

Package ‘batchelor’

April 15, 2020

Version 1.2.4

Date 2020-01-01

Title Single-Cell Batch Correction Methods

Depends SingleCellExperiment

Imports SummarizedExperiment, S4Vectors, BiocGenerics, Rcpp, stats,
methods, utils, BiocNeighbors, BiocSingular, Matrix,
DelayedArray, DelayedMatrixStats, scater, BiocParallel

Suggests testthat, BiocStyle, knitr, beachmat, scan, scRNAseq

biocViews Sequencing, RNASeq, Software, GeneExpression,
Transcriptomics, SingleCell, BatchEffect, Normalization

LinkingTo Rcpp, beachmat

Description

Implements a variety of methods for batch correction of single-cell (RNA sequencing) data. This includes methods based on detecting mutually nearest neighbors, as well as several efficient variants of linear regression of the log-expression values. Functions are also provided to perform global rescaling to remove differences in depth between batches, and to perform a principal components analysis that is robust to differences in the numbers of cells across batches.

License GPL-3

NeedsCompilation yes

VignetteBuilder knitr

SystemRequirements C++11

RoxygenNote 7.0.2

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

git_branch RELEASE_3_10

git_last_commit 24cbcd

git_last_commit_date 2020-01-01

Date/Publication 2020-04-14

Author Aaron Lun [aut, cre],
Laleh Haghverdi [ctb]

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

R topics documented:

batchCorrect	2
batchelor-restrict	5
BatchelorParam-class	6
checkBatchConsistency	7
correctExperiments	8
cosineNorm	10
divideIntoBatches	12
fastMNN	13
findMutualNN	19
mnnCorrect	21
multiBatchNorm	26
multiBatchPCA	28
noCorrect	31
reducedMNN	32
regressBatches	34
rescaleBatches	36
Index	38

batchCorrect	<i>Batch correction methods</i>
--------------	---------------------------------

Description

A common interface for single-cell batch correction methods.

Usage

```
batchCorrect(
  ...,
  batch = NULL,
  restrict = NULL,
  subset.row = NULL,
  correct.all = FALSE,
  assay.type = NULL,
  get.spikes = FALSE,
  PARAM
)

## S4 method for signature 'ClassicMnnParam'
batchCorrect(
  ...,
  batch = NULL,
  restrict = NULL,
  subset.row = NULL,
  correct.all = FALSE,
  assay.type = "logcounts",
  get.spikes = FALSE,
  PARAM
)
```

```
## S4 method for signature 'FastMnnParam'
batchCorrect(
  ...,
  batch = NULL,
  restrict = NULL,
  subset.row = NULL,
  correct.all = FALSE,
  assay.type = "logcounts",
  get.spikes = FALSE,
  PARAM
)

## S4 method for signature 'RescaleParam'
batchCorrect(
  ...,
  batch = NULL,
  restrict = NULL,
  subset.row = NULL,
  correct.all = FALSE,
  assay.type = "logcounts",
  get.spikes = FALSE,
  PARAM
)

## S4 method for signature 'RegressParam'
batchCorrect(
  ...,
  batch = NULL,
  restrict = NULL,
  subset.row = NULL,
  correct.all = FALSE,
  assay.type = "logcounts",
  get.spikes = FALSE,
  PARAM
)

## S4 method for signature 'NoCorrectParam'
batchCorrect(
  ...,
  batch = NULL,
  restrict = NULL,
  subset.row = NULL,
  correct.all = FALSE,
  assay.type = "logcounts",
  get.spikes = FALSE,
  PARAM
)
```

Arguments

... One or more matrix-like objects containing single-cell gene expression matrices. Alternatively, one or more [SingleCellExperiment](#) objects can be supplied.

	If multiple objects are supplied, each object is assumed to contain all and only cells from a single batch. Objects of different types can be mixed together. If a single object is supplied, <code>batch</code> should also be specified.
<code>batch</code>	A factor specifying the batch of origin for each cell if only one batch is supplied in <code>...</code> . This will be ignored if two or more batches are supplied.
<code>restrict</code>	A list of length equal to the number of objects in <code>...</code> . Each entry of the list corresponds to one batch and specifies the cells to use when computing the correction.
<code>subset.row</code>	A vector specifying the subset of genes to use for correction. Defaults to <code>NULL</code> , in which case all genes are used.
<code>correct.all</code>	A logical scalar indicating whether to return corrected expression values for all genes, even if <code>subset.row</code> is set. Used to ensure that the output is of the same dimensionality as the input.
<code>assay.type</code>	A string or integer scalar specifying the assay to use for correction. Only used for <code>SingleCellExperiment</code> inputs.
<code>get.spikes</code>	Deprecated, a logical scalar indicating whether to retain rows corresponding to spike-in transcripts. Only used for <code>SingleCellExperiment</code> inputs.
<code>PARAM</code>	A BatchelorParam object specifying the batch correction method to dispatch to, and the parameters with which it should be run. ClassicMnnParam will dispatch to <code>mnnCorrect</code> ; FastMnnParam will dispatch to <code>fastMNN</code> ; RescaleParam will dispatch to <code>rescaleBatches</code> ; and RegressParam will dispatch to <code>regressBatches</code> .

Details

Users can pass parameters to each method directly via `...` or via the constructors for `PARAM`. While there is no restriction on which parameters go where, we recommend only passing data-agnostic and method-specific parameters to `PARAM`. Data-dependent parameters - and indeed, the data themselves - should be passed in via `...`. This means that different data sets can be used without modifying `PARAM`, and allows users to switch to a different algorithm by only changing `PARAM`.

Note that `get.spikes=FALSE` effectively modifies `subset.row` to exclude spike-in transcripts when `SingleCellExperiment` inputs are supplied. This means that the reported `SingleCellExperiment` will not, by default, contain corrected expression values for spike-in transcripts unless `get.spikes=TRUE`.

Value

A `SingleCellExperiment` where the first assay contains corrected gene expression values for all genes. Corrected values should be returned for all genes if `subset.row=NULL` or if `correct.all=TRUE`; otherwise they should only be returned for the genes in the subset.

Cells should be reported in the same order that they are supplied. In cases with multiple batches, the cell identities are simply concatenated from successive objects in their specified order, i.e., all cells from the first object (in their provided order), then all cells from the second object, and so on. For a single input object, cells should be reported in the same order as the input.

The `colData` slot should contain `batch`, a vector specifying the batch of origin for each cell.

Author(s)

Aaron Lun

See Also

[BatchelorParam](#) classes to determine dispatch.

Examples

```
B1 <- matrix(rnorm(10000), ncol=50) # Batch 1
B2 <- matrix(rnorm(10000), ncol=50) # Batch 2

# Switching easily between batch correction methods.
m.out <- batchCorrect(B1, B2, PARAM=ClassicMnnParam())
f.out <- batchCorrect(B1, B2, PARAM=FastMnnParam(d=20))
r.out <- batchCorrect(B1, B2, PARAM=RescaleParam(pseudo.count=0))
n.out <- batchCorrect(B1, B2, PARAM=NoCorrectParam())
```

batchelor-restrict *Using restriction*

Description

Using restriction

Motivation

It is possible to compute the correction using only a subset of cells in each batch, and then extrapolate that correction to all other cells. This may be desirable in experimental designs where a control set of cells from the same source population were run on different batches. Any difference in the controls must be artificial in origin and can be directly removed without making further biological assumptions. Similarly, if certain cells are known to be of a batch-specific subpopulation, it may be desirable to exclude them to ensure that they are not inadvertently used during the batch correction.

Setting the restrict argument

To perform restriction, users should set `restrict` to specify the subset of cells in each batch to be used for correction. This should be set to a list of length equal to the number of objects passed to the `...` argument of the batch correction function. Each element of this list should be a subsetting vector to be applied to the columns of the corresponding batch. A NULL element indicates that all the cells from a batch should be used. In situations where one input object contains multiple batches, `restrict` should simply a list containing a single subsetting vector for that object.

Correction functions that support `restrict` will only use the restricted subset of cells in each batch to perform the correction. For example, `fastMNN` will only use the restricted cells to identify MNN pairs and the center of the orthogonalization. However, it will apply the correction to all cells in each batch - hence the extrapolation. This means that the output is always of the same dimensionality, regardless of whether `restrict` is specified.

As a general rule, users can expect the corrected values in the restricted cells to be the same as if the inputs were directly subsetted to only contain those cells (see Examples). This is appealing as it demonstrates that correction only uses information from the restricted subset of cells. If batch correction functions do not follow this rule, they will explicitly state so, e.g., in `?fastMNN`.

Author(s)

Aaron Lun

See Also

[rescaleBatches](#), [regressBatches](#), [fastMNN](#) and [mnnCorrect](#), as examples of batch correction methods that support restriction.

Examples

```
means <- 2^rgamma(1000, 2, 1)
A1 <- matrix(rpois(10000, lambda=means), ncol=50) # Batch 1
A2 <- matrix(rpois(10000, lambda=means*runif(1000, 0, 2)), ncol=50) # Batch 2

B1 <- log2(A1 + 1)
B2 <- log2(A2 + 1)
out <- regressBatches(B1, B2, restrict=list(1:10, 1:10))
assay(out)[,c(1:10, 50+1:10)]

# Compare to actual subsetting:
out.sub <- regressBatches(B1[,1:10], B2[,1:10])
assay(out.sub)
```

BatchelorParam-class *BatchelorParam methods*

Description

Constructors and methods for the batchelor parameter classes.

Usage

```
ClassicMnnParam(...)
FastMnnParam(...)
RescaleParam(...)
RegressParam(...)
NoCorrectParam(...)
```

Arguments

... Named arguments to pass to individual methods upon dispatch. These should not include arguments named in the [batchCorrect](#) generic.

Details

BatchelorParam objects are intended to store method-specific parameter settings to pass to the [batchCorrect](#) generic. These values should refer to data-agnostic parameters; parameters that depend on data (or the data itself) should be specified directly in the [batchCorrect](#) call.

The BatchelorParam classes are all derived from [SimpleList](#) objects and have the same available methods, e.g., `[[`, `$`. These can be used to access or modify the object after construction.

Note that the BatchelorParam class itself is not useful and should not be constructed directly. Instead, users should use the constructors shown above to create instances of the desired subclass.

Value

The constructors will return a BatchelorParam object of the specified subclass, containing parameter settings for the corresponding batch correction method.

Author(s)

Aaron Lun

See Also

[batchCorrect](#), where the BatchelorParam objects are used for dispatch to individual methods.

Examples

```
# Specifying the number of neighbors, dimensionality.
fp <- FastMnnParam(k=20, d=10)
fp

# List-like behaviour:
fp$k
fp$k <- 10
fp$k
```

checkBatchConsistency *Check batch inputs*

Description

Utilities to check inputs into batch correction functions.

Usage

```
checkBatchConsistency(batches, cells.in.columns = TRUE)
```

```
checkSpikeConsistency(batches)
```

```
checkIfSCE(batches)
```

```
checkRestrictions(batches, restrictions, cells.in.columns = TRUE)
```

Arguments

batches A list of batches, usually containing gene expression matrices or [SingleCellExperiment](#) objects.

cells.in.columns A logical scalar specifying whether batches contain cells in the columns.

restrictions A list of length equal to batches, specifying the cells in each batch that should be used for correction.

Details

These functions are intended for internal use and other package developers.

`checkBatchConsistency` will check whether the input batches are consistent with respect to the size of the dimension containing features (i.e., not cells). It will also verify that the dimension names are consistent, to avoid problems from variable ordering of rows/columns in the inputs.

`checkSpikeConsistency` will check whether the spike-in information is consistent across all batches. This only works for `SingleCellExperiment` objects, so one should only run this function if `checkIfSCE` returns TRUE.

`checkRestrictions` will check whether restrictions are consistent with the supplied batches, in terms of the length and names of the two lists. It will also check that each batch contains at least one usable cell after restriction.

Value

`checkBatchConsistency` and `checkSpikeConsistency` will return an invisible NULL if there are no errors.

`checkIfSCE` will return a logical vector specifying whether each element of batches is a `SingleCellExperiment` objects.

`checkRestrictions` will return NULL if `restrictions=NULL`. Otherwise, it will return a list by taking restrictions and converting each non-NULL element into an integer subsetting vector.

Author(s)

Aaron Lun

See Also

[divideIntoBatches](#)

Examples

```
checkBatchConsistency(list(cbind(1:5), cbind(1:5, 2:6)))
try( # fails
  checkBatchConsistency(list(cbind(1:5), cbind(1:4, 2:5)))
)
```

correctExperiments

Correct SingleCellExperiment objects

Description

Apply a correction to multiple [SingleCellExperiment](#) objects, while also combining the assay data and column metadata for easy use.

Usage

```
correctExperiments(
  ...,
  batch = NULL,
  restrict = NULL,
  subset.row = NULL,
  correct.all = FALSE,
  assay.type = "logcounts",
  PARAM = FastMnnParam(),
  combine.assays = NULL,
  combine.coldata = NULL,
  include.rowdata = TRUE
)
```

Arguments

...	One or more SingleCellExperiment objects. If multiple objects are supplied, each object is assumed to contain all and only cells from a single batch. If a single object is supplied, batch should also be specified.
batch	A factor specifying the batch of origin for each cell if only one batch is supplied in This will be ignored if two or more batches are supplied.
restrict	A list of length equal to the number of objects in Each entry of the list corresponds to one batch and specifies the cells to use when computing the correction.
subset.row	A vector specifying the subset of genes to use for correction. Defaults to NULL, in which case all genes are used.
correct.all	A logical scalar indicating whether to return corrected expression values for all genes, even if subset.row is set. Used to ensure that the output is of the same dimensionality as the input.
assay.type	A string or integer scalar specifying the assay to use for correction.
PARAM	A BatchelorParam object specifying the batch correction method to dispatch to, and the parameters with which it should be run. ClassicMnnParam will dispatch to mnnCorrect ; FastMnnParam will dispatch to fastMNN ; RescaleParam will dispatch to rescaleBatches ; and RegressParam will dispatch to regressBatches .
combine.assays	Character vector specifying the assays from each entry of ... to combine together without correction. By default, any named assay that is present in all entries of ... are combined. This can be set to character(0) to avoid combining any assays.
combine.coldata	Character vector specifying the column metadata fields from each entry of ... to combine together. By default, any column metadata field that is present in all entries of ... is combined. This can be set to character(0) to avoid combining any metadata.
include.rowdata	Logical scalar indicating whether the function should attempt to include rowRanges .

Details

This function makes it easy to retain information from the original [SingleCellExperiment](#) objects in the post-merge object. Operations like differential expression analyses can be easily performed

on the uncorrected expression values, while common annotation can be leveraged in cell-based analyses like clustering.

Additional assays may be added to the merged object, depending on `combine.assays`. This will usually contain uncorrected values from each batch that have been simply cbinded together. If `combine.assays` contains a field that overlaps with the name of the corrected assay from [batchCorrect](#), a warning will be raised and the corrected assay will be preferentially retained.

Any column metadata fields that are shared will also be included in the merged object by default (tunable by setting `combine.coldata`). If any existing field is named "batch", it will be ignored in favor of that produced by [batchCorrect](#) and a warning is emitted.

Row metadata is only included in the merged object if `include.rowdata=TRUE` and all row metadata objects are identical across objects in If not, a warning is emitted and no row metadata is attached to the merged object.

Value

A `SingleCellExperiment` containing the merged expression values in the first assay and a batch column metadata field specifying the batch of origin for each cell, as described in [batchCorrect](#).

Author(s)

Aaron Lun

See Also

[batchCorrect](#), which does the correction inside this function.

[noCorrect](#), used to combine uncorrected values for the other assays.

Examples

```
sce1 <- scater::mockSCE()
sce1 <- scater::logNormCounts(sce1)
sce2 <- scater::mockSCE()
sce2 <- scater::logNormCounts(sce2)

f.out <- correctExperiments(sce1, sce2)
colData(f.out)
assayNames(f.out)
```

cosineNorm

Cosine normalization

Description

Perform cosine normalization on the column vectors of an expression matrix.

Usage

```
cosineNorm(
  x,
  mode = c("matrix", "all", "l2norm"),
  subset.row = NULL,
  BPPARAM = SerialParam()
)
```

Arguments

<code>x</code>	A gene expression matrix with cells as columns and genes as rows.
<code>mode</code>	A string specifying the output to be returned.
<code>subset.row</code>	A vector specifying which features to use to compute the L2 norm.
<code>BPPARAM</code>	A BiocParallelParam object specifying how parallelization is to be performed. Only used when <code>x</code> is a DelayedArray object.

Details

Cosine normalization removes scaling differences between expression vectors. In the context of batch correction, this is usually applied to remove differences between batches that are normalized separately. For example, [fastMNN](#) uses this function on the log-expression vectors by default.

Technically, separate normalization introduces scaling differences in the normalized expression, which should manifest as a shift in the log-transformed expression. However, in practice, single-cell data will contain many small counts (where the log function is near-linear) or many zeroes (which remain zero when the pseudo-count is 1). In these applications, scaling differences due to separate normalization are better represented as scaling differences in the log-transformed values.

If applied to the raw count vectors, cosine normalization is similar to library size-related (i.e., L1) normalization. However, we recommend using dedicated methods for computing size factors to normalize raw count data.

While the default is to directly return the cosine-normalized matrix, it may occasionally be desirable to obtain the L2 norm, e.g., to apply an equivalent normalization to other matrices. This can be achieved by setting `mode` accordingly.

The function will return a [DelayedMatrix](#) if `x` is a [DelayedMatrix](#). This aims to delay the calculation of cosine-normalized values for very large matrices.

Value

If `mode="matrix"`, a double-precision matrix of the same dimensions as `X` is returned, containing cosine-normalized values.

If `mode="l2norm"`, a double-precision vector is returned containing the L2 norm for each cell.

If `mode="all"`, a named list is returned containing the fields `"matrix"` and `"l2norm"`, which are as described above.

Author(s)

Aaron Lun

See Also

[mnnCorrect](#) and [fastMNN](#), where this function gets used.

Examples

```
A <- matrix(rnorm(1000), nrow=10)
str(cosineNorm(A))
str(cosineNorm(A, mode="l2norm"))
```

divideIntoBatches	<i>Divide into batches</i>
-------------------	----------------------------

Description

Divide a single input object into multiple separate objects according to their batch of origin.

Usage

```
divideIntoBatches(x, batch, byrow = FALSE, restrict = NULL)
```

Arguments

x	A matrix-like object where one dimension corresponds to cells and another represents features.
batch	A factor specifying the batch to which each cell belongs.
byrow	A logical scalar indicating whether rows correspond to cells.
restrict	A subsetting vector specifying which cells should be used for correction.

Details

This function is intended for internal use and other package developers. It splits a single input object into multiple batches, allowing developers to use the same code for the scenario where batch is supplied with a single input.

Value

A list containing:

- `batches`, a named list of matrix-like objects where each element corresponds to a level of batch and contains all cells from that batch.
- `reorder`, an integer vector to be applied to the combined batches to recover the ordering of cells in `x`.
- `restricted`, a named list of integer vectors specifying which cells are to be used for correction. Set to `NULL` if the input `restrict` was also `NULL`.

Author(s)

Aaron Lun

Examples

```
X <- matrix(rnorm(1000), ncol=100)
out <- divideIntoBatches(X, sample(3, 100, replace=TRUE))
names(out)

# Recovering original order.
Y <- do.call(cbind, out$batches)
Z <- Y[,out$reorder]
all.equal(Z, X) # should be TRUE.
```

fastMNN

*Fast mutual nearest neighbors correction***Description**

Correct for batch effects in single-cell expression data using a fast version of the mutual nearest neighbors (MNN) method.

Usage

```
fastMNN(
  ...,
  batch = NULL,
  k = 20,
  prop.k = NULL,
  restrict = NULL,
  cos.norm = TRUE,
  ndist = 3,
  d = 50,
  weights = NULL,
  merge.order = NULL,
  auto.merge = FALSE,
  auto.order = NULL,
  min.batch.skip = 0,
  subset.row = NULL,
  correct.all = FALSE,
  pc.input = FALSE,
  assay.type = "logcounts",
  get.spikes = FALSE,
  use.dimred = NULL,
  BSPARAM = IrlbaParam(deferred = TRUE),
  BNPARAM = KmknnParam(),
  BPPARAM = SerialParam()
)
```

Arguments

... One or more log-expression matrices where genes correspond to rows and cells correspond to columns. Each matrix should contain the same number of rows, corresponding to the same genes in the same order.

Alternatively, one or more [SingleCellExperiment](#) objects can be supplied containing a log-expression matrix in the `assay.type` assay. Note the same restrictions described above for gene expression matrix inputs.

If multiple objects are supplied, each object is assumed to contain all and only cells from a single batch. Objects of different types can be mixed together. If a single object is supplied, `batch` should also be specified.

<code>batch</code>	A factor specifying the batch of origin for all cells when only a single object is supplied in <code>...</code> . This is ignored if multiple objects are present.
<code>k</code>	An integer scalar specifying the number of nearest neighbors to consider when identifying MNNs.
<code>prop.k</code>	A numeric scalar in (0, 1) specifying the proportion of cells in each dataset to use for mutual nearest neighbor searching. If set, <code>k</code> for the search in each batch is redefined as $\max(k, \text{prop.k} * N)$ where <code>N</code> is the number of cells in that batch.
<code>restrict</code>	A list of length equal to the number of objects in <code>...</code> . Each entry of the list corresponds to one batch and specifies the cells to use when computing the correction.
<code>cos.norm</code>	A logical scalar indicating whether cosine normalization should be performed on the input data prior to PCA.
<code>ndist</code>	A numeric scalar specifying the threshold beyond which neighbours are to be ignored when computing correction vectors. Each threshold is defined as a multiple of the number of median distances.
<code>d</code>	Number of dimensions to use for dimensionality reduction in multiBatchPCA .
<code>weights</code>	Numeric scalar of weights to use in multiBatchPCA .
<code>merge.order</code>	An integer vector containing the linear merge order of batches in <code>...</code> . Alternatively, a list of lists representing a tree structure specifying a hierarchical merge order.
<code>auto.merge</code>	Logical scalar indicating whether to automatically identify the “best” merge order.
<code>auto.order</code>	Deprecated, use <code>merge.order</code> or <code>auto.merge</code> instead.
<code>min.batch.skip</code>	Numeric scalar specifying the minimum relative magnitude of the batch effect, below which no correction will be performed at a given merge step.
<code>subset.row</code>	A vector specifying which features to use for correction.
<code>correct.all</code>	Logical scalar indicating whether a rotation matrix should be computed for genes not in <code>subset.row</code> .
<code>pc.input</code>	Deprecated, use reducedMNN instead.
<code>assay.type</code>	A string or integer scalar specifying the assay containing the log-expression values. Only used for SingleCellExperiment inputs.
<code>get.spikes</code>	Deprecated.
<code>use.dimred</code>	Deprecated, use reducedMNN instead.
<code>BSPARAM</code>	A BiocSingularParam object specifying the algorithm to use for PCA. This uses a fast approximate algorithm from irlba by default, see multiBatchPCA for details.
<code>BNPARAM</code>	A BiocNeighborParam object specifying the nearest neighbor algorithm.
<code>BPPARAM</code>	A BiocParallelParam object specifying whether the PCA and nearest-neighbor searches should be parallelized.

Details

This function provides a variant of the `mnnCorrect` function, modified for speed and more robust performance. In particular:

- It performs a multi-sample PCA via `multiBatchPCA` and subsequently performs all calculations in the PC space. This reduces computational work and provides some denoising for improved neighbour detection. As a result, though, the corrected output cannot be interpreted on a gene level and is useful only for cell-level comparisons, e.g., clustering and visualization.
- The correction vector for each cell is directly computed from its k nearest neighbours in the same batch. Specifically, only the k nearest neighbouring cells that *also* participate in MNN pairs are used. Each MNN-participating neighbour is weighted by distance from the current cell, using a tricube scheme with bandwidth equal to the median distance multiplied by `ndist`. This ensures that the correction vector only uses information from the closest cells, improving the fidelity of local correction.
- Issues with “kissing” are avoided with a two-step procedure that removes variation along the batch effect vector. First, the average correction vector across all MNN pairs is computed. Cell coordinates are adjusted such that all cells in a single batch have the same position along this vector. The correction vectors are then recalculated with the adjusted coordinates (but the same MNN pairs).

The default setting of `cos.norm=TRUE` provides some protection against differences in scaling between log-expression matrices from batches that are normalized separately (see `cosineNorm` for details). However, if possible, we recommend using the output of `multiBatchNorm` as input to `fastMNN`. This will equalize coverage on the count level before the log-transformation, which is a more accurate rescaling than cosine normalization on the log-values.

The `batch` argument allows users to easily perform batch correction when all cells have already been combined into a single object. This avoids the need to manually split the matrix or `SingleCellExperiment` object into separate objects for input into `fastMNN`. In this situation, the order of input batches is defined by the order of levels in `batch`.

Value

A `SingleCellExperiment` is returned where each row is a gene and each column is a cell. This contains:

- A corrected matrix in the `reducedDims` slot, containing corrected low-dimensional coordinates for each cell. This has number of columns equal to `d` and number of rows equal to the total number of cells in `...`
- A batch column in the `colData` slot, containing the batch of origin for each row (i.e., cell) in corrected.
- A rotation column in the `rowData` slot, containing the rotation matrix used for the PCA. This has `d` columns and number of rows equal to the number of genes to report (see the “Choice of genes” section).
- A reconstructed matrix in the `assays` slot, containing the low-rank reconstruction of the expression matrix. This can be interpreted as per-gene corrected log-expression values (after cosine normalization, if `cos.norm=TRUE`) but should not be used for quantitative analyses. This has number of rows equal to the number of input genes if `subset.row=NULL` or `correct.all=TRUE`, otherwise each row corresponds to a gene in `subset.row`.

Cells in the output object are always ordered in the same manner as supplied in `...`. For a single input object, cells will be reported in the same order as they are arranged in that object. In cases with multiple input objects, the cell identities are simply concatenated from successive objects, i.e.,

all cells from the first object (in their provided order), then all cells from the second object, and so on. This is true regardless of the value of `merge.order` and `auto.merge`, which only affects the internal merge order of the batches.

The metadata of the output object contains `merge.info`, a [DataFrame](#) of diagnostic information about each merge step. See the “Merge diagnostics” section for more details.

Controlling the merge order

By default, batches are merged in the user-supplied order, i.e., the first batch is merged with the second batch, the third batch is merged with the combined first-second batch, the fourth batch is merged with the combined first-second-third batch and so on. We refer to this approach as a progressive merge.

If `merge.order` is an integer vector, it is treated as an ordering permutation with which to perform a progressive merge. For example, if `merge.order=c(4, 1, 3, 2)`, batches 4 and 1 in . . . are merged first; batch 3 is merged with the combined 4+1 batch; and then batch 2 is merged with the combined 4+1+3 batch. This is often more convenient than changing the order manually in . . . , which would alter the order of batches in the output corrected matrix.

If `merge.order` is a character vector, it is treated as an ordering permutation for named batches.

If `merge.order` is a nested list, it is treated as a tree that specifies a hierarchical merge. Each element of the list should either be a string or integer scalar, corresponding to a leaf node that specifies a batch; or another list, corresponding to an internal node that should contain at least two children; or an integer or character vector of length 2 or more, again corresponding to an internal node.

- For example, `list(list(1,2),list(3,4))` indicates that batch 1 should be merged with batch 2; batch 3 should be merged with batch 4; and that, finally, the combined batches 1+2 and 3+4 should be merged.
- More than two children per node are supported and will result in a progressive merge within that node. For example, `list(list(1,2,3),list(4,5,6))` will merge batch 1 with 2, then 1+2 with 3; batch 4 with 5, and then 4+5 with 6; and finally, 1+2+3 with 4+5+6.
- The same approach can be used for integer or character vectors, e.g., `list(1:3,4:6)` has the same effect as above.

Note that, while batches can be specified by name (character) or index (integer), users cannot use both in the same tree.

The merge order may occasionally be important as it determines the number of MNN pairs available at each merge step. MNN pairs results in greater stability of the batch vectors and increased likelihood of identifying shared subpopulations, which are important to the precision and accuracy of the MNN-based correction, respectively.

- In a progressive merge, the reference increases in size at each step, ensuring that more cells are available to identify MNN pairs in later merges. We suggest setting the largest, most heterogeneous batch as the first reference, which favors detection of sufficient MNN pairs between the first and other batches. Conversely, if two small batches without shared populations are supplied first, the wrong MNN pairs will be detected and the result of the merge will be incorrect.
- A merge tree is useful for merging together batches that are known to be more closely related (e.g., replicates) before attempting difficult merges involving more dissimilar batches. The idea is to increase the number of cells and thus MNN pairs prior to merging batches with few shared subpopulations. By comparison, performing the more difficult merges first is more likely to introduce errors whereby distinct subpopulations are incorrectly placed together,

which is propagated to later steps as the initial merge is used as a reference for subsequent merges.

- If `auto.merge=TRUE`, merge steps are chosen to maximize the number of MNN pairs at each step. The aim is to improve the stability of the correction by first merging more similar batches with more MNN pairs. This can be somewhat time-consuming as MNN pairs need to be iteratively recomputed for all possible batch pairings.

The order of cells in the output is *never* affected by the setting of `merge.order` or `auto.order`. It depends only on the order of objects in `...` and the order of cells within each object.

Choice of genes

All genes are used with the default setting of `subset.row=NULL`. Users can set `subset.row` to subset the inputs to highly variable genes or marker genes. This improves the quality of the PCA and identification of MNN pairs by reducing the noise from irrelevant genes. Note that users should not be too restrictive with subsetting, as high dimensionality is required to satisfy the orthogonality assumption in MNN detection.

By default, only the selected genes are used to compute rotation vectors and a low-rank corrected expression matrix. However, setting `correct.all=TRUE` will return rotation vectors that span all genes in the supplied input data. This is useful for ensuring that corrected values are returned for all input genes, e.g., in `correctExperiments`. Note that this setting will not affect the corrected low-dimension coordinates or the rotation values for the selected genes.

Using restriction

See `?batchelor-restrict` for a description of the `restrict` argument. Specifically, fastMNN will only use the restricted subset of cells in each batch to identify MNN pairs and the center of the orthogonalization. It will then extrapolate the correction to all cells in each batch.

Note that *all* cells are used to perform the PCA, regardless of whether `restrict` is set. This is generally desirable in applications where `restrict` is useful. For example, constructing the projection vectors with only control cells will not guarantee resolution of unique non-control populations in each batch.

However, this also means that the corrected values for the restricted cells will differ from the output when the inputs are directly subsetting to only contain the restricted cells. If this is not desirable, users can perform the PCA manually and apply `reducedMNN` instead.

Merge diagnostics

We can consider fastMNN's operation in terms of pairwise merge steps. Each merge step involves two mutually exclusive sets of cells, a "left" set and "right" set. Each set may consist of cells from different batches if those batches were merged in a previous step. The merge will then create a new set of cells that combines the left and right sets. Iteratively repeating this process with the newly formed sets will eventually merge all batches together.

The output metadata contains `merge.info`, a DataFrame where each row corresponds to a merge step. It contains the following fields:

- `left`, a `List` of integer or character vectors. Each vector specifies the batches in the left set at a given merge step.
- `right`, a similar `List` of integer or character vectors. Each vector specifies the batches in the right set at a given merge step.

- `pairs`, a List of DataFrames specifying which pairs of cells were identified as MNNs at each step. In each DataFrame, each row corresponds to a single MNN pair and specifies the paired cells that were in the left and right sets, respectively. Note that the indices refer to those paired cells in the *output* ordering of cells, i.e., users can identify the paired cells at each step by column-indexing the output of the `fastMNN` function.
- `batch.size`, a numeric vector specifying the relative magnitude of the batch effect at each merge, see “Orthogonalization details”.
- `skipped`, a logical vector indicating whether the correction was skipped if the magnitude of the batch effect was below `min.batch.skip`.
- `lost.var`, a numeric matrix specifying the percentage of variance lost due to orthogonalization at each merge step. This is reported separately for each batch (columns, ordered according to the input order, *not* the merge order).

Specifying the number of neighbors

The threshold to define nearest neighbors is defined by `k`, which is passed to `findMutualNN` to identify MNN pairs. The size of `k` can be roughly interpreted as the anticipated minimum size of a shared subpopulation in each batch. If a batch has fewer than `k` cells of a shared subpopulation, there is an increased risk that its counterparts in other batches will form incorrect MNN pairs.

From the perspective of the algorithm, larger values allow for more MNN pairs to be obtained, which improves the stability of the correction vectors. Larger values also increase robustness against non-orthogonality, by ignoring a certain level of biological variation when identifying pairs. This can be used to avoid the kissing problem where MNN pairs are only detected on the “surface” of the distribution. However, values of `k` should not be too large, as this would result in MNN pairs being inappropriately identified between biologically distinct populations.

In practice, increasing `k` will generally result in more aggressive merging as the algorithm is more generous in matching subpopulations across batches. We suggest starting with the default `k` and increasing it if one is confident that the same cell types are not adequately merged across batches. This is better than starting with a large `k` as incorrect merging is much harder to diagnose than insufficient merging.

An additional consideration is that the effect of any given `k` will vary with the number of cells in each batch. With more cells, a larger `k` may be preferable to achieve better merging in the presence of non-orthogonality. We can achieve this by setting `prop.k`, which allows the choice of `k` to adapt to the size of each batch at each merge step. This also handles asymmetry in batch sizes via the `k1` and `k2` arguments in `findMutualNN`.

Orthogonalization details

`fastMNN` will compute the percentage of variance that is lost from each batch during orthogonalization at each merge step. This represents the variance in each batch that is parallel to the average correction vectors (and hence removed during orthogonalization) at each merge step. Large proportions suggest that there is biological structure that is parallel to the batch effect, corresponding to violations of the assumption that the batch effect is orthogonal to the biological subspace.

If `fastMNN` is called with DataFrame inputs, each DataFrame is assumed to be the result of a previous `fastMNN` call and have a set of vectors used for orthogonalization in the merge steps of that previous call. In the current call, `fastMNN` will gather all such batch vectors across all DataFrame inputs. Each batch is then re-orthogonalized with respect to each of these vectors. This ensures that the same variation is removed from each batch prior to merging. The variance lost due to this pre-correction orthogonalization is reported in the `pre.orthog` field in the output metadata.

Orthogonalization may cause problems if there is actually no batch effect, resulting in large losses of variance. To avoid this, `fastMNN` will not perform any correction if the relative magnitude of

the batch effect is less than `min.batch.skip`. The relative magnitude is defined as the L2 norm of the average correction vector divided by the root-mean-square of the L2 norms of the per-MNN pair correction vectors. This will be large when the per-pair vectors are all pointing in the same direction, and small when the per-pair vectors point in random directions due to the absence of a consistent batch effect. If a large loss of variance is observed along with a small batch effect in a given merge step, users can set `min.batch.skip` to simply skip correction in that step.

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

Lun ATL (2018). Further MNN algorithm development. <https://MarioniLab.github.io/FurtherMNN2018/theory/description.html>

See Also

`cosineNorm` and `multiBatchPCA`, to obtain the values to be corrected.

`reducedMNN`, for a version of the function that operates in low-dimensional space.

`mnnCorrect` for the “classic” version of the MNN correction algorithm.

Examples

```
B1 <- matrix(rnorm(10000, -1), ncol=50) # Batch 1
B2 <- matrix(rnorm(10000, 1), ncol=50) # Batch 2
out <- fastMNN(B1, B2)

# Corrected values for use in clustering, etc.
str(reducedDim(out))

# Extracting corrected expression values for gene 10.
summary(assay(out)[10,])
```

findMutualNN

Find mutual nearest neighbors

Description

Find mutual nearest neighbors (MNN) across two data sets.

Usage

```
findMutualNN(
  data1,
  data2,
  k1,
  k2 = k1,
  BNPARAM = KmknnParam(),
```

```

  BPPARAM = SerialParam()
)

```

Arguments

data1	A numeric matrix containing samples (e.g., cells) in the rows and variables/dimensions in the columns.
data2	A numeric matrix like data1 for another data set with the same variables/dimensions.
k1	Integer scalar specifying the number of neighbors to search for in data1.
k2	Integer scalar specifying the number of neighbors to search for in data2.
BNPARAM	A BiocNeighborParam object specifying the neighbour search algorithm to use.
BPPARAM	A BiocParallelParam object specifying how parallelization should be performed.

Details

The concept of a MNN pair can be explained by considering cells in each of two data sets. For each cell in data set 1, the set of k2 nearest cells in data set 2 is identified, based on the Euclidean distance in expression space. For each cell in data set 2, the set of k1 nearest cells in data set 1 is similarly identified. Two cells in different batches are considered to be MNNs if each cell is in the other's set.

Value

A list containing the integer vectors `first` and `second`. Corresponding entries in `first` and `second` specify a MNN pair of cells from data1 and data2, respectively.

Author(s)

Aaron Lun

See Also

[queryKNN](#) for the underlying neighbor search code.

[fastMNN](#) and [mnnCorrect](#), which call this function to identify MNNs.

Examples

```

B1 <- matrix(rnorm(10000), ncol=50) # Batch 1
B2 <- matrix(rnorm(10000), ncol=50) # Batch 2
out <- findMutualNN(B1, B2, k1=20)
head(out$first)
head(out$second)

```

mnnCorrect	<i>Mutual nearest neighbors correction</i>
------------	--

Description

Correct for batch effects in single-cell expression data using the mutual nearest neighbors method.

Usage

```
mnnCorrect(
  ...,
  batch = NULL,
  restrict = NULL,
  k = 20,
  prop.k = NULL,
  sigma = 0.1,
  cos.norm.in = TRUE,
  cos.norm.out = TRUE,
  svd.dim = 0L,
  var.adj = TRUE,
  subset.row = NULL,
  correct.all = FALSE,
  merge.order = NULL,
  auto.merge = FALSE,
  auto.order = NULL,
  assay.type = "logcounts",
  get.spikes = FALSE,
  BSPARAM = ExactParam(),
  BNPARAM = KmknnParam(),
  BPPARAM = SerialParam()
)
```

Arguments

...	Two or more log-expression matrices where genes correspond to rows and cells correspond to columns. Each matrix should contain cells from the same batch; multiple matrices represent separate batches of cells. Each matrix should contain the same number of rows, corresponding to the same genes (in the same order). Alternatively, one or more SingleCellExperiment objects can be supplied containing a log-expression matrix in the <code>assay.type</code> assay. Note the same constraints described above for matrix inputs.
batch	A factor specifying the batch of origin for all cells when only a single object is supplied in ... This is ignored if multiple objects are present.
restrict	A list of length equal to the number of objects in ... Each entry of the list corresponds to one batch and specifies the cells to use when computing the correction.
k	An integer scalar specifying the number of nearest neighbors to consider when identifying mutual nearest neighbors.

<code>prop.k</code>	A numeric scalar in (0, 1) specifying the proportion of cells in each dataset to use for mutual nearest neighbor searching. If set, <code>k</code> for the search in each batch is redefined as $\max(k, \text{prop.k} * N)$ where <code>N</code> is the number of cells in that batch.
<code>sigma</code>	A numeric scalar specifying the bandwidth of the Gaussian smoothing kernel used to compute the correction vector for each cell.
<code>cos.norm.in</code>	A logical scalar indicating whether cosine normalization should be performed on the input data prior to calculating distances between cells.
<code>cos.norm.out</code>	A logical scalar indicating whether cosine normalization should be performed prior to computing corrected expression values.
<code>svd.dim</code>	An integer scalar specifying the number of dimensions to use for summarizing biological substructure within each batch.
<code>var.adj</code>	A logical scalar indicating whether variance adjustment should be performed on the correction vectors.
<code>subset.row</code>	A vector specifying which features to use for correction.
<code>correct.all</code>	A logical scalar specifying whether correction should be applied to all genes, even if only a subset is used for the MNN calculations.
<code>merge.order</code>	An integer vector containing the linear merge order of batches in . . . Alternatively, a list of lists representing a tree structure specifying a hierarchical merge order.
<code>auto.merge</code>	Logical scalar indicating whether to automatically identify the “best” merge order.
<code>auto.order</code>	Deprecated, use <code>merge.order</code> or <code>auto.merge</code> instead.
<code>assay.type</code>	A string or integer scalar specifying the assay containing the log-expression values, if <code>SingleCellExperiment</code> objects are present in . . .
<code>get.spikes</code>	A logical scalar indicating whether to retain rows corresponding to spike-in transcripts. Only used for <code>SingleCellExperiment</code> inputs.
<code>BSPARAM</code>	A BiocSingularParam object specifying the SVD algorithm to use.
<code>BNPARAM</code>	A BiocNeighborParam object specifying the neighbor search algorithm to use.
<code>BPPARAM</code>	A BiocParallelParam object specifying the parallelization scheme to use.

Details

This function is designed for batch correction of single-cell RNA-seq data where the batches are partially confounded with biological conditions of interest. It does so by identifying pairs of mutual nearest neighbors (MNN) in the high-dimensional log-expression space. Each MNN pair represents cells in different batches that are of the same cell type/state, assuming that batch effects are mostly orthogonal to the biological manifold. Correction vectors are calculated from the pairs of MNNS and corrected (log-)expression values are returned for use in clustering and dimensionality reduction.

For each MNN pair, a pairwise correction vector is computed based on the difference in the log-expression profiles. The correction vector for each cell is computed by applying a Gaussian smoothing kernel with bandwidth `sigma` is the pairwise vectors. This stabilizes the vectors across many MNN pairs and extends the correction to those cells that do not have MNNS. The choice of `sigma` determines the extent of smoothing - a value of 0.1 is used by default, corresponding to 10% of the radius of the space after cosine normalization.

Value

A [SingleCellExperiment](#) object containing the corrected assay. This contains corrected expression values for each gene (row) in each cell (column) in each batch. A batch field is present in the column data, specifying the batch of origin for each cell.

Cells in the output object are always ordered in the same manner as supplied in For a single input object, cells will be reported in the same order as they are arranged in that object. In cases with multiple input objects, the cell identities are simply concatenated from successive objects, i.e., all cells from the first object (in their provided order), then all cells from the second object, and so on.

The metadata of the `SingleCellExperiment` contains `merge.info`, a `DataFrame` where each row corresponds to a merge step. See “Merge diagnostics” for more information.

Choosing the gene set

All genes are used with the default setting of `subset.row=NULL`. Users can set `subset.row` to subset the inputs to highly variable genes or marker genes. This may provide more meaningful identification of MNN pairs by reducing the noise from irrelevant genes. Note that users should not be too restrictive with subsetting, as high dimensionality is required to satisfy the orthogonality assumption in MNN detection.

If `subset.row` is specified and `correct.all=TRUE`, corrected values are returned for *all* genes. This is possible as `subset.row` is only used to identify the MNN pairs and other cell-based distance calculations. Correction vectors between MNN pairs can then be computed in for all genes in the supplied matrices. Note that setting `correct.all=TRUE` will not alter the corrected expression values for the subsetted genes.

Expected type of input data

The input expression values should generally be log-transformed, e.g., log-counts, see [normalize](#) for details. They should also be normalized within each data set to remove cell-specific biases in capture efficiency and sequencing depth. By default, a further cosine normalization step is performed on the supplied expression data to eliminate gross scaling differences between data sets.

- When `cos.norm.in=TRUE`, cosine normalization is performed on the matrix of expression values used to compute distances between cells. This can be turned off when there are no scaling differences between data sets.
- When `cos.norm.out=TRUE`, cosine normalization is performed on the matrix of values used to calculate correction vectors (and on which those vectors are applied). This can be turned off to obtain corrected values on the log-scale, similar to the input data.

The cosine normalization is achieved using the [cosineNorm](#) function.

Further options

The function depends on a shared biological manifold, i.e., one or more cell types/states being present in multiple batches. If this is not true, MNNS may be incorrectly identified, resulting in over-correction and removal of interesting biology. Some protection can be provided by removing components of the correction vectors that are parallel to the biological subspaces in each batch. The biological subspace in each batch is identified with a SVD on the expression matrix to obtain `svd.dim` dimensions. (By default, this option is turned off by setting `svd.dim=0`.)

If `var.adj=TRUE`, the function will adjust the correction vector to equalize the variances of the two data sets along the batch effect vector. In particular, it avoids “kissing” effects whereby MNN pairs are identified between the surfaces of point clouds from different batches. Naive correction

would then bring only the surfaces into contact, rather than fully merging the clouds together. The adjustment ensures that the cells from the two batches are properly intermingled after correction. This is done by identifying each cell's position on the correction vector, identifying corresponding quantiles between batches, and scaling the correction vector to ensure that the quantiles are matched after correction.

See `?batchelor-restrict` for a description of the `restrict` argument. Specifically, `mnnCorrect` will only use the restricted subset of cells in each batch to identify MNN pairs (and to perform variance adjustment, if `var.adj=TRUE`), and then apply the correction to all cells in each batch.

Merge diagnostics

Each merge step combines two mutually exclusive sets of cells, a “left” set and “right” set. The metadata thus contains the following fields:

- `left`, a `List` of integer or character vectors. Each vector specifies the batches in the left set at a given merge step.
- `right`, a similar `List` of integer or character vectors. Each vector specifies the batches in the right set at a given merge step.
- `pairs`, a `List` of `DataFrames` specifying which pairs of cells were identified as MNNs at each step. In each `DataFrame`, each row corresponds to a single MNN pair and specifies the paired cells that were in the left and right sets, respectively. Note that the indices refer to those paired cells in the *output* ordering of cells, i.e., users can identify the paired cells at each step by column-indexing the output of the `mnnCorrect` function.

Specifying the number of neighbors

The threshold to define nearest neighbors is defined by `k`, which is passed to `findMutualNN` to identify MNN pairs. The size of `k` can be roughly interpreted as the anticipated minimum size of a shared subpopulation in each batch. If a batch has fewer than `k` cells of a shared subpopulation, there is an increased risk that its counterparts in other batches will form incorrect MNN pairs.

From the perspective of the algorithm, larger values allow for more MNN pairs to be obtained, which improves the stability of the correction vectors. Larger values also increase robustness against non-orthogonality, by ignoring a certain level of biological variation when identifying pairs. This can be used to avoid the kissing problem where MNN pairs are only detected on the “surface” of the distribution. However, values of `k` should not be too large, as this would result in MNN pairs being inappropriately identified between biologically distinct populations.

In practice, increasing `k` will generally result in more aggressive merging as the algorithm is more generous in matching subpopulations across batches. We suggest starting with the default `k` and increasing it if one is confident that the same cell types are not adequately merged across batches. This is better than starting with a large `k` as incorrect merging is much harder to diagnose than insufficient merging.

An additional consideration is that the effect of any given `k` will vary with the number of cells in each batch. With more cells, a larger `k` may be preferable to achieve better merging in the presence of non-orthogonality. We can achieve this by setting `prop.k`, which allows the choice of `k` to adapt to the size of each batch at each merge step. This also handles asymmetry in batch sizes via the `k1` and `k2` arguments in `findMutualNN`.

Controlling the merge order

By default, batches are merged in the user-supplied order, i.e., the first batch is merged with the second batch, the third batch is merged with the combined first-second batch, the fourth batch

is merged with the combined first-second-third batch and so on. We refer to this approach as a progressive merge.

If `merge.order` is an integer vector, it is treated as an ordering permutation with which to perform a progressive merge. For example, if `merge.order=c(4, 1, 3, 2)`, batches 4 and 1 in . . . are merged first; batch 3 is merged with the combined 4+1 batch; and then batch 2 is merged with the combined 4+1+3 batch. This is often more convenient than changing the order manually in . . . , which would alter the order of batches in the output corrected matrix.

If `merge.order` is a character vector, it is treated as an ordering permutation for named batches.

If `merge.order` is a nested list, it is treated as a tree that specifies a hierarchical merge. Each element of the list should either be a string or integer scalar, corresponding to a leaf node that specifies a batch; or another list, corresponding to an internal node that should contain at least two children; or an integer or character vector of length 2 or more, again corresponding to an internal node.

- For example, `list(list(1,2),list(3,4))` indicates that batch 1 should be merged with batch 2; batch 3 should be merged with batch 4; and that, finally, the combined batches 1+2 and 3+4 should be merged.
- More than two children per node are supported and will result in a progressive merge within that node. For example, `list(list(1,2,3),list(4,5,6))` will merge batch 1 with 2, then 1+2 with 3; batch 4 with 5, and then 4+5 with 6; and finally, 1+2+3 with 4+5+6.
- The same approach can be used for integer or character vectors, e.g., `list(1:3,4:6)` has the same effect as above.

Note that, while batches can be specified by name (character) or index (integer), users cannot use both in the same tree.

The merge order may occasionally be important as it determines the number of MNN pairs available at each merge step. MNN pairs results in greater stability of the batch vectors and increased likelihood of identifying shared subpopulations, which are important to the precision and accuracy of the MNN-based correction, respectively.

- In a progressive merge, the reference increases in size at each step, ensuring that more cells are available to identify MNN pairs in later merges. We suggest setting the largest, most heterogeneous batch as the first reference, which favors detection of sufficient MNN pairs between the first and other batches. Conversely, if two small batches without shared populations are supplied first, the wrong MNN pairs will be detected and the result of the merge will be incorrect.
- A merge tree is useful for merging together batches that are known to be more closely related (e.g., replicates) before attempting difficult merges involving more dissimilar batches. The idea is to increase the number of cells and thus MNN pairs prior to merging batches with few shared subpopulations. By comparison, performing the more difficult merges first is more likely to introduce errors whereby distinct subpopulations are incorrectly placed together, which is propagated to later steps as the initial merge is used as a reference for subsequent merges.
- If `auto.merge=TRUE`, merge steps are chosen to maximize the number of MNN pairs at each step. The aim is to improve the stability of the correction by first merging more similar batches with more MNN pairs. This can be somewhat time-consuming as MNN pairs need to be iteratively recomputed for all possible batch pairings.

The order of cells in the output is *never* affected by the setting of `merge.order` or `auto.order`. It depends only on the order of objects in . . . and the order of cells within each object.

Author(s)

Laleh Haghverdi, with modifications by 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

See Also

[fastMNN](#) for a faster equivalent.

Examples

```
B1 <- matrix(rnorm(10000), ncol=50) # Batch 1
B2 <- matrix(rnorm(10000), ncol=50) # Batch 2
out <- mnnCorrect(B1, B2) # corrected values
```

multiBatchNorm

Per-batch scaling normalization

Description

Perform scaling normalization within each batch to provide comparable results to the lowest-coverage batch.

Usage

```
multiBatchNorm(
  ...,
  batch = NULL,
  assay.type = "counts",
  norm.args = list(),
  min.mean = 1,
  subset.row = NULL,
  normalize.all = FALSE,
  preserve.single = TRUE
)
```

Arguments

...	One or more SingleCellExperiment objects containing counts and size factors. Each object should contain the same number of rows, corresponding to the same genes in the same order. If multiple objects are supplied, each object is assumed to contain all and only cells from a single batch. If a single object is supplied, batch should also be specified.
batch	A factor specifying the batch of origin for all cells when only a single object is supplied in ... This is ignored if multiple objects are present.
assay.type	A string specifying which assay values contains the counts.

<code>norm.args</code>	A named list of further arguments to pass to <code>normalize</code> .
<code>min.mean</code>	A numeric scalar specifying the minimum (library size-adjusted) average count of genes to be used for normalization.
<code>subset.row</code>	A vector specifying which features to use for normalization.
<code>normalize.all</code>	A logical scalar indicating whether normalized values should be returned for all genes.
<code>preserve.single</code>	A logical scalar indicating whether to combine the results into a single matrix if only one object was supplied in <code>...</code>

Details

When performing integrative analyses of multiple batches, it is often the case that different batches have large differences in coverage. This function removes systematic differences in coverage across batches to simplify downstream comparisons. It does so by rescaling the size factors using median-based normalization on the ratio of the average counts between batches. This is roughly equivalent to the between-cluster normalization described by Lun et al. (2016).

This function will adjust the size factors so that counts in high-coverage batches are scaled *downwards* to match the coverage of the most shallow batch. The `logNormCounts` function will then add the same pseudo-count to all batches before log-transformation. By scaling downwards, we favour stronger squeezing of log-fold changes from the pseudo-count, mitigating any technical differences in variance between batches. Of course, genuine biological differences will also be shrunk, but this is less of an issue for upregulated genes with large counts.

For comparison, imagine if we ran `logNormCounts` separately in each batch prior to correction. In most cases, size factors will be computed within each batch; batch-specific application in `logNormCounts` will not account for scaling differences between batches. In contrast, `multiBatchNorm` will rescale the size factors so that they are comparable across batches. This removes at least one difference between batches to facilitate easier correction.

Only genes with library size-adjusted average counts greater than `min.mean` will be used for computing the rescaling factors. This improves precision and avoids problems with discreteness. By default, we use `min.mean=1`, which is usually satisfactory but may need to be lowered for very sparse datasets.

Users can also set `subset.row` to restrict the set of genes used for computing the rescaling factors. By default, normalized values will only be returned for genes specified in the subset. Setting `normalize.all=TRUE` will return normalized values for all genes.

Value

A list of `SingleCellExperiment` objects with normalized log-expression values in the "logcounts" assay (depending on values in `norm.args`). Each object contains cells from a single batch.

If `preserve.single=TRUE` and `...` contains only one `SingleCellExperiment`, that object is returned with an additional "logcounts" assay containing normalized log-expression values. The order of cells is not changed.

Note about spike-ins

Rescaling is only performed on endogenous genes in each `SingleCellExperiment` object. If any spike-in transcripts are present in the `altExps`, their abundances will not be rescaled here, and are no longer directly comparable to the rescaled abundances of the genes. This is usually not a major problem as spike-ins are rarely used during the batch correction itself - however, users should not attempt to perform variance modelling with the spike-ins on the output of this function.

Author(s)

Aaron Lun

References

Lun ATL (2018). Further MNN algorithm development. <https://MarioniLab.github.io/FurtherMNN2018/theory/description.html>

See Also

[mnnCorrect](#) and [fastMNN](#), for methods that can benefit from rescaling.
[normalize](#) for the calculation of log-transformed normalized expression values.

Examples

```
d1 <- matrix(rnbinom(50000, mu=10, size=1), ncol=100)
sce1 <- SingleCellExperiment(list(counts=d1))
sizeFactors(sce1) <- runif(ncol(d1))

d2 <- matrix(rnbinom(20000, mu=50, size=1), ncol=40)
sce2 <- SingleCellExperiment(list(counts=d2))
sizeFactors(sce2) <- runif(ncol(d2))

out <- multiBatchNorm(sce1, sce2)
summary(sizeFactors(out[[1]]))
summary(sizeFactors(out[[2]]))
```

multiBatchPCA

Multi-batch PCA

Description

Perform a principal components analysis across multiple gene expression matrices to project all cells to a common low-dimensional space.

Usage

```
multiBatchPCA(
  ...,
  batch = NULL,
  d = 50,
  subset.row = NULL,
  weights = NULL,
  get.all.genes = FALSE,
  rotate.all = FALSE,
  get.variance = FALSE,
  preserve.single = FALSE,
  assay.type = "logcounts",
  get.spikes = FALSE,
  BSPARAM = ExactParam(),
  BPPARAM = SerialParam()
)
```

Arguments

<code>...</code>	Two or more matrices containing expression values (usually log-normalized). Each matrix is assumed to represent one batch and should contain the same number of rows, corresponding to the same genes (in the same order). Alternatively, two or more SingleCellExperiment objects containing these matrices. Note the same restrictions described above for gene expression matrix inputs. Alternatively, one matrix or SingleCellExperiment can be supplied containing cells from all batches. This requires <code>batch</code> to also be specified.
<code>batch</code>	A factor specifying the batch identity of each cell in the input data. Ignored if <code>...</code> contains more than one argument.
<code>d</code>	An integer scalar specifying the number of dimensions to keep from the initial multi-sample PCA.
<code>subset.row</code>	A vector specifying which features to use for correction.
<code>weights</code>	Numeric vector of length equal to the number of entries in <code>...</code> , specifying the scaling of the weight of each batch. This defaults to 1 for all batches.
<code>get.all.genes</code>	A logical scalar indicating whether the reported rotation vectors should include genes that are excluded by a non-NULL value of <code>subset.row</code> .
<code>rotate.all</code>	A deprecated synonym for <code>get.all.genes</code> .
<code>get.variance</code>	A logical scalar indicating whether to return the (weighted) variance explained by each PC.
<code>preserve.single</code>	A logical scalar indicating whether to combine the results into a single matrix if only one object was supplied in <code>...</code> .
<code>assay.type</code>	A string or integer scalar specifying the assay containing the expression values, if SingleCellExperiment objects are present in <code>...</code> .
<code>get.spikes</code>	Deprecated, a logical scalar indicating whether to retain rows corresponding to spike-in transcripts. Only used for SingleCellExperiment inputs.
<code>BSPARAM</code>	A BiocSingularParam object specifying the algorithm to use for PCA, see runSVD for details.
<code>BPPARAM</code>	A BiocParallelParam object specifying whether the SVD should be parallelized.

Details

This function is roughly equivalent to `cbinding` all matrices in `...` and performing PCA on the merged matrix. The main difference is that each sample is forced to contribute equally to the identification of the rotation vectors. Specifically, the mean vector used for centering is defined as the grand mean of the mean vectors within each batch. Each batch's contribution to the gene-gene covariance matrix is also divided by the number of cells in that batch.

Our approach is to effectively weight the cells in each batch to mimic the situation where all batches have the same number of cells. This ensures that the low-dimensional space can distinguish subpopulations in smaller batches. Otherwise, batches with a large number of cells would dominate the PCA, i.e., the definition of the mean vector and covariance matrix. This may reduce resolution of unique subpopulations in smaller batches that differ in a different dimension to the subspace of the larger batches.

When `weights` is set, this will scale the weight of each batch by the specified value. For example, each batch may represent one replicate, with multiple replicates per study. In such cases, it may be more appropriate to ensure that each *study* has equal weight. This is done by assigning a value

of weights to each replicate that is inversely proportional to the number of replicates in the same study - see Examples.

If `...` contains `SingleCellExperiment` objects, any spike-in transcripts should be the same across all batches. These will be removed prior to PCA unless `get.spikes=TRUE`. If `subset.row` is specified and `get.spikes=FALSE`, only the non-spike-in specified features will be used.

Setting `get.all.genes=TRUE` will report rotation vectors that span all genes, even when only a subset of genes are used for the PCA. This is done by projecting all non-used genes into the low-dimensional “cell space” defined by the first `d` components.

If `BSPARAM` is defined with `deferred=TRUE`, the per-gene centering and per-cell scaling will be manually deferred during matrix multiplication. This can greatly improve speeds when the input matrices are sparse, as deferred operations avoids loss of sparsity (at the cost of numerical precision).

Value

A [List](#) of numeric matrices is returned where each matrix corresponds to a batch and contains the first `d` PCs (columns) for all cells in the batch (rows).

If `preserve.single=TRUE` and `...` contains a single object, the List will only contain a single matrix. This contains the first `d` PCs (columns) for all cells in the same order as supplied in the single input object.

The metadata contains `rotation`, a matrix of rotation vectors, which can be used to construct a low-rank approximation of the input matrices. This has number of rows equal to the number of genes after any subsetting, except if `rotate.all=TRUE`, where the number of rows is equal to the genes before subsetting.

If `get.variance=TRUE`, the metadata will also contain `var.explained`, the weighted variance explained by each PC; and `var.total`, the total variance after weighting.

Author(s)

Aaron Lun

See Also

[runSVD](#)

Examples

```
d1 <- matrix(rnorm(5000), ncol=100)
d1[1:10,1:10] <- d1[1:10,1:10] + 2 # unique population in d1
d2 <- matrix(rnorm(2000), ncol=40)
d2[11:20,1:10] <- d2[11:20,1:10] + 2 # unique population in d2

out <- multiBatchPCA(d1, d2)

# Examining results.
xlim <- range(c(out[[1]][,1], out[[2]][,1]))
ylim <- range(c(out[[1]][,2], out[[2]][,2]))
plot(out[[1]][,1], out[[1]][,2], col="red", xlim=xlim, ylim=ylim)
points(out[[2]][,1], out[[2]][,2], col="blue")

# Using the weighting scheme, assuming that 'd2' and 'd3'
# are replicates and should contribute the same combined
# weight as 'd1'.
```

```
d3 <- d2 + 5
out <- multiBatchPCA(d1, d2, d3, weights=c(1, 0.5, 0.5))
```

noCorrect	<i>No correction</i>
-----------	----------------------

Description

Provides a no-correction method that has the same interface as the correction functions. This allows users to easily swap function calls to examine the effect of correction.

Usage

```
noCorrect(..., batch = NULL, subset.row = NULL, assay.type = "logcounts")
```

Arguments

...	One or more log-expression matrices where genes correspond to rows and cells correspond to columns, if <code>pc.input=FALSE</code> . Each matrix should contain the same number of rows, corresponding to the same genes in the same order. Alternatively, one or more SingleCellExperiment objects can be supplied containing a log-expression matrix in the <code>assay.type</code> assay. Note the same restrictions described above for gene expression matrix inputs. If multiple objects are supplied, each object is assumed to contain all and only cells from a single batch. Objects of different types can be mixed together. If a single object is supplied, <code>batch</code> should also be specified.
<code>batch</code>	A factor specifying the batch of origin for all cells when only a single object is supplied in ... This is ignored if multiple objects are present.
<code>subset.row</code>	A vector specifying which features to retain.
<code>assay.type</code>	A string or integer scalar specifying the assay containing the log-expression values. Only used for SingleCellExperiment inputs.

Details

This function is effectively equivalent to `cbinding` the matrices together without any correction. The aim is to provide a consistent interface that allows users to simply combine batches without additional operations. This is often desirable as a negative control to see if the transformation is actually beneficial. It also allows for convenient downstream analyses that are based on the uncorrected data, e.g., differential expression.

Value

A [SingleCellExperiment](#) is returned where each row is a gene and each column is a cell. This contains:

- A merged matrix in the `assays` slot, containing the merged expression values from all elements of ...
- A batch column in the `colData` slot, containing the batch of origin for each row (i.e., cell) in corrected.

Author(s)

Aaron Lun

Examples

```

B1 <- matrix(rnorm(10000), ncol=50) # Batch 1
B2 <- matrix(rnorm(10000), ncol=50) # Batch 2
out <- noCorrect(B1, B2)

# Same as combining the expression values.
stopifnot(all(assay(out)==cbind(B1, B2)))

# Specifies which cell came from which batch:
str(out$batch)

```

reducedMNN

*MNN correction in reduced dimensions***Description**

MNN correction in reduced dimensions

Usage

```

reducedMNN(
  ...,
  batch = NULL,
  k = 20,
  prop.k = NULL,
  restrict = NULL,
  ndist = 3,
  merge.order = NULL,
  auto.merge = FALSE,
  auto.order = NULL,
  min.batch.skip = 0,
  BNPARAM = KmknnParam(),
  BPPARAM = SerialParam()
)

```

Arguments

... One or more matrices of low-dimensional representations where rows are cells and columns are dimensions. Each object should contain the same number of columns, corresponding to the same dimensions. These should have been generated by a single call to [multiBatchPCA](#).

If multiple objects are supplied, each object is assumed to contain all and only cells from a single batch. If a single object is supplied, batch should also be specified.

batch A factor specifying the batch of origin for all cells when only a single object is supplied in ... This is ignored if multiple objects are present.

<code>k</code>	An integer scalar specifying the number of nearest neighbors to consider when identifying MNNs.
<code>prop.k</code>	A numeric scalar in (0, 1) specifying the proportion of cells in each dataset to use for mutual nearest neighbor searching. If set, <code>k</code> for the search in each batch is redefined as $\max(k, \text{prop.k} \cdot N)$ where <code>N</code> is the number of cells in that batch.
<code>restrict</code>	A list of length equal to the number of objects in <code>...</code> . Each entry of the list corresponds to one batch and specifies the cells to use when computing the correction.
<code>ndist</code>	A numeric scalar specifying the threshold beyond which neighbours are to be ignored when computing correction vectors. Each threshold is defined as a multiple of the number of median distances.
<code>merge.order</code>	An integer vector containing the linear merge order of batches in <code>...</code> . Alternatively, a list of lists representing a tree structure specifying a hierarchical merge order.
<code>auto.merge</code>	Logical scalar indicating whether to automatically identify the “best” merge order.
<code>auto.order</code>	Deprecated, use <code>merge.order</code> or <code>auto.merge</code> instead.
<code>min.batch.skip</code>	Numeric scalar specifying the minimum relative magnitude of the batch effect, below which no correction will be performed at a given merge step.
<code>BNPARAM</code>	A BiocNeighborParam object specifying the nearest neighbor algorithm.
<code>BPPARAM</code>	A BiocParallelParam object specifying whether the PCA and nearest-neighbor searches should be parallelized.

Details

`reducedMNN` performs the same operations as [fastMNN](#) but assumes that the PCA has already been performed. This is useful as the PCA (via [multiBatchPCA](#)) is often the most time-consuming step. By performing the PCA once, `reducedMNN` allows the MNN correction to be quickly repeated with different parameters.

`reducedMNN` operates on the same principles as [fastMNN](#), so users are referred to the documentation for the latter for more details on the effect of each of the arguments. Obviously, any arguments pertaining to gene-based steps in [fastMNN](#) are not relevant here.

Note that [multiBatchPCA](#) will not perform cosine-normalization, so it is the responsibility of the user to cosine-normalize each batch beforehand with [cosineNorm](#) to recapitulate results of [fastMNN](#) with `cos.norm=TRUE`. In addition, [multiBatchPCA](#) must be run on all samples at once, to ensure that all cells are projected to the same low-dimensional space.

Value

A [DataFrame](#) is returned where each row corresponds to a cell, containing:

- `corrected`, the matrix of corrected low-dimensional coordinates for each cell.
- `batch`, the Rle specifying the batch of origin for each row.

Cells in the output object are always ordered in the same manner as supplied in `...`. The metadata on this object is the same as that in the output of [fastMNN](#).

Author(s)

Aaron Lun

See Also

[multiBatchPCA](#), to obtain the values to be corrected.

[fastMNN](#), for the version that operates on gene-expression values.

Examples

```
B1 <- matrix(rnorm(10000), nrow=50) # Batch 1
B2 <- matrix(rnorm(10000), nrow=50) # Batch 2

# Equivalent to fastMNN().
cB1 <- cosineNorm(B1)
cB2 <- cosineNorm(B2)
pcs <- multiBatchPCA(cB1, cB2)
out2 <- reducedMNN(pcs[[1]], pcs[[2]])
```

regressBatches

Regress out batch effects

Description

Fit a linear model to regress out uninteresting factors of variation.

Usage

```
regressBatches(
  ...,
  batch = NULL,
  restrict = NULL,
  subset.row = NULL,
  assay.type = "logcounts"
)
```

Arguments

...	Two or more log-expression matrices where genes correspond to rows and cells correspond to columns. Each matrix should contain the same number of rows, corresponding to the same genes (in the same order). Alternatively, one or more SingleCellExperiment objects can be supplied containing a count matrix in the <code>assay.type</code> assay. Note the same restrictions described above for gene expression matrix inputs. If multiple objects are supplied, each object is assumed to contain all and only cells from a single batch. Objects of different types can be mixed together. If a single object is supplied, <code>batch</code> should also be specified.
<code>batch</code>	A factor specifying the batch of origin for all cells when only a single object is supplied in ... This is ignored if multiple objects are present.
<code>restrict</code>	A list of length equal to the number of objects in ... Each entry of the list corresponds to one batch and specifies the cells to use when computing the correction.
<code>subset.row</code>	A vector specifying which features to use for correction.
<code>assay.type</code>	A string or integer scalar specifying the assay containing the log-expression values, if SingleCellExperiment objects are present in ...

Details

This function fits a linear model to the log-expression values for each gene and returns the residuals. The model is parameterized as a one-way layout with the batch of origin, so the residuals represent the expression values after correcting for the batch effect.

The novelty of this function is that it returns a [ResidualMatrix](#) in as the "corrected" assay. This avoids explicitly computing the residuals, which would result in a loss of sparsity or similar problems. Rather, the residuals are either computed as needed or are never explicitly computed as all (e.g., during matrix multiplication).

All genes are used with the default setting of `subset.row=NULL`. Users can set `subset.row` to subset the inputs, though this is purely for convenience as each gene is processed independently of other genes.

See `?batchelor-restrict` for a description of the `restrict` argument. Specifically, this function will compute the model coefficients using only the specified subset of cells. The regression will then be applied to all cells in each batch.

Value

A [SingleCellExperiment](#) object containing the corrected assay. This contains corrected log-expression values for each gene (row) in each cell (column) in each batch. A batch field is present in the column data, specifying the batch of origin for each cell.

Cells in the output object are always ordered in the same manner as supplied in `...`. For a single input object, cells will be reported in the same order as they are arranged in that object. In cases with multiple input objects, the cell identities are simply concatenated from successive objects, i.e., all cells from the first object (in their provided order), then all cells from the second object, and so on.

Author(s)

Aaron Lun

See Also

[rescaleBatches](#), for another approach to regressing out the batch effect.

Examples

```
means <- 2^rgamma(1000, 2, 1)
A1 <- matrix(rpois(10000, lambda=means), ncol=50) # Batch 1
A2 <- matrix(rpois(10000, lambda=means*runif(1000, 0, 2)), ncol=50) # Batch 2

B1 <- log2(A1 + 1)
B2 <- log2(A2 + 1)
out <- regressBatches(B1, B2)
```

rescaleBatches *Scale counts across batches*

Description

Scale counts so that the average count within each batch is the same for each gene.

Usage

```
rescaleBatches(
  ...,
  batch = NULL,
  restrict = NULL,
  log.base = 2,
  pseudo.count = 1,
  subset.row = NULL,
  assay.type = "logcounts",
  get.spikes = FALSE
)
```

Arguments

...	Two or more log-expression matrices where genes correspond to rows and cells correspond to columns. Each matrix should contain the same number of rows, corresponding to the same genes (in the same order). Alternatively, one or more SingleCellExperiment objects can be supplied containing a count matrix in the <code>assay.type</code> assay. Note the same restrictions described above for gene expression matrix inputs. If multiple objects are supplied, each object is assumed to contain all and only cells from a single batch. Objects of different types can be mixed together. If a single object is supplied, <code>batch</code> should also be specified.
<code>batch</code>	A factor specifying the batch of origin for all cells when only a single object is supplied in ... This is ignored if multiple objects are present.
<code>restrict</code>	A list of length equal to the number of objects in ... Each entry of the list corresponds to one batch and specifies the cells to use when computing the correction.
<code>log.base</code>	A numeric scalar specifying the base of the log-transformation.
<code>pseudo.count</code>	A numeric scalar specifying the pseudo-count used for the log-transformation.
<code>subset.row</code>	A vector specifying which features to use for correction.
<code>assay.type</code>	A string or integer scalar specifying the assay containing the log-expression values, if SingleCellExperiment objects are present in ...
<code>get.spikes</code>	Deprecated, a logical scalar indicating whether to retain rows corresponding to spike-in transcripts. Only used for SingleCellExperiment inputs.

Details

This function assumes that the log-expression values were computed by a log-transformation of normalized count data, plus a pseudo-count. It reverses the log-transformation and scales the underlying counts in each batch so that the average (normalized) count is equal across batches. The

assumption here is that each batch contains the same population composition. Thus, any scaling difference between batches is technical and must be removed.

This function is equivalent to centering in log-expression space, the simplest application of linear regression methods for batch correction. However, by scaling the raw counts, it avoids loss of sparsity that would otherwise result from centering. It also mitigates issues with artificial differences in variance due to log-transformation.

The output values are always re-log-transformed with the same `log.base` and `pseudo.count`. These can be used directly in place of the input values for downstream operations.

All genes are used with the default setting of `subset.row=NULL`. Users can set `subset.row` to subset the inputs, though this is purely for convenience as each gene is processed independently of other genes.

See [?"batchelor-restrict"](#) for a description of the `restrict` argument. Specifically, the function will compute the scaling differences using only the specified subset of cells, and then apply the re-scaling to all cells in each batch.

Value

A `SingleCellExperiment` object containing the corrected assay. This contains corrected log-expression values for each gene (row) in each cell (column) in each batch. A `batch` field is present in the column data, specifying the batch of origin for each cell.

Cells in the output object are always ordered in the same manner as supplied in `...`. For a single input object, cells will be reported in the same order as they are arranged in that object. In cases with multiple input objects, the cell identities are simply concatenated from successive objects, i.e., all cells from the first object (in their provided order), then all cells from the second object, and so on.

Author(s)

Aaron Lun

Examples

```
means <- 2^rgamma(1000, 2, 1)
A1 <- matrix(rpois(10000, lambda=means), ncol=50) # Batch 1
A2 <- matrix(rpois(10000, lambda=means*runif(1000, 0, 2)), ncol=50) # Batch 2

B1 <- log2(A1 + 1)
B2 <- log2(A2 + 1)
out <- rescaleBatches(B1, B2)
```

Index

- altExps, [27](#)

- batchCorrect, [2, 6, 7, 10](#)
- batchCorrect, ClassicMnnParam-method
(batchCorrect), [2](#)
- batchCorrect, FastMnnParam-method
(batchCorrect), [2](#)
- batchCorrect, NoCorrectParam-method
(batchCorrect), [2](#)
- batchCorrect, RegressParam-method
(batchCorrect), [2](#)
- batchCorrect, RescaleParam-method
(batchCorrect), [2](#)
- batchelor-restrict, [5](#)
- BatchelorParam, [4, 9](#)
- BatchelorParam-class, [6](#)
- BiocNeighborParam, [14, 20, 22, 33](#)
- BiocParallelParam, [11, 14, 20, 22, 29, 33](#)
- BiocSingularParam, [14, 22, 29](#)

- checkBatchConsistency, [7](#)
- checkIfSCE (checkBatchConsistency), [7](#)
- checkRestrictions
(checkBatchConsistency), [7](#)
- checkSpikeConsistency
(checkBatchConsistency), [7](#)
- ClassicMnnParam, [4, 9](#)
- ClassicMnnParam (BatchelorParam-class),
[6](#)
- ClassicMnnParam-class
(BatchelorParam-class), [6](#)
- correctExperiments, [8, 17](#)
- cosineNorm, [10, 15, 19, 23, 33](#)

- DataFrame, [16, 33](#)
- DelayedArray, [11](#)
- DelayedMatrix, [11](#)
- divideIntoBatches, [8, 12](#)

- fastMNN, [4–6, 9, 11, 13, 20, 26, 28, 33, 34](#)
- FastMnnParam, [4, 9](#)
- FastMnnParam (BatchelorParam-class), [6](#)
- FastMnnParam-class
(BatchelorParam-class), [6](#)

- findMutualNN, [18, 19, 24](#)

- List, [17, 24, 30](#)
- logNormCounts, [27](#)

- mnnCorrect, [4, 6, 9, 11, 15, 19, 20, 21, 28](#)
- multiBatchNorm, [15, 26](#)
- multiBatchPCA, [14, 15, 19, 28, 32–34](#)

- noCorrect, [10, 31](#)
- NoCorrectParam (BatchelorParam-class), [6](#)
- NoCorrectParam-class
(BatchelorParam-class), [6](#)
- normalize, [23, 27, 28](#)

- queryKNN, [20](#)

- reducedMNN, [14, 17, 19, 32](#)
- regressBatches, [4, 6, 9, 34](#)
- RegressParam, [4, 9](#)
- RegressParam (BatchelorParam-class), [6](#)
- RegressParam-class
(BatchelorParam-class), [6](#)
- rescaleBatches, [4, 6, 9, 35, 36](#)
- RescaleParam, [4, 9](#)
- RescaleParam (BatchelorParam-class), [6](#)
- RescaleParam-class
(BatchelorParam-class), [6](#)
- ResidualMatrix, [35](#)
- rowRanges, [9](#)
- runSVD, [29, 30](#)

- SimpleList, [6](#)
- SingleCellExperiment, [3, 7–9, 14, 15, 21, 23, 26, 29, 31, 34–37](#)