replaces the proportion-based span for defining each window. Instead, the window for each point must be of a minimum width and is extended until it contains a minimum number of points. Setting this to ‘TRUE’ ensures that sensitivity is maintained in the trend fit at low-density regions for the distribution of means, e.g., at high abundances. It also avoids overfitting from very small windows in high-density intervals.
min.width	Minimum width of the window to use when use.min.width=TRUE. The default value is chosen based on the typical range of means in single-cell RNA-seq data.
min.window.count	Minimum number of observations in each window. This ensures that each window contains at least a given number of observations for a stable fit. If the minimum width window contains fewer observations, it is extended using the standard LOWESS logic until the minimum number is achieved. Only used if use.min.width=TRUE.
num.threads	Number of threads to use.

Value

List containing `fitted`, a numeric vector containing the fitted values of the trend for each gene; and `residuals`, a numeric vector containing the residuals from the trend.

Author(s)

Aaron Lun

See Also

[modelGeneVariances](#), to compute the means and variances on which the trend is fitted.

The `fit_variance_trend` function in https://libscrان.github.io/scrان_variances/.

Examples

```
# Setting up some single-cell-like data.
mu <- 2^runif(1000, -10, 10)
counts <- matrix(rpois(20 * length(mu), lambda=mu), ncol=20)

sf <- centerSizeFactors(colSums(counts))
normalized <- normalizeCounts(counts, size.factors=sf)
stats <- modelGeneVariances(normalized)

out <- fitVarianceTrend(stats$statistics$means, stats$statistics$variances)
plot(stats$statistics$means, stats$statistics$variances)
curve(approxfun(stats$statistics$means, out$fitted)(x), col="red", add=TRUE)
```

getTestData.se

Get datasets for testing

Description

Get single-cell datasets from the **scRNAseq** package with varying levels of processing. This is primarily intended for testing other **scraper** functions, e.g., in their Examples section.

Usage

```
getTestRnaData.se(at = c("start", "qc", "norm", "hvg", "pca", "cluster"))

getTestAdtData.se(at = c("start", "qc", "norm", "hvg", "pca"))

getTestCrisprData.se(at = c("start", "qc"))
```

Arguments

<code>at</code>	String specifying the level of processing. For "start", no processing was performed. Otherwise, the dataset is returned after quality control ("qc"), normalization ("norm"), feature selection ("hvg"), PCA ("pca") or graph-based clustering ("cluster").
-----------------	---

Details

For `getTestRnaData`, this is a scRNA-seq dataset of the mouse brain, where the main experiment contains RNA counts and the alternative experiments contain ERCC and repeat element counts. This is obtained with `fetchDataset("zeisel-brain-2015", "2023-12-14")`.

For `getTestAdtData`, this is a CITE-seq dataset of human PBMCs, where the main experiment contains RNA counts and the alternative experiment contains ADT counts. This is obtained with `fetchDataset("kotliarov-pbmc-2020", "2024-04-18")`. Only the first 5000 cells are loaded for speed.

For `getTestCrisprData`, this is a Perturb-seq dataset of a pancreatic beta cell line, where the main experiment contains RNA counts and the alternative experiment contains CRISPR guide counts. This is obtained with `fetchDataset("cao-pancreas-2025", "2025-10-10", "rqc")`. Only the first 5000 cells are loaded for speed.

Value

A `SingleCellExperiment` containing a dataset at the specified level of processing.

Author(s)

Aaron Lun

See Also

`fetchDataset`, used to obtain each dataset.

Examples

```
getTestRnaData.se()
getTestAdtData.se()
getTestCrisprData.se()
```

LogNormalizedMatrix *Delayed log-normalization of a matrix*

Description

Delayed calculation of log-normalized expression values, typically returned by `normalizeCounts`.

Usage

```
LogNormalizedMatrix(x, size.factors, pseudo.count = 1, log.base = 2)
```

```
LogNormalizedMatrixSeed(x, size.factors, pseudo.count = 1, log.base = 2)
```

Arguments

<code>x</code>	Count matrix to be normalized.
<code>size.factors</code>	Numeric vector of size factors, of length equal to the number of columns of <code>x</code> .
<code>pseudo.count</code>	Number specifying the pseudo-count to add prior to log-transformation.
<code>log.base</code>	Number specifying the base of the log-transformation.

Details

This is based on the **DelayedArray** framework and

Value

An instance of a `LogNormalizedMatrix(Seed)`.

Author(s)

Aaron Lun

Examples

```
mat <- matrix(rpois(1000, lambda=2), ncol=10)
sf <- centerSizeFactors(colSums(mat))
norm <- LogNormalizedMatrix(mat, sf)
norm

# Also works with sparse matrices.
library(Matrix)
smat <- abs(rsparsematrix(50, 20, density=0.21))
ssf <- centerSizeFactors(colSums(smat))
snorm <- LogNormalizedMatrix(smat, ssf)
snorm
```

modelGeneVariances *Model per-gene variances in expression*

Description

Model the per-gene variances as a function of the mean in single-cell expression data. Highly variable genes can then be selected for downstream analyses.

Usage

```
modelGeneVariances(
  x,
  block = NULL,
  block.average.policy = c("mean", "quantile"),
  block.weight.policy = c("variable", "equal", "none"),
  variable.block.weight = c(0, 1000),
  block.quantile = 0.5,
  mean.filter = TRUE,
  min.mean = 0.1,
  transform = TRUE,
  span = 0.3,
  use.min.width = FALSE,
  min.width = 1,
  min.window.count = 200,
  num.threads = 1
)
```

Arguments

x A matrix-like object where rows correspond to genes or genomic features and columns correspond to cells. It is typically expected to contain log-expression values, e.g., from [normalizeCounts](#).

<code>block</code>	Factor specifying the block of origin (e.g., batch, sample) for each cell in <code>x</code> . If provided, calculation of means/variances and trend fitting are performed within each block to ensure that block effects do not confound the estimates. The weighted average of each statistic across all blocks is reported for each gene. Alternatively NULL, if all cells are from the same block.
<code>block.average.policy</code>	String specifying the policy to use for average statistics across blocks. This can either be a (weighted) "mean" or a "quantile". Only used if <code>block</code> is not NULL.
<code>block.weight.policy</code>	String specifying the policy to use for weighting different blocks when computing the average for each statistic. See the argument of the same name in computeBlockWeights for more detail. Only used if <code>block</code> is not NULL and <code>block.average.policy="mean"</code> .
<code>variable.block.weight</code>	Numeric vector of length 2, specifying the parameters for variable block weighting. See the argument of the same name in computeBlockWeights for more detail. Only used if <code>block</code> is not NULL, <code>block.average.policy="mean"</code> and <code>block.weight.policy="variable"</code> .
<code>block.quantile</code>	Number specifying the probability of the quantile of statistics across blocks. Defaults to 0.5, i.e., the median of per-block statistics. Only used if <code>block</code> is not NULL and <code>block.average.policy="quantile"</code> .
<code>mean.filter</code>	Logical scalar indicating whether to filter on the means before trend fitting. The assumption is that there is a bulk of low-abundance genes that are uninteresting and should be removed to avoid skewing the windows of the LOWESS smoother.
<code>min.mean</code>	Numeric scalar specifying the minimum mean of genes to use in trend fitting. Genes with lower means do not participate in the LOWESS fit, to ensure that windows are not skewed towards the majority of low-abundance genes. Instead, the fitted values for these genes are defined by extrapolating the left edge of the fitted trend is extrapolated to the origin. The default value is chosen based on the typical distribution of means of log-expression values across genes. Only used if <code>mean.filter=TRUE</code> .
<code>transform</code>	Logical scalar indicating whether a quarter-root transformation should be applied before trend fitting. This transformation is copied from <code>limma::voom</code> and shrinks all values towards 1, flattening any sharp gradients in the trend for an easier fit. The default of TRUE assumes that the variances are computed from log-expression values, in which case there is typically a strong "hump" in the mean-variance relationship.
<code>span</code>	Numeric scalar specifying the span of the LOWESS smoother, as a proportion of the total number of points. Larger values improve stability at the cost of sensitivity to changes in low-density regions. Ignored if <code>use.min.width=TRUE</code> .
<code>use.min.width</code>	Logical scalar indicating whether a minimum width constraint should be applied to the LOWESS smoother. This replaces the proportion-based span for defining each window. Instead, the window for each point must be of a minimum width and is extended until it contains a minimum number of points. Setting this to 'TRUE' ensures that sensitivity is maintained in the trend fit at low-density regions for the distribution of means, e.g., at high abundances. It also avoids overfitting from very small windows in high-density intervals.
<code>min.width</code>	Minimum width of the window to use when <code>use.min.width=TRUE</code> . The default value is chosen based on the typical range of means in single-cell RNA-seq data.

<code>min.window.count</code>	Minimum number of observations in each window. This ensures that each window contains at least a given number of observations for a stable fit. If the minimum width window contains fewer observations, it is extended using the standard LOWESS logic until the minimum number is achieved. Only used if <code>use.min.width=TRUE</code> .
<code>num.threads</code>	Integer scalar specifying the number of threads to use.

Details

We compute the mean and variance for each gene and fit a trend to the variances with respect to the means using `fitVarianceTrend`. We assume that most genes at any given abundance are not highly variable, such that the fitted value of the trend is interpreted as the “uninteresting” variance - this is mostly attributed to technical variation like sequencing noise, but can also represent constitutive biological noise like transcriptional bursting. Under this assumption, the residual can be treated as a measure of biologically interesting variation. Genes with large residuals can then be selected for downstream analyses, e.g., with `chooseHighlyVariableGenes`.

Value

A list containing statistics, a `DataFrame` with number of rows equal to the number of genes. This contains the columns `means`, `variances`, `fitted` and `residuals`, each of which is a numeric vector containing the statistic of the same name across all genes.

If `block` is supplied, each of the column vectors described above contains the average across all blocks. The list will also contain `per.block`, a list of `DataFrames` containing the equivalent statistics for each block; and `block.ids`, a vector containing the identities of the unique blocks in the same order as `per.block`.

Author(s)

Aaron Lun

See Also

The `model_gene_variances` function in https://libscran.github.io/scran_variances/. [chooseRnaHvgs.se](#), which computes the variances and trend from a `SummarizedExperiment`.

Examples

```
library(Matrix)
x <- abs(rsparsematrix(1000, 100, 0.1) * 10)
out <- modelGeneVariances(x)
out

# Throwing in some blocking.
block <- sample(letters[1:4], ncol(x), replace=TRUE)
out <- modelGeneVariances(x, block=block)
out
```

normalizeAdtCounts.se *Normalize ADT counts in a SummarizedExperiment*

Description

Compute (log-)normalized expression values after performing scaling normalization of an ADT count matrix. This calls `computeClrm1Factors` on an assay of a `SingleCellExperiment`, centering the subsequent size factors with `centerSizeFactors`, and then computing normalized log-expression values with `normalizeCounts`.

Usage

```
normalizeAdtCounts.se(
  x,
  size.factors = NULL,
  num.threads = 1,
  center = TRUE,
  block = NULL,
  mode = "lowest",
  log = TRUE,
  pseudo.count = 1,
  more.norm.args = list(),
  assay.type = "counts",
  output.name = "logcounts",
  factor.name = "sizeFactor"
)
```

Arguments

<code>x</code>	A <code>SummarizedExperiment</code> object or one of its subclasses. Rows correspond to antibody-derived tags (ADTs) and columns correspond to cells.
<code>size.factors</code>	Numeric vector of length equal to the number of columns of <code>x</code> , containing the size factor for each cell in <code>x</code> . If <code>NULL</code> , this defaults to the output of <code>computeClrm1Factors</code> .
<code>num.threads</code>	Number of threads, passed to <code>computeClrm1Factors</code> .
<code>center</code>	Logical scalar indicating whether to center the <code>size.factors</code> , see <code>?centerSizeFactors</code> for more details.
<code>block</code>	Block assignments for each cell, passed to <code>centerSizeFactors</code> .
<code>mode</code>	How to center size factors in different blocks, see <code>?centerSizeFactors</code> for more details.
<code>log</code>	Whether to log-transform the normalized expression values, see <code>?normalizeCounts</code> for more details.
<code>pseudo.count</code>	The pseudo-count for log-transformation, see <code>?normalizeCounts</code> for more details.
<code>more.norm.args</code>	Named list of additional arguments to pass to <code>normalizeCounts</code> .
<code>assay.type</code>	Integer or string specifying the assay of <code>x</code> with the count matrix.
<code>output.name</code>	String containing the name of the assay to store the normalized matrix.
<code>factor.name</code>	String containing the name of the <code>colData</code> column in which to store the size factors in the output object. If <code>NULL</code> , the size factors are not stored.

Value

x is returned with a new assay containing the (log-)normalized matrix. Size factors are also stored in the `colData`.

Author(s)

Aaron Lun

Examples

```
library(SingleCellExperiment)
sce <- altExp(getTestAdtData.se("qc"), "ADT")
sce <- normalizeAdtCounts.se(sce)
assayNames(sce)
summary(sizeFactors(sce))
```

normalizeCounts	<i>Normalize the count matrix</i>
-----------------	-----------------------------------

Description

Apply scaling normalization and log-transformation to a count matrix. This yields normalized expression values that can be used in downstream procedures like PCA.

Usage

```
normalizeCounts(
  x,
  size.factors,
  log = TRUE,
  pseudo.count = 1,
  log.base = 2,
  preserve.sparsity = FALSE,
  delayed = TRUE
)
```

Arguments

- | | |
|--------------|--|
| x | A matrix-like object where rows correspond to genes or genomic features and columns correspond to cells. Values are expected to be non-negative counts. Alternatively, an external pointer created by initializeCpp . |
| size.factors | A numeric vector of length equal to the number of cells in x, containing positive size factors for all cells. Any invalid values should be replaced with sanitizeSizeFactors . For most applications, these size factors should also be centered with centerSizeFactors . |
| log | Logical scalar indicating whether log-transformation should be performed. This ensures that downstream analyses like t-tests and distance calculations focus on relative fold-changes rather than absolute differences. The log-transformation also provides some measure of variance stabilization so that the downstream analyses are not dominated by sampling noise at large counts. |

pseudo.count	Numeric scalar specifying the positive pseudo-count to add before any log-transformation. Larger values shrink the differences between cells towards zero, reducing spurious differences (but also signal) at low counts - see choosePseudoCount for comments. Ignored if log=FALSE.
log.base	Numeric scalar specifying the base of the log-transformation. Ignored if log=FALSE.
preserve.sparsity	Logical scalar indicating whether to preserve sparsity for pseudo.count != 1. If TRUE, users should manually add $\log(\text{pseudo.count}, \text{log.base})$ to the returned matrix to obtain the desired log-transformed expression values. Ignored if log = FALSE or pseudo.count = 1.
delayed	Logical scalar indicating whether operations on a matrix-like x should be delayed. This improves memory efficiency at the cost of some speed in downstream operations.

Value

If x is a matrix-like object and delayed=TRUE, a [DelayedArray](#) is returned containing the (log-transformed) normalized expression matrix. If delayed=FALSE, the type of the (log-)normalized matrix will depend on the operations applied to x.

If x is an external pointer produced by [initializeCpp](#), a new external pointer is returned containing the normalized expression matrix.

Author(s)

Aaron Lun

See Also

The `normalize_counts` function in https://libscran.github.io/scran_norm/.

[normalizeRnaCounts.se](#) and related functions, which compute normalized expression values from a [SummarizedExperiment](#).

Examples

```
# Mocking a matrix:
library(Matrix)
x <- round(abs(rsparsematrix(1000, 100, 0.1) * 100))
sf <- centerSizeFactors(colSums(x))
normed <- normalizeCounts(x, size.factors=sf)
normed

# Passing a pointer.
ptr <- beachmat::initializeCpp(x)
optr <- normalizeCounts(ptr, sf)
optr
```

```
normalizeCrisprCounts.se
```

Normalize CRISPR counts in a SummarizedExperiment

Description

Compute (log-)normalized expression values after performing scaling normalization of an CRISPR count matrix. This calls `normalizeCounts` on an assay of a `SummarizedExperiment`, after centering the size factors with `centerSizeFactors`.

Usage

```
normalizeCrisprCounts.se(
  x,
  size.factors = NULL,
  center = TRUE,
  block = NULL,
  mode = "lowest",
  log = TRUE,
  pseudo.count = 1,
  more.norm.args = list(),
  assay.type = "counts",
  output.name = "logcounts",
  factor.name = "sizeFactor",
  num.threads = 1
)
```

Arguments

<code>x</code>	A <code>SummarizedExperiment</code> object or one of its subclasses. Rows correspond to CRISPR guides and columns correspond to cells.
<code>size.factors</code>	Numeric vector of length equal to the number of columns of <code>x</code> , containing the size factor for each cell in <code>x</code> . If <code>NULL</code> , this defaults to the column sums of the count matrix in <code>x</code> .
<code>center</code>	Logical scalar indicating whether to center the <code>size.factors</code> , see <code>?centerSizeFactors</code> for more details.
<code>block</code>	Block assignments for each cell, passed to <code>centerSizeFactors</code> .
<code>mode</code>	How to center size factors in different blocks, see <code>?centerSizeFactors</code> for more details.
<code>log</code>	Whether to log-transform the normalized expression values, see <code>?normalizeCounts</code> for more details.
<code>pseudo.count</code>	The pseudo-count for log-transformation, see <code>?normalizeCounts</code> for more details.
<code>more.norm.args</code>	Named list of additional arguments to pass to <code>normalizeCounts</code> .
<code>assay.type</code>	Integer or string specifying the assay of <code>x</code> with the count matrix.
<code>output.name</code>	String containing the name of the assay to store the normalized matrix.
<code>factor.name</code>	String containing the name of the <code>colData</code> column in which to store the size factors in the output object. If <code>NULL</code> , the size factors are not stored.
<code>num.threads</code>	Integer specifying the number of threads for computing column sums. Only used if <code>size.factors=NULL</code> .

Value

x is returned with a new assay containing the (log-)normalized matrix. Size factors are also stored in the `colData`.

Author(s)

Aaron Lun

Examples

```
library(SingleCellExperiment)
sce <- altExp(getTestCrisprData.se("qc"), "CRISPR Guide Capture")
sce <- normalizeCrisprCounts.se(sce, size.factors=sce$sum)
assayNames(sce)
summary(sizeFactors(sce))
```

normalizeRnaCounts.se *Normalize RNA counts in a SummarizedExperiment*

Description

Compute (log-)normalized expression values after performing scaling normalization of an RNA count matrix. This calls `normalizeCounts` on an assay of a `SummarizedExperiment`, after centering the size factors with `centerSizeFactors`.

Usage

```
normalizeRnaCounts.se(  
  x,  
  size.factors = NULL,  
  center = TRUE,  
  block = NULL,  
  mode = "lowest",  
  log = TRUE,  
  pseudo.count = 1,  
  more.norm.args = list(),  
  assay.type = "counts",  
  output.name = "logcounts",  
  factor.name = "sizeFactor",  
  num.threads = 1  
)
```

Arguments

- | | |
|--------------|--|
| x | A <code>SummarizedExperiment</code> object or one of its subclasses. Rows correspond to genes and columns correspond to cells. |
| size.factors | Numeric vector of length equal to the number of columns of x, containing the size factor for each cell in x. If NULL, this defaults to the column sums of the count matrix in x. |

center	Logical scalar indicating whether to center the size.factors, see ?centerSizeFactors for more details.
block	Block assignments for each cell, passed to centerSizeFactors .
mode	How to center size factors in different blocks, see ?centerSizeFactors for more details.
log	Whether to log-transform the normalized expression values, see ?normalizeCounts for more details.
pseudo.count	The pseudo-count for log-transformation, see ?normalizeCounts for more details.
more.norm.args	Named list of additional arguments to pass to normalizeCounts .
assay.type	Integer or string specifying the assay of x with the count matrix.
output.name	String containing the name of the assay to store the normalized matrix.
factor.name	String containing the name of the colData column in which to store the size factors in the output object. If NULL, the size factors are not stored.
num.threads	Integer specifying the number of threads for computing column sums. Only used if size.factors=NULL.

Value

x is returned with a new assay containing the (log-)normalized matrix. Size factors are also stored in the [colData](#).

Author(s)

Aaron Lun

Examples

```
library(SingleCellExperiment)
sce <- getTestRnaData.se("qc")
sce <- normalizeRnaCounts.se(sce, size.factors=sce$sum)
assayNames(sce)
summary(sizeFactors(sce))
```

quickAdtQc.se

Quick quality control for ADT data in a SummarizedExperiment

Description

Quickly compute quality control (QC) metrics, thresholds and filters from ADT data in a [SummarizedExperiment](#). This calls [computeAdtQcMetrics](#) on an assay in a SummarizedExperiment, followed by [suggestAdtQcThresholds](#) and [filterAdtQcMetrics](#) to identify high-quality cells.

Usage

```

quickAdtQc.se(
  x,
  subsets,
  num.threads = 1,
  thresholds = NULL,
  block = NULL,
  more.suggest.args = list(),
  assay.type = "counts",
  output.prefix = NULL,
  meta.name = "qc",
  flatten = TRUE
)

formatComputeAdtQcMetricsResult(compute.res, flatten = TRUE)

```

Arguments

x	A SummarizedExperiment object or one of its subclasses. Rows correspond to antibody-derived tags (ADTs) and columns correspond to cells.
subsets	List of subsets of control tags, see ?computeAdtQcMetrics for more details.
num.threads	Number of threads, to pass to computeAdtQcMetrics .
thresholds	List containing pre-defined thresholds for each QC metric, see the return value of suggestAdtQcThresholds for the expected format.
block	Block assignment for each cell, to pass to suggestAdtQcThresholds and filterAdtQcMetrics .
more.suggest.args	Named list of additional arguments to pass to suggestAdtQcThresholds .
assay.type	Integer or string specifying the assay of x containing the ADT count matrix.
output.prefix	String containing a prefix to add to the names of the <code>link[SummarizedExperiment]{colData}</code> columns containing the output statistics.
meta.name	String containing the name of the metadata entry containing additional outputs like the filtering thresholds. If NULL, additional outputs are not reported.
flatten	Logical scalar indicating whether to flatten the subset proportions into separate columns of the <code>link[SummarizedExperiment]{colData}</code> . If FALSE, the subset proportions are stored in a nested DataFrame .
compute.res	DataFrame returned by computeAdtQcMetrics .

Value

For `quickAdtQc.se`, x is returned with additional columns added to its `colData`. Each column contains per-cell values for one of the QC metrics, see [computeAdtQcMetrics](#) for details. The suggested thresholds are stored as a list in [metadata](#). The `colData` also contains a `keep` column, specifying which cells are to be retained.

For `formatComputeAdtQcMetricsResult`, a [DataFrame](#) is returned with the per-cell QC metrics.

Author(s)

Aaron Lun

Examples

```
library(SingleCellExperiment)
sce <- altExp(getTestAdtData.se(), "ADT")
sce <- quickAdtQc.se(sce, subsets=list(igg=grepl("IgG", rownames(sce))))
colData(sce)[,c("sum", "detected", "subset.sum.igg")]
metadata(sce)$qc$thresholds
summary(sce$keep)
```

quickCrisprQc.se

Quick quality control for CRISPR data in a SummarizedExperiment

Description

Quickly compute quality control (QC) metrics, thresholds and filters from CRISPR data in a [SummarizedExperiment](#). This calls [computeCrisprQcMetrics](#) on an assay in a [SummarizedExperiment](#), followed by [suggestCrisprQcThresholds](#) and [filterCrisprQcMetrics](#) to identify high-quality cells.

Usage

```
quickCrisprQc.se(
  x,
  num.threads = 1,
  thresholds = NULL,
  block = NULL,
  more.suggest.args = list(),
  assay.type = "counts",
  output.prefix = NULL,
  meta.name = "qc"
)

formatComputeCrisprQcMetricsResult(compute.res)
```

Arguments

x	A SummarizedExperiment object or one of its subclasses. Rows correspond to CRISPR guides and columns correspond to cells.
num.threads	Number of threads, to pass to computeCrisprQcMetrics .
thresholds	List containing pre-defined thresholds for each QC metric, see the return value of suggestRnaQcThresholds for the expected format.
block	Block assignment for each cell, to pass to suggestCrisprQcThresholds and filterCrisprQcMetrics .
more.suggest.args	Named list of additional arguments to pass to suggestCrisprQcThresholds .
assay.type	Integer or string specifying the assay of x containing the CRISPR count matrix.
output.prefix	String containing a prefix to add to the names of the <code>link[SummarizedExperiment]{colData}</code> columns containing the output statistics.
meta.name	String containing the name of the metadata entry containing additional outputs like the filtering thresholds. If NULL, additional outputs are not reported.
compute.res	DataFrame returned by computeCrisprQcMetrics .

Value

For `quickCrisprQc.se`, `x` is returned with additional columns added to its `colData`. Each column contains per-cell values for one of the QC metrics, see `computeCrisprQcMetrics` for details. The suggested thresholds are stored as a list in `metadata`. The `colData` also contains a `keep` column, specifying which cells are to be retained.

For `formatComputeCrisprQcMetricsResult`, a `DataFrame` is returned with the per-cell QC metrics.

Author(s)

Aaron Lun

Examples

```
library(SingleCellExperiment)
sce <- altExp(getTestCrisprData.se(), "CRISPR Guide Capture")
sce <- quickCrisprQc.se(sce)
colData(sce)[,c("sum", "detected", "max.value", "max.index")]
metadata(sce)$qc$thresholds
summary(sce$keep)
```

quickRnaQc.se

Quick quality control for RNA data in a SummarizedExperiment

Description

Quickly compute quality control (QC) metrics, thresholds and filters from RNA data in a `SummarizedExperiment`. This calls `computeRnaQcMetrics` on an assay in a `SummarizedExperiment`, followed by `suggestRnaQcThresholds` and `filterRnaQcMetrics` to identify high-quality cells.

Usage

```
quickRnaQc.se(
  x,
  subsets,
  num.threads = 1,
  thresholds = NULL,
  block = NULL,
  more.suggest.args = list(),
  altexp.proportions = NULL,
  assay.type = "counts",
  output.prefix = NULL,
  meta.name = "qc",
  flatten = TRUE
)

computeRnaQcMetricsWithAltExps(
  x,
  subsets,
  altexp.proportions,
```

```

num.threads = 1,
assay.type = "counts"
)

formatComputeRnaQcMetricsResult(compute.res, flatten = TRUE)

```

Arguments

x	A SummarizedExperiment object or one of its subclasses. Rows correspond to genes and columns correspond to cells.
subsets	List of subsets of control genes, see ?computeRnaQcMetrics for more details.
num.threads	Number of threads, to pass to computeRnaQcMetrics .
thresholds	List containing pre-defined thresholds for each QC metric, see the return value of suggestRnaQcThresholds for the expected format.
block	Block assignment for each cell, to pass to suggestRnaQcThresholds and filterRnaQcMetrics .
more.suggest.args	Named list of additional arguments to pass to suggestRnaQcThresholds .
altexp.proportions	<p>Alternative experiments for which to compute QC metrics. This is typically used to refer to alternative experiments holding spike-in data. For each alternative experiment, the proportion is defined as $X/(X + Y)$ where X is the alternative experiment's total and Y is the RNA total. These proportions will be used for filtering in the same manner as the proportions computed from subsets.</p> <p>More specifically, <code>altexp.proportions</code> should be an unnamed integer or character vector containing the names/indices of the alternative experiments of interest. The assay to use from each alternative experiment is determined by <code>assay.type</code>.</p> <p>Alternatively, <code>altexp.proportions</code> may be a named integer or character vector. Each name specifies an alternative experiment while each value is the index/name of the assay to use from that experiment.</p> <p>Only relevant if <code>x</code> is a SingleCellExperiment.</p>
assay.type	Integer or string specifying the assay of <code>x</code> containing the RNA count matrix.
output.prefix	String containing a prefix to add to the names of the <code>link[SummarizedExperiment]{colData}</code> columns containing the output statistics.
meta.name	String containing the name of the <code>metadata</code> entry containing the additional outputs such as the filtering thresholds. If NULL, additional outputs are not reported.
flatten	Logical scalar indicating whether to flatten the subset proportions into separate columns of the <code>link[SummarizedExperiment]{colData}</code> . If FALSE, the subset proportions are stored in a nested DataFrame .
compute.res	DataFrame returned by computeRnaQcMetrics .

Value

For `quickRnaQc.se`, `x` is returned with additional columns added to its `colData`. Each column contains per-cell values for one of the QC metrics, see [computeRnaQcMetrics](#) for details. The suggested thresholds are stored as a list in `metadata`. The `colData` also contains a `keep` column, specifying which cells are to be retained. If `altexp.proportions` is provided, QC metrics are added to the `colData` of the specified alternative experiments in the output object.

For `computeRnaQcMetricsWithAltExps`, a list is returned containing:

- `main`, the result of calling `computeRnaQcMetrics` on the RNA count matrix in `x`. The proportion of counts in each alternative experiment is added to the subsets.
- `altexp`, a named list of length equal to `altexp.proportions`. Each inner list is the result of calling `computeRnaQcMetrics` on the RNA count matrix of the corresponding alternative experiment of `x`.

For `formatComputeRnaQcMetricsResult`, a `DataFrame` is returned containing the per-cell QC metrics.

Author(s)

Aaron Lun

Examples

```
library(SingleCellExperiment)
sce <- getTestRnaData.se()
sce <- quickRnaQc.se(sce, subsets=list(mito=grepl("^mt", rownames(sce))))
colData(sce)[,c("sum", "detected", "subset.proportion.mito")]
metadata(sce)$qc$thresholds
summary(sce$keep)

# Computing spike-in proportions, if available.
sce <- getTestRnaData.se()
sce <- quickRnaQc.se(
  sce,
  subsets=list(mito=grepl("^mt", rownames(sce))),
  altexp.proportions="ERCC"
)
colData(sce)[,c("sum", "detected", "subset.proportion.mito", "subset.proportion.ERCC")]
colData(altExp(sce, "ERCC"))[,c("sum", "detected")]
```

reportGroupMarkerStatistics

Report marker statistics for a single group

Description

Combine all marker statistics for a single group into a data frame for easy inspection. Users can pick one of the columns for sorting potential marker genes.

Usage

```
reportGroupMarkerStatistics(
  results,
  group,
  effect.sizes = NULL,
  summaries = NULL,
  include.mean = TRUE,
  include.detected = TRUE
)
```

Arguments

<code>results</code>	Named list of marker statistics, typically generated by <code>scoreMarkers</code> with all <code>.pairwise=FALSE</code> .
<code>group</code>	String or integer scalar specifying the group of interest.
<code>effect.sizes</code>	Character vector specifying the effect sizes of interest. If NULL, all effect sizes are reported in the returned data frame.
<code>summaries</code>	Character vector specifying the summary statistics of interest. If NULL, all summaries are reported in the returned data frame.
<code>include.mean</code>	Logical scalar indicating whether the mean expression should be reported in the returned data frame.
<code>include.detected</code>	Logical scalar indicating whether the proportion of detected cells should be reported in the returned data frame.

Value

Data frame where each row corresponds to a gene. Each column contains the requested statistics for group. Effect size summary columns are named as `<EFFECT>.<SUMMARY>`.

Author(s)

Aaron Lun

See Also

[scoreMarkers](#), to generate results.

[summarizeEffects](#), for the trade-offs between effect size summaries.

rna_quality_control *Quality control for RNA count data*

Description

Compute per-cell QC metrics from an initialized matrix of RNA counts, and use the metrics to suggest filter thresholds to retain high-quality cells.

Usage

```
computeRnaQcMetrics(x, subsets, num.threads = 1)

suggestRnaQcThresholds(metrics, block = NULL, num.mads = 3)

filterRnaQcMetrics(thresholds, metrics, block = NULL)
```

Arguments

x	A matrix-like object where rows are genes and columns are cells. Values are expected to be counts.
subsets	Named list of vectors specifying gene subsets of interest, typically for control-like features like mitochondrial genes or spike-in transcripts. Each vector may be logical (whether to keep each row), integer (row indices) or character (row names). For character vectors, strings not present in rownames(x) are ignored.
num.threads	Integer scalar specifying the number of threads to use.
metrics	DataFrame of per-cell QC metrics. This should have the same structure as the return value of <code>computeRnaQcMetrics</code> .
block	Factor specifying the block of origin (e.g., batch, sample) for each cell in <code>metrics</code> . Alternatively NULL if all cells are from the same block. For <code>filterRnaQcMetrics</code> , a blocking factor should be provided if <code>block</code> was used to construct thresholds.
num.mads	Number of median from the median, to define the threshold for outliers in each metric.
thresholds	List with the same structure as produced by <code>suggestRnaQcThresholds</code> .

Value

For `computeRnaQcMetrics`, a [DataFrame](#) is returned with one row per cell in `x`. This contains the following columns:

- `sum`, a numeric vector containing the total RNA count for each cell. This represents the efficiency of library preparation and sequencing. Low totals indicate that the library was not successfully captured.
- `detected`, an integer vector containing the number of detected genes per cell. This also quantifies library preparation efficiency but with greater focus on capturing transcriptional complexity.
- `subsets`, a nested `DataFrame` where each column corresponds to a feature subset and is a numeric vector containing the proportion of counts in that subset. The exact interpretation of which depends on the nature of the subset. For example, if one subset contains all genes on the mitochondrial chromosome, higher proportions are representative of cell damage; the assumption is that cytoplasmic transcripts leak through tears in the cell membrane while the mitochondria are still trapped inside. The proportion of spike-in transcripts can be interpreted in a similar manner, where the loss of endogenous transcripts results in higher spike-in proportions.

Each vector is of length equal to the number of cells.

For `suggestRnaQcThresholds`, a named list is returned.

- If `block=NULL`, the list contains:
 - `sum`, a numeric scalar containing the lower bound on the sum. This is defined as `num.mads` MADs below the median of the log-transformed metrics across all cells.
 - `detected`, a numeric scalar containing the lower bound on the number of detected genes. This is defined as `num.mads` MADs below the median of the log-transformed metrics across all cells.
 - `subsets`, a numeric vector containing the upper bound on the sum of counts in each feature subset. This is defined as `num.mads` MADs above the median across all cells.

- Otherwise, if `block` is supplied, the list contains:
 - `sum`, a numeric vector containing the lower bound on the sum for each blocking level. Here, the threshold is computed independently for each block, using the same method as the unblocked case.
 - `detected`, a numeric vector containing the lower bound on the number of detected genes for each blocking level. Here, the threshold is computed independently for each block, using the same method as the unblocked case.
 - `subsets`, a list of numeric vectors containing the upper bound on the sum of counts in each feature subset for each blocking level. Here, the threshold is computed independently for each block, using the same method as the unblocked case.
 - `block.ids`, a vector containing the identities of the unique blocks.

Each vector is of length equal to the number of levels in `block` and is named accordingly.

For `filterRnaQcMetrics`, a logical vector of length `ncol(x)` is returned indicating which cells are of high quality. High-quality cells are defined as those with sums and detected genes above their respective thresholds and subset proportions below the `subsets` threshold.

Author(s)

Aaron Lun

See Also

The `compute_rna_qc_metrics`, `compute_rna_qc_filters` and `compute_rna_qc_filters_blocked` functions in https://libscrn.github.io/scrn_qc/.

[quickRnaQc.se](#), to run all of the RNA-related QC functions on a [SummarizedExperiment](#).

Examples

```
# Mocking a matrix:
library(Matrix)
x <- round(abs(rsparsematrix(1000, 100, 0.1) * 100))

# Mocking up a control set.
sub <- list(mito=rbinom(nrow(x), 1, 0.1) > 0)

qc <- computeRnaQcMetrics(x, sub)
qc

filt <- suggestRnaQcThresholds(qc)
str(filt)

keep <- filterRnaQcMetrics(filt, qc)
summary(keep)
```

runAllNeighborSteps *Run all neighbor-related steps*

Description

Run all steps that require a nearest-neighbor search. This includes [runUmap](#), [runTsne](#) and [buildSnnGraph](#) with [clusterGraph](#). The idea is to build the index once, perform the neighbor search, and run each task in parallel to save time.

Usage

```
runAllNeighborSteps(
  x,
  runUmap.args = list(),
  runTsne.args = list(),
  buildSnnGraph.args = list(),
  clusterGraph.args = list(),
  BNPARAM = AnnoyParam(),
  return.graph = FALSE,
  collapse.search = TRUE,
  num.threads = 3
)
```

Arguments

x	Numeric matrix where rows are dimensions and columns are cells, typically containing a low-dimensional representation from, e.g., runPca . Alternatively, an index constructed by buildIndex .
runUmap.args	Named list of further arguments to pass to runUmap . This can be set to NULL to omit the UMAP.
runTsne.args	Named list of further arguments to pass to runTsne . This can be set to NULL to omit the t-SNE.
buildSnnGraph.args	Named list of further arguments to pass to buildSnnGraph . Ignored if <code>clusterGraph.args=NULL</code> .
clusterGraph.args	Named list of further arguments to pass to clusterGraph . This can be set to NULL to omit the clustering.
BNPARAM	A BiocNeighborParam instance specifying the nearest-neighbor search algorithm to use.
return.graph	Logical scalar indicating whether to return the output of buildSnnGraph . By default, only the output of clusterGraph is returned.
collapse.search	Logical scalar indicating whether to collapse the nearest-neighbor search for each step into a single search. Steps that need fewer neighbors will take a subset of the neighbors from the collapsed search. Setting this to TRUE is faster but may not give the same results as separate searches for some nearest-neighbor algorithms (e.g., approximate methods).
num.threads	Integer scalar specifying the number of threads to use. At least one thread should be available for each step.

Value

A named list containing the results of each step. See each individual function for the format of the results.

Author(s)

Aaron Lun

See Also

[runAllNeighborSteps.se](#), to run each neighbor-related step on a [SingleCellExperiment](#).

Examples

```
x <- t(as.matrix(iris[,1:4]))
# (Turning down the number of threads so that R CMD check is happy.)
res <- runAllNeighborSteps(x, num.threads=2)
str(res)
```

runAllNeighborSteps.se

Run all nearest neighbor steps on a SummarizedExperiment

Description

Concurrently run all steps involving a nearest-neighbor search (t-SNE, UMAP and graph-based clustering) using the same nearest-neighbor index, by calling [runAllNeighborSteps](#) on a reduced dimension entry of a [SingleCellExperiment](#).

Usage

```
runAllNeighborSteps.se(
  x,
  umap.output.name = "UMAP",
  more.umap.args = list(),
  tsne.output.name = "TSNE",
  more.tsne.args = list(),
  build.graph.name = NULL,
  more.build.graph.args = list(),
  cluster.output.name = "clusters",
  cluster.meta.name = NULL,
  more.cluster.graph.args = list(),
  BNPARAM = AnnoyParam(),
  num.threads = 3,
  more.neighbor.args = list(),
  reddim.type = "PCA"
)
```

Arguments

<code>x</code>	A SingleCellExperiment object or one of its subclasses. Rows correspond to genomic features and columns correspond to cells.
<code>umap.output.name</code>	String containing the name of the reducedDim entry to store the UMAP coordinates. If NULL, the UMAP is not computed.
<code>more.umap.args</code>	Named list of additional arguments to pass to runAllNeighborSteps as <code>runUmap.args</code> .
<code>tsne.output.name</code>	String containing the name of the reducedDim entry to store the t-SNE coordinates. If NULL, the t-SNE is not computed.
<code>more.tsne.args</code>	Named list of additional arguments to pass to runAllNeighborSteps as <code>runTsne.args</code> .
<code>build.graph.name</code>	String containing the name of the metadata entry in which to store the nearest neighbor graph. If NULL, the graph is not stored.
<code>more.build.graph.args</code>	Named list of additional arguments to pass to runAllNeighborSteps as <code>buildSnnGraph.args</code> .
<code>cluster.output.name</code>	String containing the name of the colData column in which to store the cluster assignments. If NULL, graph-based clustering is not performed.
<code>cluster.meta.name</code>	String containing the name of the metadata entry in which to store additional clustering outputs. If NULL, these additional outputs are not stored.
<code>more.cluster.graph.args</code>	Named list of additional arguments to pass to runAllNeighborSteps as <code>clusterGraph.args</code> .
<code>BNPARAM, num.threads</code>	Arguments to pass to runAllNeighborSteps .
<code>more.neighbor.args</code>	Named list of additional arguments to pass to runAllNeighborSteps .
<code>reddim.type</code>	String or integer specifying the reducedDim entry on which to perform a nearest neighbor search. Alternatively, a named integer or character vector of length 1, where the name specifies an alternative experiment of <code>x</code> and the value is the name/index of a reducedDim entry in that alternative experiment.

Value

`x` is returned with additional coordinates stored in its [reducedDims](#) and clustering output in its [colData](#). Additional information may also be stored in its [metadata](#).

Author(s)

Aaron Lun

Examples

```
library(SingleCellExperiment)
sce <- getTestRnaData.se("pca")
sce <- runAllNeighborSteps.se(
  sce,
  more.tsne.args=list(max.iterations=50),
  more.umap.args=list(num.epochs=50),
```

```

    num.threads=2 # to keep R CMD check happy
  )
  reducedDimNames(sce)
  table(sce$clusters)

```

runPca

Principal components analysis

Description

Run a PCA on the gene-by-cell log-expression matrix and extract the top principal components (PCs). This yields a low-dimensional representation that reduces noise and compute time in downstream analyses. For efficiency, the PCA itself is approximated using IRLBA.

Usage

```

runPca(
  x,
  number = 25,
  scale = FALSE,
  block = NULL,
  block.weight.policy = c("variable", "equal", "none"),
  variable.block.weight = c(0, 1000),
  components.from.residuals = FALSE,
  subset = NULL,
  extra.work = 7,
  iterations = 1000,
  seed = 5489,
  realized = TRUE,
  warn = TRUE,
  num.threads = 1
)

```

Arguments

x	A matrix-like object where rows correspond to genes or genomic features and columns correspond to cells. Typically, the matrix is expected to contain log-expression values (see normalizeCounts) for “interesting” genes (see chooseHighlyVariableGenes).
number	Integer scalar specifying the number of top PCs to retain. More PCs will capture more biological signal at the cost of increasing noise and compute time. If this is greater than the maximum number of PCs (i.e., the smaller dimension of x), only the maximum number of PCs will be reported in the results.
scale	Logical scalar indicating whether to scale all genes to have the same variance. This ensures that each gene contributes equally to the PCA, favoring consistent variation across many genes rather than large variation in a few genes. If block is specified, each gene’s variance is calculated as a weighted sum of the variances from each block. Genes with zero variance are ignored.
block	Factor specifying the block of origin (e.g., batch, sample) for each cell in x. The PCA will be performed on the residuals after regressing out the block effect, ensuring that differences between block do not dominate the variation in the dataset. Alternatively NULL if all cells are from the same block.

<code>block.weight.policy</code>	String specifying the policy to use for weighting the contribution of different blocks to the PCA. See the argument of the same name in computeBlockWeights for more detail. Only used if <code>block</code> is not NULL.
<code>variable.block.weight</code>	Numeric vector of length 2, specifying the parameters for variable block weighting. See the argument of the same name in computeBlockWeights for more detail. Only used if <code>block</code> is not NULL and <code>block.weight.policy = "variable"</code> .
<code>components.from.residuals</code>	Logical scalar indicating whether to compute the PC scores from the residuals in the presence of a blocking factor. By default, the residuals are only used to compute the rotation matrix, and the original expression values of the cells are projected onto this new space (see Details). Only used if <code>block</code> is not NULL.
<code>subset</code>	Integer, logical or character vector specifying the rows of <code>x</code> to use for the PCA. This yields the same results as <code>runPca</code> on <code>x[subset,]</code> , except that entries of the rotation matrix will also be computed for rows outside of the subset. If NULL, all rows of <code>x</code> are used.
<code>extra.work</code>	Integer scalar specifying the extra dimensions for the IRLBA workspace. Larger values improve accuracy at the cost of compute time.
<code>iterations</code>	Integer scalar specifying the maximum number of restart iterations for IRLBA. Larger values improve accuracy at the cost of compute time.
<code>seed</code>	Integer scalar specifying the seed for the initial random vector in IRLBA.
<code>realized</code>	Logical scalar indicating whether to realize <code>x</code> into an optimal memory layout for IRLBA. This speeds up computation at the cost of increased memory usage.
<code>warn</code>	Boolean specifying whether a warning should be emitted if IRLBA failed to converge.
<code>num.threads</code>	Number of threads to use.

Details

When `block` is specified, the nature of the reported PC scores depends on the choice of `components.from.residuals`:

- If TRUE, the PC scores are computed from the matrix of residuals. This yields a low-dimensional space where inter-block differences have been removed, assuming that all blocks have the same subpopulation composition and the inter-block differences are consistent for all cell subpopulations. Under these assumptions, we could use these components for downstream analysis without any concern for block-wise effects.
- If FALSE, the rotation vectors are first computed from the matrix of residuals. To obtain PC scores, each cell is then projected onto the associated subspace using its original expression values. This approach ensures that inter-block differences do not contribute to the PCA but does not attempt to explicitly remove them.

In complex datasets, the assumptions mentioned for TRUE not hold and more sophisticated batch correction methods like MNN correction are required. Functions like [correctMnn](#) will accept a low-dimensional embedding of cells that can be created as described above with FALSE.

Value

List containing:

- `components`, a matrix of PC scores. Rows are dimensions (i.e., PCs) and columns are cells.

- `rotation`, the rotation matrix. Rows are genes and columns are dimensions.
- `variance.explained`, the vector of variances explained by each PC.
- `total.variance`, the total variance in the dataset. This can be used to divide `variance.explained` to obtain the proportion of variance explained by each PC.
- `center`, a numeric vector containing the mean for each gene. If `block` is provided, this is instead a matrix containing the mean for each gene (column) in each block (row).
- `block.ids`, a vector containing the identities of the unique blocks in the same order as the rows of `center`. Only reported if `block` is provided.
- `scale`, a numeric vector containing the scaling for each gene. Only reported if `scale=TRUE`.
- `metrics`, a named list containing IRLBA metrics. This includes `converged`, whether the algorithm converged successfully; `iterations`, the number of IRLBA restart iterations; and `multiplications`, the number of matrix multiplications.

Author(s)

Aaron Lun

See Also

The `simple_pca` and `blocked_pca` functions for https://libscran.github.io/scran_pca/, [runPca.se](#), to run a PCA on a [SummarizedExperiment](#).

Examples

```
library(Matrix)
x <- abs(rsparsematrix(1000, 100, 0.1) * 10)
y <- normalizeCounts(x, size.factors=centerSizeFactors(colSums(x)))

# A simple PCA:
out <- runPca(y)
str(out)

# Blocking on uninteresting factors:
block <- sample(LETTERS[1:3], ncol(y), replace=TRUE)
bout <- runPca(y, block=block)
str(bout)
```

runPca.se

Principal components analysis of a Summarizedexperiment

Description

Compact and denoise the dataset by performing PCA on the (log-)normalized expression matrix, by calling [runPca](#) on an assay of a [SummarizedExperiment](#).

Usage

```
runPca.se(
  x,
  features,
  number = 25,
  block = NULL,
  num.threads = 1,
  more.pca.args = list(),
  assay.type = "logcounts",
  output.name = "PCA",
  meta.name = "PCA",
  delayed.transpose = FALSE
)
```

Arguments

x	A SummarizedExperiment object or one of its subclasses. Rows correspond to genomic features and columns correspond to cells.
features	Integer, logical or character vector containing the features of interest to use in the PCA. For RNA data, this is typically the hvg vector added by chooseRnaHvgs.se . If NULL, all available features are used.
number	Number of PCs to retain, passed to runPca .
block	Block assignment for each cell, passed to runPca .
num.threads	Number of threads for the PCA, passed to runPca .
more.pca.args	Named list of additional arguments to pass to runPca .
assay.type	Integer or string specifying the assay of x to be used for PCA. This is typically the log-normalized expression matrix created by normalizeRnaCounts.se .
output.name	String containing the name of the reducedDim entry in which to store the PC scores.
meta.name	String containing the name of the <code>link[S4Vectors]{metadata}</code> entry in which to store other PCA statistics.
delayed.transpose	Logical scalar indicating whether to delay the transposition when storing coordinates in the reducedDims .

Value

x is returned with the principal component scores in the `reducedDim`. (This is converted to a [SingleCellExperiment](#) if it wasn't one already.) Additional outputs (e.g., rotation matrix, variance explained) are stored in the `metadata`.

Author(s)

Aaron Lun

Examples

```
library(SingleCellExperiment)
sce <- getTestRnaData.se("hvg")
sce <- runPca.se(sce, rowData(sce)$hvg)
dim(reducedDim(sce, "PCA"))
```

```
plot(metadata(sce)$PCA$variance.explained / metadata(sce)$PCA$total.variance)
```

```
runTsne          t-stochastic neighbor embedding
```

Description

Compute t-SNE coordinates to visualize similarities between cells.

Usage

```
runTsne(
  x,
  perplexity = 30,
  num.neighbors = tsnePerplexityToNeighbors(perplexity),
  theta = 1,
  early.exaggeration.iterations = 250,
  exaggeration.factor = 12,
  momentum.switch.iterations = 250,
  start.momentum = 0.5,
  final.momentum = 0.8,
  eta = 200,
  max.depth = 7,
  leaf.approximation = FALSE,
  max.iterations = 500,
  seed = 42,
  num.threads = 1,
  BNPARAM = AnnoyParam()
)

tsnePerplexityToNeighbors(perplexity)
```

Arguments

- | | |
|------------|--|
| x | <p>Numeric matrix where rows are dimensions and columns are cells, typically containing a low-dimensional representation from, e.g., runPca.</p> <p>Alternatively, a named list of nearest-neighbor search results like that returned by findKNN. This should contain <code>index</code>, an integer matrix where rows are neighbors and columns are cells; and <code>distance</code>, a numeric matrix of the same dimensions containing the distances to each neighbor. Each column contains 1-based indices for the nearest neighbors of the corresponding cell, ordered by increasing distance. The number of neighbors should be the same as <code>num.neighbors</code>, otherwise a warning is raised.</p> <p>Alternatively, an index constructed by buildIndex.</p> |
| perplexity | <p>Numeric scalar specifying the perplexity to use in the t-SNE algorithm. Higher perplexities will focus on global structure, at the cost of increased runtime and decreased local resolution.</p> |

<code>num.neighbors</code>	Integer scalar specifying the number of neighbors, typically derived from <code>perplexity</code> . If <code>x</code> contains pre-computed neighbor search results with a different number of neighbors than <code>num.neighbors</code> , an error is thrown; this can be suppressed by setting <code>num.neighbors = NULL</code> .
<code>theta</code>	Numeric scalar specifying the approximation level for the Barnes-Hut calculation of repulsive forces. Lower values increase accuracy at the cost of increased compute time. All values should be non-negative.
<code>early.exaggeration.iterations</code>	Integer scalar specifying the number of iterations of the early exaggeration phase, where clusters are artificially compacted to leave more empty space so that cells can easily relocate to find a good global organization. Larger values improve convergence within this phase at the cost of reducing the remaining iterations in <code>max.iterations</code> .
<code>exaggeration.factor</code>	Numeric scalar containing the exaggeration factor for the early exaggeration phase (see <code>early.exaggeration.iterations</code>). Larger values increase the attraction between nearest neighbors to favor local structure.
<code>momentum.switch.iterations</code>	Integer scalar specifying the number of iterations to perform before switching from the starting momentum to the final momentum. Higher momentums can improve convergence by increasing the step size and smoothing over local oscillations, at the risk of potentially skipping over relevant minima.
<code>start.momentum</code>	Numeric scalar containing the starting momentum, to be used in the iterations before the momentum switch at <code>momentum.switch.iterations</code> . This is usually lower than <code>final.momentum</code> to avoid skipping over suitable local minima.
<code>final.momentum</code>	Numeric scalar containing the final momentum, to be used in the iterations after the momentum switch at <code>momentum.switch.iterations</code> . This is usually higher than <code>start.momentum</code> to accelerate convergence to the local minima once the observations are moderately well-organized.
<code>eta</code>	Numeric scalar containing the learning rate, used to scale the updates for each cell. Larger values can speed up convergence at the cost of skipping over local minima.
<code>max.depth</code>	Integer scalar specifying the maximum depth of the Barnes-Hut quadtree. If neighboring cells cannot be separated before the maximum depth is reached, they will be assigned to the same leaf node of the quadtree. Smaller values (7-10) improve speed by bounding the recursion depth at the cost of accuracy.
<code>leaf.approximation</code>	Logical scalar indicating whether to use the “leaf approximation”. If TRUE, repulsive forces are computed between leaf nodes and re-used across all cells assigned to that leaf node. This sacrifices some accuracy for greater speed, assuming that <code>max.depth</code> is small enough for multiple cells to be assigned to the same leaf.
<code>max.iterations</code>	Integer scalar specifying the maximum number of iterations to perform. Larger values improve convergence at the cost of compute time.
<code>seed</code>	Integer scalar specifying the seed to use for generating the initial coordinates.
<code>num.threads</code>	Integer scalar specifying the number of threads to use.
<code>BNPARAM</code>	A <code>BiocNeighborParam</code> object specifying the algorithm to use. Only used if <code>x</code> is not a prebuilt index or a list of existing nearest-neighbor search results.

Value

For runTsne, a numeric matrix where rows are cells and columns are the two dimensions of the embedding.

For tsnePerplexityToNeighbors, an integer scalar specifying the number of neighbors to use for a given perplexity.

Author(s)

Aaron Lun

References

van der Maaten LJP and Hinton GE (2008). Visualizing high-dimensional data using t-SNE. *Journal of Machine Learning Research* 9, 2579-2605.

van der Maaten LJP (2014). Accelerating t-SNE using tree-based algorithms. *Journal of Machine Learning Research* 15, 3221-3245.

See Also

<https://libscan.github.io/qdtsne/>, for an explanation of the approximations.

[runTsne.se](#), to run t-SNE on a [SingleCellExperiment](#).

Examples

```
x <- t(as.matrix(iris[,1:4]))
embedding <- runTsne(x)
plot(embedding[,1], embedding[,2], col=iris[,5])
```

runTsne.se

t-SNE on a SummarizedExperiment

Description

Generate a t-SNE visualization from an existing embedding, by calling [runUmap](#) on a reduced dimension entry in [SingleCellExperiment](#).

Usage

```
runTsne.se(
  x,
  perplexity = 30,
  num.threads = 1,
  more.tsne.args = list(),
  reddim.type = "PCA",
  output.name = "TSNE"
)
```

Arguments

x	A SingleCellExperiment object or one of its subclasses. Rows correspond to genomic features and columns correspond to cells.
perplexity	Perplexity to use in the t-SNE algorithm, passed to runTsne .
num.threads	Number of threads for the neighbor search and optimization, passed to runTsne .
more.tsne.args	Named list of further arguments to pass to runTsne .
reddim.type	Integer or string specifying the existing embedding in the reducedDim of x. Alternatively, a named integer or character vector of length 1, where the name specifies an alternative experiment of x and the value is the name/index of a reducedDim entry in that alternative experiment.
output.name	String containing the name of the output reducedDim .

Value

x is returned with the t-SNE coordinates stored in the [reducedDim](#).

Author(s)

Aaron Lun

Examples

```
library(SingleCellExperiment)
sce <- getTestRnaData.se("pca")
# Using fewer iterations for a faster-running example.
sce <- runTsne.se(sce, more.tsne.args=list(max.iterations=50))
head(reducedDim(sce, "TSNE"))
```

runUmap

Uniform manifold approximation and projection

Description

Compute UMAP coordinates to visualize similarities between cells.

Usage

```
runUmap(
  x,
  num.dim = 2,
  local.connectivity = 1,
  bandwidth = 1,
  mix.ratio = 1,
  spread = 1,
  min.dist = 0.1,
  a = NULL,
  b = NULL,
  repulsion.strength = 1,
  initialize.method = c("spectral", "random", "none"),
```

```

initial.coordinates = NULL,
initialize.random.on.spectral.fail = TRUE,
initialize.spectral.scale = 10,
initialize.spectral.jitter = FALSE,
initialize.spectral.jitter.sd = 1e-04,
initialize.random.scale = 10,
initialize.seed = 9876543210,
num.epochs = NULL,
learning.rate = 1,
negative.sample.rate = 5,
num.neighbors = 15,
optimize.seed = 1234567890,
num.threads = 1,
parallel.optimization = FALSE,
BNPARAM = AnnoyParam()
)

```

Arguments

<code>x</code>	<p>Numeric matrix where rows are dimensions and columns are cells, typically containing a low-dimensional representation from, e.g., runPca.</p> <p>Alternatively, a named list of nearest-neighbor search results like that returned by findKNN. This should contain <code>index</code>, an integer matrix where rows are neighbors and columns are cells; and <code>distance</code>, a numeric matrix of the same dimensions containing the distances to each neighbor. Each column contains 1-based indices for the nearest neighbors of the corresponding cell, ordered by increasing distance. The number of neighbors should be the same as <code>num.neighbors</code>, otherwise a warning is raised.</p> <p>Alternatively, an index constructed by buildIndex.</p>
<code>num.dim</code>	Integer scalar specifying the number of dimensions of the output embedding.
<code>local.connectivity</code>	Numeric scalar specifying the number of nearest neighbors that are assumed to be always connected, with maximum membership confidence. Larger values increase the connectivity of the embedding and reduce the focus on local structure. This may be a fractional number of neighbors, in which case interpolation is performed when computing the membership confidence.
<code>bandwidth</code>	Numeric scalar specifying the effective bandwidth of the kernel when converting the distance to a neighbor into a fuzzy set membership confidence. Larger values reduce the decay in confidence with respect to distance, increasing connectivity and favoring global structure.
<code>mix.ratio</code>	Numeric scalar between 0 and 1 specifying the mixing ratio when combining fuzzy sets. A mixing ratio of 1 will take the union of confidences, a ratio of 0 will take the intersection, and intermediate values will interpolate between them. Larger values favor connectivity and more global structure.
<code>spread</code>	Numeric scalar specifying the scale of the coordinates of the final low-dimensional embedding. Ignored if <code>a</code> and <code>b</code> are provided.
<code>min.dist</code>	Numeric scalar specifying the minimum distance between observations in the final low-dimensional embedding. Smaller values will increase local clustering while larger values favor a more even distribution of observations throughout the low-dimensional space. This is interpreted relative to <code>spread</code> . Ignored if <code>a</code> and <code>b</code> are provided.

- a** Numeric scalar specifying the a parameter for the fuzzy set membership strength calculations. Larger values yield a sharper decay in membership strength with increasing distance between observations. If this or **b** are NULL, a suitable value for this parameter is automatically determined from `spread` and `min.dist`.
- b** Numeric scalar specifying the b parameter for the fuzzy set membership strength calculations. Larger values yield an earlier decay in membership strength with increasing distance between observations. If this or **a** are NULL, a suitable value for this parameter is automatically determined from `spread` and `min.dist`.
- `repulsion.strength` Numeric scalar specifying the modifier for the repulsive force. Larger values increase repulsion and favor local structure.
- `initialize.method` String specifying how to initialize the embedding. This should be one of:
- `spectral`: spectral decomposition of the normalized graph Laplacian. Specifically, the initial coordinates are defined from the eigenvectors corresponding to the smallest non-zero eigenvalues. This fails in the presence of multiple graph components or if the approximate SVD fails to converge.
 - `random`: fills the embedding with random draws from a normal distribution.
 - `none`: uses initial values from `initial.coordinates`.
- `initial.coordinates` Numeric matrix of initial coordinates, with number of rows equal to the number of observations and number of columns equal to `num.dim`. Only relevant if `initialize.method = "none"`; or `initialize.method = "spectral"` and spectral initialization fails and `initialize.random.on.spectral.fail = FALSE`.
- `initialize.random.on.spectral.fail` Logical scalar indicating whether to fall back to random sampling (i.e., same as `random`) if spectral initialization fails due to the presence of multiple components in the graph. If `FALSE`, the values in `initial.coordinates` will be used instead, i.e., same as `none`. Only relevant if `initialize.method = "spectral"` and spectral initialization fails.
- `initialize.spectral.scale` Numeric scalar specifying the maximum absolute magnitude of the coordinates after spectral initialization. All initial coordinates are scaled such that the maximum of the absolute values is equal to `initialize.spectral.scale`. This ensures that outlier observations will not have large absolute distances that may interfere with optimization. Only relevant if `initialize.method = "spectral"` and spectral initialization does not fail.
- `initialize.spectral.jitter` Logical scalar indicating whether to jitter coordinates after spectral initialization to separate duplicate observations (e.g., to avoid overplotting). This is done using normally-distributed noise of mean zero and standard deviation of `initialize.spectral.jitter.sd`. Only relevant if `initialize.method = "spectral"` and spectral initialization does not fail.
- `initialize.spectral.jitter.sd` Numeric scalar specifying the standard deviation of the jitter to apply after spectral initialization. Only relevant if `initialize.method = "spectral"` and spectral initialization does not fail and `initialize.spectral.jitter = TRUE`.
- `initialize.random.scale` Numeric scalar specifying the scale of the randomly generated initial coordinates. Coordinates are sampled from a uniform distribution from $[-x, x]$

where x is `initialize.random.scale`. Only relevant if `initialize.method = "random"`, or `initialize.method = "spectral"` and spectral initialization fails and `initialize.random.on.spectral.fail = TRUE`.

<code>initialize.seed</code>	Numeric scalar specifying the seed for the random number generation during initialization. Only relevant if <code>initialize.method = "random"</code> , or <code>initialize.method = "spectral"</code> and <code>initialize.spectral.jitter = TRUE</code> ; or <code>initialize.method = "spectral"</code> and spectral initialization fails and <code>initialize.random.on.spectral.fail = TRUE</code> .
<code>num.epochs</code>	Integer scalar specifying the number of epochs for the gradient descent, i.e., optimization iterations. Larger values improve accuracy at the cost of increased compute time. If NULL, a value is automatically chosen based on the size of the dataset: <ul style="list-style-type: none"> • For datasets with no more than 10000 observations, the default number of epochs is set to 500. • For larger datasets, the number of epochs is inversely proportional to the number of cells, starting from 500 and decreasing asymptotically to a lower limit of 200. This choice aims to reduce computational work for very large datasets.
<code>learning.rate</code>	Numeric scalar specifying the initial learning rate used in the gradient descent. Larger values can accelerate convergence but at the risk of skipping over suitable local optima.
<code>negative.sample.rate</code>	Numeric scalar specifying the rate of sampling negative observations to compute repulsive forces. Greater values will improve accuracy but increase compute time.
<code>num.neighbors</code>	Integer scalar specifying the number of neighbors to use to define the fuzzy sets. Larger values improve connectivity and favor preservation of global structure, at the cost of increased compute time. Ignored if <code>x</code> contains pre-computed neighbor search results.
<code>optimize.seed</code>	Numeric scalar specifying the seed to use for the optimization epochs.
<code>num.threads</code>	Integer scalar specifying the number of threads to use.
<code>parallel.optimization</code>	Logical scalar specifying whether to parallelize the optimization step.
BNPARAM	A <code>BiocNeighborParam</code> object specifying the algorithm to use. Only used if <code>x</code> is not a prebuilt index or a list of existing nearest-neighbor search results.

Value

A numeric matrix where rows are cells and columns are the two dimensions of the embedding.

Author(s)

Aaron Lun

References

McInnes L, Healy J, Melville J (2020). UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. *arXiv*, <https://arxiv.org/abs/1802.03426>

See Also

<https://libscrn.github.io/umapp/>, for details on the underlying implementation.
[runUmap.se](#), to run UMAP on a [SingleCellExperiment](#).

Examples

```
x <- t(as.matrix(iris[,1:4]))
embedding <- runUmap(x)
plot(embedding[,1], embedding[,2], col=iris[,5])
```

runUmap.se

*UMAP on a SummarizedExperiment***Description**

Generate a UMAP visualization from an existing embedding, by calling [runUmap](#) on a reduced dimension entry in [SingleCellExperiment](#).

Usage

```
runUmap.se(
  x,
  num.dim = 2,
  min.dist = 0.1,
  num.neighbors = 15,
  num.threads = 1,
  more.umap.args = list(),
  reddim.type = "PCA",
  output.name = "UMAP"
)
```

Arguments

<code>x</code>	A SingleCellExperiment object or one of its subclasses. Rows correspond to genomic features and columns correspond to cells.
<code>num.dim</code>	Number of dimensions in the output embedding, passed to runUmap .
<code>min.dist</code>	Minimum distance between observations, passed to runUmap .
<code>num.neighbors</code>	Number of neighbors for constructing the fuzzy sets, passed to runUmap .
<code>num.threads</code>	Number of threads for the UMAP, passed to runUmap .
<code>more.umap.args</code>	Named list of further arguments to pass to runUmap .
<code>reddim.type</code>	Integer or string specifying the existing embedding in the <code>reducedDim</code> of <code>x</code> . Alternatively, a named integer or character vector of length 1, where the name specifies an alternative experiment of <code>x</code> and the value is the name/index of a <code>reducedDim</code> entry in that alternative experiment.
<code>output.name</code>	String containing the name of the output <code>reducedDim</code> .

Value

`x` is returned with the UMAP coordinates stored in the `reducedDim`.

Author(s)

Aaron Lun

Examples

```
library(SingleCellExperiment)
sce <- getTestRnaData.se("pca")
# Using fewer epochs for a faster-running example.
sce <- runUmap.se(sce, more.umap.args=list(num.epochs=50))
head(reducedDim(sce, "UMAP"))
```

sanitizeSizeFactors *Sanitize size factors*

Description

Replace invalid size factors, i.e., zero, negative, infinite or NaN values. Such size factors can occasionally arise if, e.g., insufficient quality control was performed upstream. Removing them ensures that the normalized values from `normalizeCounts` remain finite for sensible downstream processing.

Usage

```
sanitizeSizeFactors(
  size.factors,
  replace.zero = TRUE,
  replace.negative = TRUE,
  replace.infinite = TRUE,
  replace.nan = TRUE
)
```

Arguments

<code>size.factors</code>	Numeric vector of size factors across cells.
<code>replace.zero</code>	Logical scalar indicating whether to replace size factors of zero with the lowest positive factor in <code>size.factors</code> . This ensures that the normalized values will be large to reflect the extremity of the scaling, but still finite for sensible downstream processing. If FALSE, zeros are retained.
<code>replace.negative</code>	Logical scalar indicating whether to replace negative size factors with the lowest positive factor in <code>size.factors</code> . This ensures that the normalized values will be large to reflect the extremity of the scaling, but still finite for sensible downstream processing. If FALSE, negative values are retained.
<code>replace.infinite</code>	Logical scalar indicating whether to replace infinite size factors with the largest positive factor in <code>size.factors</code> . This ensures that any normalized values will be, at least, finite; the choice of a relatively large replacement value reflects the extremity of the scaling. If FALSE, infinite values are retained.
<code>replace.nan</code>	Logical scalar indicating whether to replace NaN size factors with unity, e.g., scaling normalization is a no-op. If FALSE, NaN values are retained.

Value

Numeric vector of length equal to `size.factors`, containing the sanitized size factors.

Author(s)

Aaron Lun

See Also

The `sanitize_size_factors` function in https://libscran.github.io/scran_norm/.

Examples

```
sf <- 2^rnorm(100)
sf[1] <- 0
sf[2] <- -1
sf[3] <- Inf
sf[4] <- NaN
sanitizeSizeFactors(sf)
```

scaleByNeighbors

Scale and combine multiple embeddings

Description

Scale multiple embeddings (usually derived from different modalities for the same cells) so that their within-population variances are comparable, and then combine them into a single embedding matrix for further analyses like clustering, t-SNE, etc. The aim is to equalize uninteresting variance across modalities so that high technical variance in one modality does not drown out interesting biology in another modality.

Usage

```
scaleByNeighbors(
  x,
  num.neighbors = 20,
  block = NULL,
  block.weight.policy = c("variable", "equal", "none"),
  variable.block.weight = c(0, 1000),
  num.threads = 1,
  weights = NULL,
  BNPARAM = AnnoyParam()
)
```

Arguments

`x` List of numeric matrices of principal components or other embeddings, one for each modality. For each entry, rows are dimensions and columns are cells. All entries should have the same number of columns but may have different numbers of rows.

num.neighbors	Integer scalar specifying the number of neighbors to use to define the scaling factor.
block	Factor specifying the block of origin (e.g., batch, sample) for each cell in <i>x</i> . If provided, the scaling factor is computed as a weighted average across blocks to ensure that block effects do not inflate the within-population variance. Alternatively NULL, if all cells are from the same block.
block.weight.policy	String specifying the policy to use for weighting different blocks when computing the average scaling factor. See the argument of the same name in computeBlockWeights for more detail. Only used if <code>block</code> is not NULL.
variable.block.weight	Numeric vector of length 2, specifying the parameters for variable block weighting. See the argument of the same name in computeBlockWeights for more detail. Only used if <code>block</code> is not NULL and <code>block.weight.policy = "variable"</code> .
num.threads	Integer scalar specifying the number of threads to use.
weights	Numeric vector of length equal to that of <i>x</i> , specifying the weights to apply to each modality. Each value represents a multiplier of the within-population variance of its modality, i.e., larger values increase the contribution of that modality in the combined output matrix. NULL is equivalent to an all-1 vector, i.e., all modalities are scaled to have the same within-population variance.
BNPARAM	A BiocNeighborParam object specifying how to perform the neighbor search.

Value

List containing `scaling`, a vector of scaling factors to be applied to each embedding; and `combined`, a numeric matrix created by scaling each entry of *x* by `scaling` and then `rbinding` them together.

Author(s)

Aaron Lun

See Also

<https://libscrn.github.io/mumosa/>, for the basis and caveats of this approach.
[scaleByNeighbors.se](#), to combine embeddings in a [SingleCellExperiment](#).

Examples

```
pcs <- list(
  gene = matrix(rnorm(10000), ncol=200),
  protein = matrix(rnorm(1000, sd=3), ncol=200),
  guide = matrix(rnorm(2000, sd=5), ncol=200)
)

out <- scaleByNeighbors(pcs)
out$scaling
dim(out$combined)
```

scaleByNeighbors.se *Scale and combine multiple embeddings in a SingleCellExperiment*

Description

Scale embeddings for different modalities to equalize their intra-population variance, and combine them into a single embedding for downstream analysis. This calls `scaleByNeighbors` on the reduced dimensions of the main/alternative experiments in a `SingleCellExperiment`.

Usage

```
scaleByNeighbors.se(
  x,
  altexp.reddims,
  main.reddims = "PCA",
  num.neighbors = 20,
  block = NULL,
  BNPARAM = AnnoyParam(),
  num.threads = 1,
  more.scale.args = list(),
  output.name = "combined",
  meta.name = "combined",
  delayed.transpose = FALSE
)
```

Arguments

<code>x</code>	A <code>SingleCellExperiment</code> object or one of its subclasses. Rows correspond to genomic features and columns correspond to cells.
<code>altexp.reddims</code>	Named list of character or integer vectors. Each entry is named after an alternative experiment. Each vector contains the names/indices of the <code>reducedDim</code> embeddings from that experiment to be combined.
<code>main.reddims</code>	Character or integer vector specifying the names/indices of the <code>reducedDim</code> entries from <code>x</code> to be combined.
<code>num.neighbors</code>	Number of neighbors used to define the scaling factor, passed to <code>scaleByNeighbors</code> .
<code>block</code>	Block assignment for each cell, passed to <code>scaleByNeighbors</code> .
<code>BNPARAM</code>	Algorithm for the nearest neighbor search, passed to <code>scaleByNeighbors</code> .
<code>num.threads</code>	Number of threads for the neighbor search, passed to <code>scaleByNeighbors</code> .
<code>more.scale.args</code>	Named list of additional arguments to pass to <code>scaleByNeighbors</code> .
<code>output.name</code>	String containing the name of the <code>reducedDim</code> entry in which to store the combined embeddings.
<code>meta.name</code>	String containing the name of the <code>metadata</code> entry in which to store additional metrics. If <code>NULL</code> , additional metrics are not stored.
<code>delayed.transpose</code>	Logical scalar indicating whether to delay the transposition when storing coordinates in the <code>reducedDims</code> .

Value

x is returned with the combined embeddings stored in its reducedDims. The scaling factors for all embeddings are stored in the metadata.

Author(s)

Aaron Lun

Examples

```
library(SingleCellExperiment)
sce <- getTestAdtData.se("pca")
sce <- scaleByNeighbors.se(sce, altexp.reddims=list(ADT="PCA"))
reducedDimNames(sce)
metadata(sce)$combined
```

scoreGeneSet	<i>Score gene set activity for each cell</i>
--------------	--

Description

Compute per-cell scores for a gene set, defined as the column sums of a rank-1 approximation to the submatrix for the gene set. This uses the same approach as the **GSDecon** package by Jason Hackney, adapted to use an approximate PCA (via IRLBA) and to support blocking.

Usage

```
scoreGeneSet(
  x,
  set,
  rank = 1,
  scale = FALSE,
  block = NULL,
  block.weight.policy = c("variable", "equal", "none"),
  variable.block.weight = c(0, 1000),
  extra.work = 7,
  iterations = 1000,
  seed = 5489,
  realized = TRUE,
  num.threads = 1
)
```

Arguments

x A matrix-like object where rows correspond to genes or genomic features and columns correspond to cells. Typically, the matrix is expected to contain log-expression values.

set	Vector specifying the rows of <code>x</code> that belong to the gene set. This may be an integer vector of row indices, a logical vector of length equal to the number of rows, or a character vector of row names. For integer and character vectors, duplicate entries are ignored. For a character vector, any string not present in <code>rownames(x)</code> is ignored.
rank	Integer scalar specifying the rank of the approximation. The default value of 1 assumes that each gene set only describes a single coordinated biological function.
scale	Logical scalar indicating whether to scale all genes to have the same variance. This ensures that each gene contributes equally to the PCA, favoring consistent variation across many genes rather than large variation in a few genes. If <code>block</code> is specified, each gene's variance is calculated as a weighted sum of the variances from each block. Genes with zero variance are ignored.
block	Factor specifying the block of origin (e.g., batch, sample) for each cell in <code>x</code> . The PCA will be performed on the residuals after regressing out the block effect, ensuring that differences between block do not dominate the variation in the dataset. Alternatively NULL if all cells are from the same block.
block.weight.policy	String specifying the policy to use for weighting the contribution of different blocks to the PCA. See the argument of the same name in computeBlockWeights for more detail. Only used if <code>block</code> is not NULL.
variable.block.weight	Numeric vector of length 2, specifying the parameters for variable block weighting. See the argument of the same name in computeBlockWeights for more detail. Only used if <code>block</code> is not NULL and <code>block.weight.policy = "variable"</code> .
extra.work	Integer scalar specifying the extra dimensions for the IRLBA workspace. Larger values improve accuracy at the cost of compute time.
iterations	Integer scalar specifying the maximum number of restart iterations for IRLBA. Larger values improve accuracy at the cost of compute time.
seed	Integer scalar specifying the seed for the initial random vector in IRLBA.
realized	Logical scalar indicating whether to realize <code>x</code> into an optimal memory layout for IRLBA. This speeds up computation at the cost of increased memory usage.
num.threads	Number of threads to use.

Value

List containing:

- `scores`, a numeric vector of per-cell scores for each column in `x`.
- `weights`, a [DataFrame](#) containing `row`, an integer vector of ordered and unique row indices corresponding to the genes in `set`; and `weight`, a numeric vector of per-gene weights for each gene in `row`.

Author(s)

Aaron Lun

See Also

The `compute` and `compute_blocked` functions in <https://libscrans.github.io/guide/>, to compute gene set scores from a [SummarizedExperiment](#).

Examples

```
library(Matrix)
x <- round(abs(rsparsematrix(1000, 100, 0.1) * 100))
normed <- normalizeCounts(x, size.factors=centerSizeFactors(colSums(x)))
scoreGeneSet(normed, set=c(1,3,5,10,20,100))
```

scoreGeneSet.se *Score a gene set in a SummarizedExperiment*

Description

Compute a gene set activity score for each cell based on the expression values of the genes in the set, by calling `scoreGeneSet` on an assay of a `SummarizedExperiment`.

Usage

```
scoreGeneSet.se(
  x,
  set,
  block = NULL,
  num.threads = 1,
  more.score.args = list(),
  assay.type = "logcounts"
)
```

Arguments

<code>x</code>	A <code>SummarizedExperiment</code> object or one of its subclasses. Rows correspond to genes and columns correspond to cells.
<code>set</code>	Vector containing the gene set, see <code>?scoreGeneSet</code> for details.
<code>block</code>	Block assignment for each cell, passed to <code>scoreGeneSet</code> .
<code>num.threads</code>	Number of threads for <code>scoreGeneSet</code> .
<code>more.score.args</code>	Named list of further arguments to pass to <code>scoreGeneSet</code> .
<code>assay.type</code>	Integer or string specifying the relevant assay in <code>x</code> , usually containing log-normalized expression values.

Value

List containing scores, a numeric vector of the gene set scores across all cells in `x`; and weights, a numeric vector of weights for all genes in `set`.

Author(s)

Aaron Lun

Examples

```
# Defining a gene set of oligodendrocyte genes.
library(org.Mm.eg.db)
oligo.set <- select(org.Mm.eg.db, keytype="GO", keys="GO:0048709", columns="SYMBOL")
oligo.set <- unique(oligo.set$SYMBOL)

sce <- getTestRnaData.se("norm")
oligo.scores <- scoreGeneSet.se(sce, oligo.set)
summary(oligo.scores$scores)
```

scoreMarkers

Score marker genes

Description

Score marker genes for each group using a variety of effect sizes from pairwise comparisons between groups. This includes Cohen's *d*, the area under the curve (AUC), the difference in the means (delta-mean) and the difference in the proportion of detected cells (delta-detected). For each group, the strongest markers are those genes with the largest effect sizes (i.e., upregulated) when compared to all other groups.

Usage

```
scoreMarkers(
  x,
  groups,
  block = NULL,
  block.average.policy = c("mean", "quantile"),
  block.weight.policy = c("variable", "equal", "none"),
  variable.block.weight = c(0, 1000),
  block.quantile = 0.5,
  compute.group.mean = TRUE,
  compute.group.detected = TRUE,
  compute.delta.mean = TRUE,
  compute.delta.detected = TRUE,
  compute.cohens.d = TRUE,
  compute.auc = TRUE,
  compute.summary.min = TRUE,
  compute.summary.mean = TRUE,
  compute.summary.median = TRUE,
  compute.summary.max = TRUE,
  compute.summary.quantiles = NULL,
  compute.summary.min.rank = TRUE,
  threshold = 0,
  all.pairwise = FALSE,
  top.index.only = FALSE,
  min.rank.limit = 500,
  num.threads = 1
)
```

Arguments

<code>x</code>	A matrix-like object where rows correspond to genes or genomic features and columns correspond to cells. It is typically expected to contain log-expression values, e.g., from normalizeCounts .
<code>groups</code>	A vector specifying the group assignment for each cell in <code>x</code> .
<code>block</code>	Factor specifying the block of origin (e.g., batch, sample) for each cell in <code>x</code> . If provided, comparisons are performed within each block to ensure that block effects do not confound the estimates. The weighted average of the effect sizes across all blocks is reported for each gene. Alternatively NULL, if all cells are from the same block.
<code>block.average.policy</code>	String specifying the policy to use for average statistics across blocks. This can either be a (weighted) "mean" or a "quantile". Only used if <code>block</code> is not NULL.
<code>block.weight.policy</code>	String specifying the policy to use for weighting different blocks when computing the average for each statistic. See the argument of the same name in computeBlockWeights for more detail. Only used if <code>block</code> is not NULL.
<code>variable.block.weight</code>	Numeric vector of length 2, specifying the parameters for variable block weighting. See the argument of the same name in computeBlockWeights for more detail. Only used if <code>block</code> is not NULL and <code>block.weight.policy = "variable"</code> .
<code>block.quantile</code>	Number specifying the probability of the quantile of statistics across blocks. Defaults to 0.5, i.e., the median of per-block statistics. Only used if <code>block</code> is not NULL and <code>block.average.policy="quantile"</code> .
<code>compute.group.mean</code>	Logical scalar indicating whether to compute the group-wise mean expression for each gene.
<code>compute.group.detected</code>	Logical scalar indicating whether to compute the group-wise proportion of detected cells for each gene.
<code>compute.delta.mean</code>	Logical scalar indicating whether to compute the delta-means, i.e., the log-fold change when <code>x</code> contains log-expression values.
<code>compute.delta.detected</code>	Logical scalar indicating whether to compute the delta-detected, i.e., differences in the proportion of cells with detected expression.
<code>compute.cohens.d</code>	Logical scalar indicating whether to compute Cohen's d.
<code>compute.auc</code>	Logical scalar indicating whether to compute the AUC. Setting this to FALSE can improve speed and memory efficiency.
<code>compute.summary.min</code>	Boolean specifying whether to compute the minimum as a summary statistic for each effect size. Only used if <code>all.pairwise=FALSE</code> .
<code>compute.summary.mean</code>	Boolean specifying whether to compute the mean as a summary statistic for each effect size. Only used if <code>all.pairwise=FALSE</code> .
<code>compute.summary.median</code>	Boolean specifying whether to compute the median as a summary statistic for each effect size. Only used if <code>all.pairwise=FALSE</code> .

<code>compute.summary.max</code>	Boolean specifying whether to compute the maximum as a summary statistic for each effect size. Only used if <code>all.pairwise=FALSE</code> .
<code>compute.summary.quantiles</code>	Numeric scalars containing the probabilities of quantiles to compute as summary statistics for each effect size. If NULL, no quantiles are computed. Only used if <code>all.pairwise=FALSE</code> .
<code>compute.summary.min.rank</code>	Boolean specifying whether to compute the minimum rank as a summary statistic for each effect size. If NULL, no quantiles are computed. Only used if <code>all.pairwise=FALSE</code> .
<code>threshold</code>	Non-negative numeric scalar specifying the minimum threshold on the differences in means (i.e., the log-fold change, if <code>x</code> contains log-expression values). This is incorporated into the effect sizes for Cohen's <i>d</i> and the AUC. Larger thresholds will favor genes with large differences at the expense of genes with low variance that would otherwise have comparable effect sizes.
<code>all.pairwise</code>	Logical scalar indicating whether to report the effect sizes for every pairwise comparison between groups. Alternatively, an integer scalar indicating the number of top markers to report from each pairwise comparison between groups. If FALSE, only the summary statistics are reported.
<code>top.index.only</code>	Boolean indicating whether to only report the indices of the top genes when <code>all.pairwise</code> is an integer. This is more efficient when the effect sizes of the top genes are not required. Ignored for all other values of <code>all.pairwise</code> .
<code>min.rank.limit</code>	Integer scalar specifying the maximum value of the min-rank to report. Lower values improve memory efficiency at the cost of discarding information about lower-ranked genes. Only used if <code>all.pairwise=FALSE</code> .
<code>num.threads</code>	Integer scalar specifying the number of threads to use.

Value

A named list containing:

- `nrow`, integer specifying the number of rows in `x`.
- `row.names`, character vector or NULL containing the row names of `x`.
- `group.ids`, vector containing the identities of the unique groups.
- `mean`, a numeric matrix containing the mean expression for each group. Each row is a gene and each column is a group in `group.ids`. Omitted if `compute.group.mean=FALSE`.
- `detected`, a numeric matrix containing the proportion of detected cells in each group. Each row is a gene and each column is a group in `group.ids`. Omitted if `compute.group.detected=FALSE`.

If `all.pairwise=FALSE`, the list also contains:

- `cohens.d`, a list of [DataFrames](#) where each DataFrame corresponds to a group in `group.ids`. Each row of a DataFrame represents a gene, while each column contains a summary of Cohen's *d* from pairwise comparisons to all other groups. This includes `min`, `mean`, `median`, `max`, `quantile.*` and `min.rank` - check out [?summarizeEffects](#) for details. Omitted if `compute.cohens.d=FALSE`.
- `auc`, a list like `cohens.d` but containing the summaries of the AUCs from each pairwise comparison. Omitted if `compute.auc=FALSE`.
- `delta.mean`, a list like `cohens.d` but containing the summaries of the delta-mean from each pairwise comparison. Omitted if `compute.delta.mean=FALSE`.

- `delta.detected`, a list like `cohens.d` but containing the summaries of the delta-detected from each pairwise comparison. Omitted if `compute.delta.detected=FALSE`.

If `all.pairwise=TRUE`, the list also contains:

- `cohens.d`, a 3-dimensional numeric array containing the Cohen's d from each pairwise comparison between groups. The extents of the first two dimensions are equal to the number of groups in `group.ids`, while the extent of the final dimension is equal to the number of genes. The entry `cohens.d[i, j, k]` represents Cohen's d from the comparison of group `group.ids[j]` over group `group.ids[i]` for gene `k`. Omitted if `compute.cohens.d=FALSE`.
- `auc`, an array like `cohens.d` but containing the AUCs from each pairwise comparison. Omitted if `compute.auc=FALSE`.
- `delta.mean`, an array like `cohens.d` but containing the delta-mean from each pairwise comparison. Omitted if `compute.delta.mean=FALSE`.
- `delta.detected`, an array like `cohens.d` but containing the delta-detected from each pairwise comparison. Omitted if `compute.delta.detected=FALSE`.

If `all.pairwise` is an integer and `top.index.only=FALSE`, the list also contains:

- `cohens.d`, a list of list of [DataFrames](#) containing the top genes with the largest Cohen's d for each pairwise comparison. Specifically, `cohens.d[[i]][[j]]` is a [DataFrame](#) that contains the top `all.pairwise` genes from the comparison of group `group.ids[i]` over group `group.ids[j]`. Each [DataFrame](#) contains an `index` column, the row index of the gene; and an `effect` column, the Cohen's d for that gene. Omitted if `compute.cohens.d=FALSE`.
- `auc`, a list of list of [DataFrames](#) like `cohens.d` but containing the AUCs from each pairwise comparison. Omitted if `compute.auc=FALSE`.
- `delta.mean`, a list of list of [DataFrames](#) like `cohens.d` but containing the delta-mean from each pairwise comparison. Omitted if `compute.delta.mean=FALSE`.
- `delta.detected`, a list of list of [DataFrames](#) like `cohens.d` but containing the delta-detected from each pairwise comparison. Omitted if `compute.delta.detected=FALSE`.

If `all.pairwise` is an integer and `top.index.only=TRUE`, the list also contains:

- `cohens.d`, a list of list of integer vectors containing the row indices of the top genes with the largest Cohen's d for each pairwise comparison. Specifically, `cohens.d[[i]][[j]]` is a vector that contains the top `all.pairwise` genes from the comparison of group `group.ids[i]` over group `group.ids[j]`. Omitted if `compute.cohens.d=FALSE`.
- `auc`, a list of list of [DataFrames](#) like `cohens.d` but containing the AUCs from each pairwise comparison. Omitted if `compute.auc=FALSE`.
- `delta.mean`, a list of list of [DataFrames](#) like `cohens.d` but containing the delta-mean from each pairwise comparison. Omitted if `compute.delta.mean=FALSE`.
- `delta.detected`, a list of list of [DataFrames](#) like `cohens.d` but containing the delta-detected from each pairwise comparison. Omitted if `compute.delta.detected=FALSE`.

Choice of effect size

The delta-mean is the difference in the mean expression between groups. This is fairly straightforward to interpret - a positive delta-mean corresponds to increased expression in the first group compared to the second. The delta-mean can also be treated as the log-fold change if the input matrix contains log-transformed normalized expression values.

The delta-detected is the difference in the proportion of cells with detected expression between groups. This lies between 1 and -1, with the extremes occurring when a gene is silent in one group

and detected in all cells of the other group. For this interpretation, we assume that the input matrix contains non-negative expression values, where a value of zero corresponds to lack of detectable expression.

Cohen's d is the standardized difference between two groups. This is defined as the difference in the mean for each group scaled by the average standard deviation across the two groups. (Technically, we should use the pooled variance; however, this introduces some unintuitive asymmetry depending on the variance of the larger group, so we take a simple average instead.) A positive value indicates that the gene has increased expression in the first group compared to the second. Cohen's d is analogous to the t -statistic in a two-sample t -test and avoids spuriously large effect sizes from comparisons between highly variable groups. We can also interpret Cohen's d as the number of standard deviations between the two group means.

The area under the curve (AUC) is the probability that a randomly chosen observation in one group is greater than a randomly chosen observation in the other group. Values greater than 0.5 indicate that a gene is upregulated in the first group. The AUC is closely related to the U -statistic used in the Wilcoxon rank sum test. The key difference between the AUC and Cohen's d is that the former is less sensitive to the variance within each group, e.g., if two distributions exhibit no overlap, the AUC is the same regardless of the variance of each distribution. This may or may not be desirable as it improves robustness to outliers but reduces the information available to obtain a fine-grained ranking.

With a minimum change threshold

Setting a minimum change threshold (i.e., threshold) prioritizes genes with large shifts in expression instead of those with low variances. Currently, only positive thresholds are supported, which focuses on genes that are upregulated in the first group compared to the second. The effect size definitions are generalized when testing against a non-zero threshold:

- Cohen's d is redefined as the standardized difference between the difference in means and the specified threshold, analogous to the TREAT method from the **limma** package. Large positive values are only obtained when the observed difference in means is significantly greater than the threshold. For example, if we had a threshold of 2 and we obtained a Cohen's d of 3, this means that the observed difference in means was 3 standard deviations greater than 2. Note that a negative Cohen's d cannot be interpreted as downregulation, as the difference in means may still be positive but less than the threshold.
- The AUC is generalized to the probability of obtaining a random observation in one group that is greater than a random observation plus the threshold in the other group. For example, if we had a threshold of 2 and we obtained an AUC of 0.8, this means that, 80% of the random observations from the first group would be greater than a random observation from the second group by 2 or more. Again, AUCs below 0.5 cannot be interpreted as downregulation, as it may be caused by a positive shift that is less than the threshold.

See Also

The `score_markers_summary`, `score_markers_pairwise` and `score_markers_best` functions in https://libscran.github.io/scran_markers/. See their blocked equivalents (e.g., `score_markers_summary_block`) when `block` is specified.

[summarizeEffects](#), to summarize the pairwise effects returned when `all.pairwise=TRUE`.

[reportGroupMarkerStatistics](#), to consolidate the statistics for a single group into its own data frame.

[scoreMarkers.se](#), to score markers from a [SummarizedExperiment](#).

Examples

```
# Mocking a matrix:
library(Matrix)
x <- round(abs(rsparsematrix(1000, 100, 0.1) * 100))
normed <- normalizeCounts(x, size.factors=centerSizeFactors(colSums(x)))

# Compute marker summaries for each group:
g <- sample(letters[1:4], ncol(x), replace=TRUE)
scores <- scoreMarkers(normed, g)
names(scores)
head(scores$mean)
head(scores$cohens.d[["a"]])

# Report marker statistics for a single group:
reportGroupMarkerStatistics(scores, "b")
```

scoreMarkers.se

Score marker genes in a SummarizedExperiment

Description

Identify candidate marker genes based on effect sizes from pairwise comparisons between groups of cells, by calling `scoreMarkers` on an assay of a [SummarizedExperiment](#).

Usage

```
scoreMarkers.se(
  x,
  groups,
  block = NULL,
  num.threads = 1,
  more.marker.args = list(),
  assay.type = "logcounts",
  extra.columns = NULL,
  order.by = TRUE
)

formatScoreMarkersResult(marker.res, extra.columns = NULL, order.by = TRUE)

previewMarkers(
  marker.df,
  columns = c("mean", "detected", lfc = "delta.mean.mean"),
  pre.columns = NULL,
  post.columns = NULL,
  rows = 10,
  order.by = NULL,
  include.order.by = !is.null(order.by)
)
```

Arguments

x	A SummarizedExperiment object or one of its subclasses. Rows correspond to genes and columns correspond to cells.
groups	Group assignment for each cell, passed to scoreMarkers .
block	Block assignment for each cell, passed to scoreMarkers .
num.threads	Number of threads for marker scoring, passed to scoreMarkers .
more.marker.args	Named list of additional arguments to pass to scoreMarkers .
assay.type	Integer or string specifying the assay to use for differential comparisons, usually containing log-normalized expression values.
extra.columns	DataFrame containing extra columns to add each DataFrame . This should have the same number of rows as x. For scoreMarkers.se , this may also be a character vector specifying the columns of rowData to be added.
order.by	String specifying the column to order each DataFrame by. Alternatively TRUE, a column is automatically chosen from the effect size summaries. If NULL or FALSE, no ordering is performed.
marker.res	List containing the result of scoreMarkers .
marker.df	DataFrame containing the marker statistics for a single group.
columns	Character vector of the names of columns to retain in the preview. This may be named, in which the names are used as the column names.
pre.columns, post.columns	Character vector of the names of additional columns to retain in the preview. These are added before or after the columns in <code>columns</code> , for <code>pre.columns</code> and <code>post.columns</code> respectively.
rows	Integer specifying the number of rows to show. If NULL, all rows are returned.
include.order.by	Boolean indicating whether the column specified by <code>order.by</code> should be included in the output DataFrame . A string may also be supplied and will be treated as TRUE; the value of the string will be used as the column name in the output DataFrame .

Value

For `scoreMarkers.se` and `formatScoreMarkersResult`, a [List](#) of [DataFrames](#) is returned. Each [DataFrame](#) corresponds to a unique group in `groups`. Each row contains statistics for a gene in `x`, with the following columns:

- mean, the mean expression in the current group.
- detected, the proportion of cells with detected expression in the current group.
- `<effect>.<summary>`, a summary statistic for an effect size, e.g., `cohens.d.mean` contains the mean Cohen's d across comparisons involving the current group.

For `previewMarkers`, a [DataFrame](#) is returned containing the specified columns and rows.

Author(s)

Aaron Lun

Examples

```
sce <- getTestRnaData.se("cluster")
markers <- scoreMarkers.se(sce, sce$clusters)
previewMarkers(markers[["1"]])
```

subsampleByNeighbors *Subsample cells based on their neighbors*

Description

Subsample a dataset by selecting cells to represent all of their nearest neighbors. The aim is to preserve the relative density of the original dataset while guaranteeing representation of low-frequency subpopulations.

Usage

```
subsampleByNeighbors(
  x,
  num.neighbors = 20,
  min.remaining = 10,
  num.threads = 1,
  BNPARAM = AnnoyParam()
)
```

Arguments

- | | |
|---------------|--|
| x | <p>A numeric matrix where rows are dimensions and columns are cells, typically containing a low-dimensional representation from, e.g., runPca. Alternatively, an index constructed by buildIndex. Alternatively, a list containing existing nearest-neighbor search results. This should contain:</p> <ul style="list-style-type: none"> • <code>index</code>, an integer matrix where rows are neighbors and columns are cells. Each column contains 1-based indices for the nearest neighbors of the corresponding cell, ordered by increasing distance. • <code>distance</code>, a numeric matrix of the same dimensions as <code>index</code>, containing the distances to each of the nearest neighbors. <p>The number of neighbors should be equal to <code>num.neighbors</code>, otherwise a warning is raised.</p> |
| num.neighbors | Integer scalar specifying the number of neighbors to use. Larger values result in stronger downsampling. Ignored if <code>x</code> contains pre-computed neighbor search results. |
| min.remaining | Integer scalar specifying the minimum number of remaining neighbors that a cell must have in order to be considered for selection. This should be less than or equal to <code>num.neighbors</code> . Larger values result in stronger downsampling. |
| num.threads | Integer scalar specifying the number of threads to use for the nearest-neighbor search. Only used if <code>x</code> does not contain existing nearest-neighbor results. |
| BNPARAM | A BiocNeighborParam object specifying the algorithm to use. Only used if <code>x</code> does not contain existing nearest-neighbor results. |

Details

Starting from the densest region in the high-dimensional space, we select an observation for inclusion into the subsampled dataset. Every time we select an observation, we remove it and all of its nearest neighbors from the dataset. We then select the next observation with the most remaining neighbors, with ties broken by density; this is repeated until there are no more observations.

The premise is that each selected observation serves as a representative for its nearest neighbors. This ensures that the subsampled points are well-distributed across the original dataset. Low-frequency subpopulations will always have at least a few representatives if they are sufficiently distant from other subpopulations. We also preserve the relative density of the original dataset as more representatives will be generated from high-density regions.

Value

Integer vector with the indices of the selected cells in the subsample.

Author(s)

Aaron Lun

See Also

<https://libscrans.github.io/nenesub/>, for more details on the underlying algorithm.

Examples

```
x <- matrix(rnorm(10000), nrow=2)
keep <- subsampleByNeighbors(x, 10)
plot(x[1,], x[2,])
points(x[1,keep], x[2,keep], col="red")
legend('topright', col=c('black', 'red'), legend=c('all', 'subsample'), pch=1)
```

summarizeEffects

Summarize pairwise effect sizes for each group

Description

For each group, summarize the effect sizes for all pairwise comparisons to other groups. This yields a set of summary statistics that can be used to rank marker genes for each group.

Usage

```
summarizeEffects(
  effects,
  compute.summary.min = TRUE,
  compute.summary.mean = TRUE,
  compute.summary.median = TRUE,
  compute.summary.max = TRUE,
  compute.summary.quantiles = NULL,
  compute.summary.min.rank = TRUE,
  num.threads = 1
)
```

Arguments

<code>effects</code>	A 3-dimensional numeric containing the effect sizes from each pairwise comparison between groups. The extents of the first two dimensions are equal to the number of groups, while the extent of the final dimension is equal to the number of genes. The entry <code>[i, j, k]</code> represents the effect size from the comparison of group <code>j</code> against group <code>i</code> for gene <code>k</code> . See also the output of <code>scoreMarkers</code> with <code>all.pairwise=TRUE</code> .
<code>compute.summary.min</code>	Boolean specifying whether to compute the minimum as a summary statistic.
<code>compute.summary.mean</code>	Boolean specifying whether to compute the mean as a summary statistic.
<code>compute.summary.median</code>	Boolean specifying whether to compute the median as a summary statistic.
<code>compute.summary.max</code>	Boolean specifying whether to compute the maximum as a summary statistic.
<code>compute.summary.quantiles</code>	Numeric scalars containing the probabilities of quantiles to compute as summary statistics. If <code>NULL</code> , no quantiles are computed.
<code>compute.summary.min.rank</code>	Boolean specifying whether to compute the minimum rank as a summary statistic.
<code>num.threads</code>	Integer scalar specifying the number of threads to use.

Details

Each summary statistic can be used to prioritize different sets of marker genes for the group of interest, by ranking them in decreasing order according to said statistic:

- `min` contains the minimum effect size across all comparisons involving the group of interest. Genes with large values are upregulated in all comparisons. As such, it is the most stringent summary as markers will only have large values if they are uniquely upregulated in the group of interest compared to every other group.
- `mean` contains the mean effect size across all comparisons involving the group of interest. Genes with large values are upregulated on average compared to the other groups. This is a good general-purpose summary statistic.
- `median` contains the median effect size across all comparisons involving the group of interest. Genes with large values are upregulated compared to most (i.e., at least 50). Compared to the mean, this is more robust to outlier effects but less sensitive to strong effects in a minority of comparisons.
- `max` contains the maximum effect size across all comparisons involving the group of interest. Using this to define markers will focus on genes that are upregulated in at least one comparison. As such, it is the least stringent summary as markers can achieve large values if they are upregulated in the group of interest compared to any one other group.
- `quantile[[P]]` contains the quantile `P` across all comparisons involving the group of interest. This is a generalization of the minimum, median and maximum for arbitrary quantile probabilities. For example, a large `quantile[["20"]]` would mean that the gene is upregulated in the group of interest compared to 80

The exact definition of “large” depends on the choice of effect size. For signed effects like Cohen’s `d`, delta-mean and delta-detected, the value must be positive to be considered “large”. For the AUC,

a value greater than 0.5 is considered “large”. This interpretation is also affected by the choice of `threshold=` used to compute each effect size in `scoreMarkers`, e.g., a negative Cohen’s *d* cannot be interpreted as downregulation when the threshold is positive.

The `min.rank` is a more exotic summary statistic, containing the minimum rank for each gene across all comparisons involving the group of interest. This is defined by ranking the effect sizes across genes within each comparison, and then taking the minimum of these ranks across comparisons. Taking all genes with `min.rank <= T` will yield a set containing the top *T* genes from each comparison. The idea is to ensure that there are at least *T* genes that can distinguish the group of interest from any other group.

NaN effect sizes are allowed, e.g., if two groups do not exist in the same block for a blocked analysis in `scoreMarkers` with `block=`. This function will ignore NaN values when computing each summary. If all effects are NaN for a particular group, the summary statistic will also be NaN.

Value

List of `DataFrames` containing summary statistics for the effect sizes. Each `DataFrame` corresponds to a group, each row corresponds to a gene, and each column contains a summary statistic. If `compute.summary.quantiles` is provided, the “quantile” column is a nested `DataFrame` where each column corresponds to a probability in `compute.summary.quantiles`.

Author(s)

Aaron Lun

See Also

The `summarize_effects` function in https://libscran.github.io/scran_markers/. `scoreMarkers`, to compute the pairwise effects in the first place.

Examples

```
# Mocking a matrix:
library(Matrix)
x <- round(abs(rsparsematrix(1000, 100, 0.1) * 100))
normed <- normalizeCounts(x, size.factors=centerSizeFactors(colSums(x)))

g <- sample(letters[1:4], ncol(x), replace=TRUE)
effects <- scoreMarkers(normed, g, all.pairwise=TRUE)

summarized <- summarizeEffects(effects$cohens.d)
str(summarized)
```

testEnrichment

Test for gene set enrichment

Description

Perform a hypergeometric test for enrichment of gene sets in a list of interesting genes (e.g., markers).

Usage

```
testEnrichment(x, sets, universe = NULL, log = FALSE, num.threads = 1)
```

Arguments

x	Vector of identifiers for some interesting genes, e.g., symbols or Ensembl IDs. This is usually derived from a selection of top markers, e.g., from scoreMarkers .
sets	List of vectors of identifiers for the pre-defined gene sets. Each inner vector corresponds to a gene set and should contain the same type of identifiers as x.
universe	Vector of identifiers for the universe of genes in the dataset. x and each vector in sets will be subsetting to only include those genes in universe. If NULL, the universe is defined as the union of all genes in x and sets. Alternatively, an integer scalar specifying the number of genes in the universe. This is assumed to be greater than or equal to the number of unique genes in x and sets.
log	Logical scalar indicating whether to report log-transformed p-values. This may be desirable to avoid underflow at near-zero p-values.
num.threads	Integer scalar specifying the number of threads to use.

Value

[DataFrame](#) with one row per gene set and the following columns:

- overlap, the overlap between x and each entry of sets, i.e., the number of genes in the intersection.
- size, the set of each entry of sets.
- p.value, the (possibly log-transformed) p-value for overrepresentation of the gene set in x.

Author(s)

Aaron Lun

See Also

[phyper](#) and <https://libscrn.github.io/phyper/>, which is the basis for the underlying calculation.

Examples

```
testEnrichment(
  x=LETTERS[1:5],
  sets=list(
    first=LETTERS[1:10],
    second=LETTERS[1:5 * 2],
    third=LETTERS[10:20]
  ),
  universe=LETTERS
)
```

Index

- adt_quality_control, 3
- aggregateAcrossCells, 6, 7, 8, 10
- aggregateAcrossCells.se, 7, 7
- aggregateAcrossGenes, 7, 9, 11
- aggregateAcrossGenes.se, 10, 11
- aggregateColData
 - (aggregateAcrossCells.se), 7
- analyze, 12, 39
- analyze.se, 17

- BiocNeighborParam, 15, 21, 23, 41, 67, 75, 80, 84, 96
- buildIndex, 23, 67, 74, 78, 96
- buildSnnGraph, 14, 16, 22, 29, 31, 67

- centerSizeFactors, 14–16, 24, 27, 38, 53, 54, 56–58
- chooseHighlyVariableGenes, 14, 16, 25, 27, 28, 52, 70
- choosePseudoCount, 26, 55
- chooseRnaHvgs.se, 20, 21, 26, 27, 52, 73
- clusterGraph, 14, 16, 23, 24, 29, 31, 67
- clusterGraph.se, 24, 30, 31
- clusterKmeans, 14, 16, 32, 35
- clusterKmeans.se, 21, 34, 35
- colData, 8, 11, 19, 31, 35, 53, 54, 56–59, 61, 62, 69
- combineFactors, 6, 36
- computeAdtQcMetrics, 13, 15, 58, 59
- computeAdtQcMetrics
 - (adt_quality_control), 3
- computeBlockWeights, 37, 51, 71, 84, 87, 90
- computeClrm1Factors, 14, 16, 38, 53
- computeCrisprQcMetrics, 15, 60, 61
- computeCrisprQcMetrics
 - (crispr_quality_control), 44
- computeRnaQcMetrics, 13, 15, 61–63
- computeRnaQcMetrics
 - (rna_quality_control), 64
- computeRnaQcMetricsWithAltExps
 - (quickRnaQc.se), 61
- convertAnalyzeResults, 17, 39
- correctMnn, 14, 16, 40, 42, 71
- correctMnn.se, 20, 21, 41, 42

- countGroupsByBlock, 43
- crispr_quality_control, 44

- DataFrame, 4, 6, 8, 9, 22, 28, 36, 39, 44, 45, 52, 59–63, 65, 87, 91, 92, 95, 99, 100
- DataFrameList, 11
- DelayedArray, 55
- DelayedArray, LogNormalizedMatrixSeed-method
 - (LogNormalizedMatrix), 49
- dim, LogNormalizedMatrixSeed-method
 - (LogNormalizedMatrix), 49
- dimnames, LogNormalizedMatrixSeed-method
 - (LogNormalizedMatrix), 49

- extract_array, LogNormalizedMatrixSeed-method
 - (LogNormalizedMatrix), 49
- extract_sparse_array, LogNormalizedMatrixSeed-method
 - (LogNormalizedMatrix), 49

- fetchDataset, 48, 49
- filterAdtQcMetrics, 15, 58, 59
- filterAdtQcMetrics
 - (adt_quality_control), 3
- filterCrisprQcMetrics, 15, 60
- filterCrisprQcMetrics
 - (crispr_quality_control), 44
- filterRnaQcMetrics, 15, 61, 62
- filterRnaQcMetrics
 - (rna_quality_control), 64
- findKNN, 74, 78
- fitVarianceTrend, 46, 52
- formatComputeAdtQcMetricsResult
 - (quickAdtQc.se), 58
- formatComputeCrisprQcMetricsResult
 - (quickCrisprQc.se), 60
- formatComputeRnaQcMetricsResult
 - (quickRnaQc.se), 61
- formatModelGeneVariancesResult
 - (chooseRnaHvgs.se), 27
- formatScoreMarkersResult
 - (scoreMarkers.se), 94

- getTestAdtData.se (getTestData.se), 48
- getTestCrisprData.se (getTestData.se), 48

- getTestData.se, 48
- getTestRnaData.se (getTestData.se), 48
- igraph, 29
- initializeCpp, 32, 54, 55
- initializeCpp, LogNormalizedMatrixSeed-method (LogNormalizedMatrix), 49
- is_sparse, LogNormalizedMatrixSeed-method (LogNormalizedMatrix), 49
- kmeans, 33
- List, 11, 36, 95
- LogNormalizedMatrix, 49
- LogNormalizedMatrix-class (LogNormalizedMatrix), 49
- LogNormalizedMatrixSeed (LogNormalizedMatrix), 49
- LogNormalizedMatrixSeed-class (LogNormalizedMatrix), 49
- matrixClass, LogNormalizedMatrixSeed-method (LogNormalizedMatrix), 49
- mcols, 11
- metadata, 8, 31, 35, 59–62, 69, 73, 85
- modelGeneVariances, 14, 16, 26–28, 48, 50
- normalizeAdtCounts.se, 19, 21, 38, 53
- normalizeCounts, 14, 16, 24, 26, 27, 49, 50, 53, 54, 56–58, 70, 82, 90
- normalizeCrisprCounts.se, 20, 21, 56
- normalizeRnaCounts.se, 19, 21, 25, 55, 57, 73
- phyper, 100
- previewMarkers (scoreMarkers.se), 94
- quickAdtQc.se, 5, 19, 21, 58
- quickCrisprQc.se, 19, 21, 46, 60
- quickRnaQc.se, 19, 21, 61, 66
- reducedDim, 31, 35, 42, 69, 73, 77, 81, 85
- reducedDims, 42, 69, 73, 85
- reportGroupMarkerStatistics, 63, 93
- rna_quality_control, 64
- rowData, 9, 11, 95
- runAllNeighborSteps, 14, 16, 67, 68, 69
- runAllNeighborSteps.se, 20, 21, 68, 68
- runPca, 14, 16, 23, 25, 32, 40, 67, 70, 72–74, 78, 96
- runPca.se, 20, 21, 72, 72
- runTsne, 14, 16, 67, 74, 77
- runTsne.se, 76, 76
- runUmap, 14, 16, 67, 76, 77, 81
- runUmap.se, 81, 81
- sanitizeSizeFactors, 54, 82
- scaleByNeighbors, 14, 16, 83, 85
- scaleByNeighbors.se, 20, 21, 84, 85
- scoreGeneSet, 86, 88
- scoreGeneSet.se, 87, 88
- scoreMarkers, 15–17, 64, 89, 94, 95, 98–100
- scoreMarkers.se, 21, 93, 94
- scraper (scraper-package), 3
- scraper-package, 3
- SingleCellExperiment, 8, 17, 22, 24, 30, 31, 34, 35, 39, 41, 42, 49, 53, 62, 68, 69, 73, 76, 77, 81, 84, 85
- subsampleByNeighbors, 96
- suggestAdtQcThresholds, 13, 15, 58, 59
- suggestAdtQcThresholds (adt_quality_control), 3
- suggestCrisprQcThresholds, 14, 15, 60
- suggestCrisprQcThresholds (crispr_quality_control), 44
- suggestRnaQcThresholds, 13, 15, 60–62
- suggestRnaQcThresholds (rna_quality_control), 64
- SummarizedExperiment, 5, 7, 8, 10, 11, 13, 15, 18, 25–28, 46, 52, 53, 55–62, 66, 72, 73, 87, 88, 93–95
- summarizeEffects, 64, 91, 93, 97
- table, 43
- testEnrichment, 99
- tsnePerplexityToNeighbors (runTsne), 74
- type, LogNormalizedMatrixSeed-method (LogNormalizedMatrix), 49