

Package ‘tidySummarizedExperiment’

November 14, 2025

Type Package

Title Brings SummarizedExperiment to the Tidyverse

Version 1.20.1

Author Stefano Mangiola [aut, cre] <mangiolastefano@gmail.com>

Maintainer Stefano Mangiola <mangiolastefano@gmail.com>

Description The tidySummarizedExperiment package provides a set of tools for creating and manipulating tidy data representations of SummarizedExperiment objects. SummarizedExperiment is a widely used data structure in bioinformatics for storing high-throughput genomic data, such as gene expression or DNA sequencing data. The tidySummarizedExperiment package introduces a tidy framework for working with SummarizedExperiment objects. It allows users to convert their data into a tidy format, where each observation is a row and each variable is a column. This tidy representation simplifies data manipulation, integration with other tidyverse packages, and enables seamless integration with the broader ecosystem of tidy tools for data analysis.

License GPL-3

Depends R (>= 4.3.0), SummarizedExperiment, tsservice (>= 0.5.0)

Imports dplyr, tibble (>= 3.0.4), magrittr, tidyr, ggplot2, rlang, purrr, lifecycle, methods, utils, S4Vectors, tidyselect, ellipsis, vctrs, pillar, stringr, cli, fansi, stats, pkgconfig, plyxp, generics

Suggests BiocStyle, testthat, knitr, markdown, rmarkdown, plotly, rbibutils, prettydoc, airway

VignetteBuilder knitr

RdMacros lifecycle

Biarch true

biocViews AssayDomain, Infrastructure, RNASeq, DifferentialExpression, GeneExpression, Normalization, Clustering, QualityControl, Sequencing, Transcription, Transcriptomics

Encoding UTF-8

LazyData true

RoxygenNote 7.3.3

Roxygen list(markdown = TRUE)

LazyDataCompression xz

URL <https://github.com/stemangiola/tidySummarizedExperiment>

BugReports <https://github.com/stemangiola/tidySummarizedExperiment/issues>

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

git_branch RELEASE_3_22

git_last_commit 6626935

git_last_commit_date 2025-11-09

Repository Bioconductor 3.22

Date/Publication 2025-11-13

Contents

tidySummarizedExperiment-package	3
append_samples	4
as_tibble	4
bind_rows	6
count	8
distinct	9
extract	10
filter	11
full_join	13
ggplot	16
group_by	17
group_split	19
inner_join	20
left_join	23
mutate	26
nest	28
pasilla	29
pivot_longer	30
pivot_wider	32
plot_ly	35
pull	38
rename	39
right_join	40
rowwise	43
sample_n	44
se	45
select	46
separate	50
slice	52
summarise	53
tidy	55
unite	56
unnest	57
%>%	59

Index

`tidySummarizedExperiment-package`*tidySummarizedExperiment: Bringing SummarizedExperiment to the Tidyverse*

Description

The tidySummarizedExperiment package provides a bridge between Bioconductor SummarizedExperiment objects and the tidyverse. It creates an invisible layer that enables viewing the Bioconductor SummarizedExperiment object as a tidyverse tibble, and provides SummarizedExperiment-compatible dplyr, tidyr, ggplot2 and plotly functions. This allows users to get the best of both Bioconductor and tidyverse worlds.

Details

The main functions and methods are:

- `tidy` - Convert SummarizedExperiment to tidy format
- `as_tibble` - Convert to tibble representation
- `filter`, `select`, `mutate`, `arrange` - dplyr verbs
- `pivot_longer`, `pivot_wider`, `nest`, `unnest` - tidyr functions
- `ggplot` - ggplot2 visualization
- `left_join`, `inner_join`, `right_join`, `full_join` - dplyr joins

For detailed information on usage, see the package vignette, by typing `vignette("introduction", package = "tidySummarizedExperiment")`.

All software-related questions should be posted to the Bioconductor Support Site:

<https://support.bioconductor.org>

The code can be viewed at the GitHub repository:

<https://github.com/stemangiola/tidySummarizedExperiment>

Author(s)

Stefano Mangiola

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. Nat Methods 21, 1166–1170 (2024). doi:10.1038/s41592024022992

See Also

Useful links:

- <https://github.com/tidyomics/tidySummarizedExperiment>
- Report bugs at <https://github.com/tidyomics/tidySummarizedExperiment/issues>

append_samples	<i>Append samples from multiple SummarizedExperiment objects</i>
----------------	--

Description

Append samples from multiple SummarizedExperiment objects by column-binding them. This function is equivalent to cbind but provides a tidyverse-like interface.

Usage

```
## S3 method for class 'SummarizedExperiment'
append_samples(x, ..., .id = NULL)
```

Arguments

x	First SummarizedExperiment object to combine
...	Additional SummarizedExperiment objects to combine by samples
.id	Object identifier (currently not used)

Value

A combined SummarizedExperiment object

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. Nat Methods 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Examples

```
data(se)
append_samples(se, se)
```

as_tibble	<i>Coerce lists, matrices, and more to data frames</i>
-----------	--

Description

as_tibble() turns an existing object, such as a data frame or matrix, into a so-called tibble, a data frame with class `tbl_df`. This is in contrast with `tibble()`, which builds a tibble from individual columns. as_tibble() is to `tibble()` as `base::as.data.frame()` is to `base::data.frame()`.

as_tibble() is an S3 generic, with methods for:

- `data.frame`: Thin wrapper around the list method that implements tibble's treatment of rownames.
- `matrix`, `poly`, `ts`, `table`
- Default: Other inputs are first coerced with `base::as.data.frame()`.

`as_tibble_row()` converts a vector to a tibble with one row. If the input is a list, all elements must have size one.

`as_tibble_col()` converts a vector to a tibble with one column.

Usage

```
## S3 method for class 'SummarizedExperiment'
as_tibble(
  x,
  ...,
  .name_repair = c("check_unique", "unique", "universal", "minimal"),
  rownames = pkgconfig::get_config("tibble::rownames", NULL)
)
```

Arguments

- | | |
|---------------------------|--|
| <code>x</code> | A data frame, list, matrix, or other object that could reasonably be coerced to a tibble. |
| <code>...</code> | Unused, for extensibility. |
| <code>.name_repair</code> | Treatment of problematic column names: <ul style="list-style-type: none"> • "minimal": No name repair or checks, beyond basic existence, • "unique": Make sure names are unique and not empty, • "check_unique": (default value), no name repair, but check they are unique, • "universal": Make the names unique and syntactic • "unique_quiet": Same as "unique", but "quiet" • "universal_quiet": Same as "universal", but "quiet" • a function: apply custom name repair (e.g., <code>.name_repair = make.names</code> for names in the style of base R). • A purrr-style anonymous function, see rlang::as_function() |

This argument is passed on as `repair` to [vctrs::vec_as_names\(\)](#). See there for more details on these terms and the strategies used to enforce them.

- | | |
|-----------------------|--|
| <code>rownames</code> | How to treat existing row names of a data frame or matrix: <ul style="list-style-type: none"> • NULL: remove row names. This is the default. • NA: keep row names. • A string: the name of a new column. Existing rownames are transferred into this column and the <code>row.names</code> attribute is deleted. No name repair is applied to the new column name, even if <code>x</code> already contains a column of that name. Use <code>as_tibble(rownames_to_column(...))</code> to safeguard against this case. |
|-----------------------|--|

Read more in [rownames](#).

Value

tibble

Row names

The default behavior is to silently remove row names.

New code should explicitly convert row names to a new column using the `rownames` argument.

For existing code that relies on the retention of row names, call `pkgconfig::set_config("tibble::rownames" = NA)` in your script or in your package's `.onLoad()` function.

Life cycle

Using `as_tibble()` for vectors is superseded as of version 3.0.0, prefer the more expressive `as_tibble_row()` and `as_tibble_col()` variants for new code.

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. *Nat Methods* 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Müller, K., Wickham, H. (2023). *tibble: Simple Data Frames*. R package version 3.2.1, <https://CRAN.R-project.org/package=tibble>

See Also

`tibble()` constructs a tibble from individual columns. `enframe()` converts a named vector to a tibble with a column of names and column of values. Name repair is implemented using `vctrs::vec_as_names()`.

Examples

```
tidySummarizedExperiment::pasilla %>%
  as_tibble()

tidySummarizedExperiment::pasilla %>%
  as_tibble(.subset=-c(condition, type))
```

bind_rows

Efficiently bind multiple data frames by row and column

Description

This is an efficient implementation of the common pattern of `'do.call(rbind, dfs)'` or `'do.call(cbind, dfs)'` for binding many data frames into one.

This is an efficient implementation of the common pattern of `'do.call(rbind, dfs)'` or `'do.call(cbind, dfs)'` for binding many data frames into one.

Usage

```
## S3 method for class 'SummarizedExperiment'
bind_rows(..., .id = NULL, add.cell.ids = NULL)

## S3 method for class 'SummarizedExperiment'
bind_cols(..., .id = NULL)

## S3 method for class 'RangedSummarizedExperiment'
bind_cols(..., .id = NULL)
```

Arguments

<code>...</code>	Data frames to combine. Each argument can either be a data frame, a list that could be a data frame, or a list of data frames. When row-binding, columns are matched by name, and any missing columns will be filled with NA. When column-binding, rows are matched by position, so all data frames must have the same number of rows. To match by value, not position, see <code>mutate_joins</code>.
<code>.id</code>	Data frame identifier. When <code>'id'</code> is supplied, a new column of identifiers is created to link each row to its original data frame. The labels are taken from the named arguments to <code>'bind_rows()'</code>. When a list of data frames is supplied, the labels are taken from the names of the list. If no names are found a numeric sequence is used instead.
<code>add.cell.ids</code>	Appends the corresponding values to

Details

The output of `'bind_rows()'` will contain a column if that column appears in any of the inputs.

The output of `'bind_rows()'` will contain a column if that column appears in any of the inputs.

Value

`'bind_rows()'` and `'bind_cols()'` return the same type as the first input, either a data frame, `'tbl_df'`, or `'grouped_df'`.

`'bind_rows()'` and `'bind_cols()'` return the same type as the first input, either a data frame, `'tbl_df'`, or `'grouped_df'`.

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. *Nat Methods* 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. *Nat Methods* 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Examples

```
print("small_pbmc |> bind_rows(small_pbmc)")
```

```
print("small_pbmc |> bind_cols(annotation_column)")
```

count

*Count the observations in each group***Description**

`count()` lets you quickly count the unique values of one or more variables: `df %>% count(a, b)` is roughly equivalent to `df %>% group_by(a, b) %>% summarise(n = n())`. `count()` is paired with `tally()`, a lower-level helper that is equivalent to `df %>% summarise(n = n())`. Supply `wt` to perform weighted counts, switching the summary from `n = n()` to `n = sum(wt)`.

`add_count()` and `add_tally()` are equivalents to `count()` and `tally()` but use `mutate()` instead of `summarise()` so that they add a new column with group-wise counts.

Usage

```
## S3 method for class 'SummarizedExperiment'
count(
  x,
  ...,
  wt = NULL,
  sort = FALSE,
  name = NULL,
  .drop = group_by_drop_default(x)
)
```

Arguments

<code>x</code>	A data frame, data frame extension (e.g. a tibble), or a lazy data frame (e.g. from <code>dbplyr</code> or <code>dtplyr</code>).
<code>...</code>	<data-masking> Variables to group by.
<code>wt</code>	<data-masking> Frequency weights. Can be <code>NULL</code> or a variable: • If <code>NULL</code> (the default), counts the number of rows in each group. • If a variable, computes <code>sum(wt)</code> for each group.
<code>sort</code>	If <code>TRUE</code> , will show the largest groups at the top.
<code>name</code>	The name of the new column in the output. If omitted, it will default to <code>n</code> . If there's already a column called <code>n</code> , it will use <code>nn</code> . If there's a column called <code>n</code> and <code>nn</code> , it'll use <code>nnn</code> , and so on, adding <code>ns</code> until it gets a new name.
<code>.drop</code>	Handling of factor levels that don't appear in the data, passed on to group_by() . For <code>count()</code> : if <code>FALSE</code> will include counts for empty groups (i.e. for levels of factors that don't exist in the data). [Deprecated] For <code>add_count()</code> : deprecated since it can't actually affect the output.

Value

An object of the same type as `.data`. `count()` and `add_count()` group transiently, so the output has the same groups as the input.

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. Nat Methods 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., François, R., Henry, L., Müller, K., Vaughan, D. (2023). dplyr: A Grammar of Data Manipulation. R package version 2.1.4, <https://CRAN.R-project.org/package=dplyr>

Examples

```
data(se)
se |> count(dex)
```

distinct	Keep distinct/unique rows
----------	---------------------------

Description

Keep only unique/distinct rows from a data frame. This is similar to `unique.data.frame()` but considerably faster.

Usage

```
## S3 method for class 'SummarizedExperiment'
distinct(.data, ..., .keep_all = FALSE)
```

Arguments

.data	A data frame, data frame extension (e.g. a tibble), or a lazy data frame (e.g. from dbplyr or dtplyr). See <i>Methods</i> , below, for more details.
...	<data-masking> Optional variables to use when determining uniqueness. If there are multiple rows for a given combination of inputs, only the first row will be preserved. If omitted, will use all variables in the data frame.
.keep_all	If TRUE, keep all variables in .data. If a combination of ... is not distinct, this keeps the first row of values.

Value

An object of the same type as .data. The output has the following properties:

- Rows are a subset of the input but appear in the same order.
- Columns are not modified if ... is empty or .keep_all is TRUE. Otherwise, distinct() first calls mutate() to create new columns.
- Groups are not modified.
- Data frame attributes are preserved.

Methods

This function is a **generic**, which means that packages can provide implementations (methods) for other classes. See the documentation of individual methods for extra arguments and differences in behaviour.

The following methods are currently available in loaded packages: no methods found.

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. *Nat Methods* 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., François, R., Henry, L., Müller, K., Vaughan, D. (2023). *dplyr: A Grammar of Data Manipulation*. R package version 2.1.4, <https://CRAN.R-project.org/package=dplyr>

Examples

```
data(pasilla)
pasilla |> distinct(.sample)
```

extract	<i>Extract a character column into multiple columns using regular expression groups</i>
---------	---

Description

[Superseded]

`extract()` has been superseded in favour of `separate_wider_regex()` because it has a more polished API and better handling of problems. Superseded functions will not go away, but will only receive critical bug fixes.

Given a regular expression with capturing groups, `extract()` turns each group into a new column. If the groups don't match, or the input is NA, the output will be NA.

Usage

```
## S3 method for class 'SummarizedExperiment'
extract(
  data,
  col,
  into,
  regex = "[[:alnum:]]+",
  remove = TRUE,
  convert = FALSE,
  ...
)
```

Arguments

data	A data frame.
col	<tidy-select> Column to expand.
into	Names of new variables to create as character vector. Use NA to omit the variable in the output.
regex	A string representing a regular expression used to extract the desired values. There should be one group (defined by ()) for each element of into.
remove	If TRUE, remove input column from output data frame.
convert	If TRUE, will run <code>type.convert()</code> with <code>as.is = TRUE</code> on new columns. This is useful if the component columns are integer, numeric or logical. NB: this will cause string "NA"s to be converted to NAs.
...	Additional arguments passed on to methods.

Value

tidySummarizedExperiment

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. Nat Methods 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., Vaughan, D. (2023). tidyr: Tidy Messy Data. R package version 2.0.0, <https://CRAN.R-project.org/package=tidyr>

See Also

`separate()` to split up by a separator.

Examples

```
tidySummarizedExperiment::pasilla |>
  extract(type, into="sequencing", regex="([a-z]*)_end", convert=TRUE)
```

filter	<i>Keep rows that match a condition</i>
--------	---

Description

The `filter()` function is used to subset a data frame, retaining all rows that satisfy your conditions. To be retained, the row must produce a value of TRUE for all conditions. Note that when a condition evaluates to NA the row will be dropped, unlike base subsetting with `[]`.

Usage

```
## S3 method for class 'SummarizedExperiment'
filter(.data, ..., .preserve = FALSE)
```

Arguments

<code>.data</code>	A data frame, data frame extension (e.g. a tibble), or a lazy data frame (e.g. from <code>dbplyr</code> or <code>dtplyr</code>). See <i>Methods</i> , below, for more details.
<code>...</code>	<data-masking> Expressions that return a logical value, and are defined in terms of the variables in <code>.data</code> . If multiple expressions are included, they are combined with the <code>&</code> operator. Only rows for which all conditions evaluate to <code>TRUE</code> are kept.
<code>.preserve</code>	Relevant when the <code>.data</code> input is grouped. If <code>.preserve = FALSE</code> (the default), the grouping structure is recalculated based on the resulting data, otherwise the grouping is kept as is.

Details

The `filter()` function is used to subset the rows of `.data`, applying the expressions in `...` to the column values to determine which rows should be retained. It can be applied to both grouped and ungrouped data (see [group_by\(\)](#) and [ungroup\(\)](#)). However, `dplyr` is not yet smart enough to optimise the filtering operation on grouped datasets that do not need grouped calculations. For this reason, filtering is often considerably faster on ungrouped data.

Value

An object of the same type as `.data`. The output has the following properties:

- Rows are a subset of the input, but appear in the same order.
- Columns are not modified.
- The number of groups may be reduced (if `.preserve` is not `TRUE`).
- Data frame attributes are preserved.

Useful filter functions

There are many functions and operators that are useful when constructing the expressions used to filter the data:

- `==`, `>`, `>=` etc
- `&`, `|`, `!`, [xor\(\)](#)
- [is.na\(\)](#)
- [between\(\)](#), [near\(\)](#)

Grouped tibbles

Because filtering expressions are computed within groups, they may yield different results on grouped tibbles. This will be the case as soon as an aggregating, lagging, or ranking function is involved. Compare this ungrouped filtering:

```
starwars %>% filter(mass > mean(mass, na.rm = TRUE))
```

With the grouped equivalent:

```
starwars %>% group_by(gender) %>% filter(mass > mean(mass, na.rm = TRUE))
```

In the ungrouped version, `filter()` compares the value of `mass` in each row to the global average (taken over the whole data set), keeping only the rows with `mass` greater than this global average. In contrast, the grouped version calculates the average `mass` separately for each gender group, and keeps rows with `mass` greater than the relevant within-gender average.

Methods

This function is a **generic**, which means that packages can provide implementations (methods) for other classes. See the documentation of individual methods for extra arguments and differences in behaviour.

The following methods are currently available in loaded packages: no methods found.

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. Nat Methods 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., François, R., Henry, L., Müller, K., Vaughan, D. (2023). dplyr: A Grammar of Data Manipulation. R package version 2.1.4, <https://CRAN.R-project.org/package=dplyr>

See Also

Other single table verbs: [arrange\(\)](#), [mutate\(\)](#), [reframe\(\)](#), [rename\(\)](#), [select\(\)](#), [slice\(\)](#), [summarise\(\)](#)

Examples

```
data(pasilla)
pasilla |> filter(.sample == "untrt1")

# Learn more in ?dplyr_tidy_eval
```

full_join	<i>Mutating joins</i>
-----------	-----------------------

Description

Mutating joins add columns from y to x, matching observations based on the keys. There are four mutating joins: the inner join, and the three outer joins.

Inner join:

An `inner_join()` only keeps observations from x that have a matching key in y. The most important property of an inner join is that unmatched rows in either input are not included in the result. This means that generally inner joins are not appropriate in most analyses, because it is too easy to lose observations.

- Outer joins:**
- The three outer joins keep observations that appear in at least one of the data frames:
- A `left_join()` keeps all observations in x.
 - A `right_join()` keeps all observations in y.
 - A `full_join()` keeps all observations in x and y.

Usage

```
## S3 method for class 'SummarizedExperiment'
full_join(x, y, by = NULL, copy = FALSE, suffix = c(".x", ".y"), ...)
```

Arguments

<code>x, y</code>	A pair of data frames, data frame extensions (e.g. a tibble), or lazy data frames (e.g. from <code>dbplyr</code> or <code>dtplyr</code>). See <i>Methods</i> , below, for more details.
<code>by</code>	A join specification created with <code>join_by()</code>, or a character vector of variables to join by. If <code>NULL</code>, the default, <code>*_join()</code> will perform a natural join, using all variables in common across <code>x</code> and <code>y</code>. A message lists the variables so that you can check they're correct; suppress the message by supplying <code>by</code> explicitly. To join on different variables between <code>x</code> and <code>y</code>, use a <code>join_by()</code> specification. For example, <code>join_by(a == b)</code> will match <code>x\$a</code> to <code>y\$b</code>. To join by multiple variables, use a <code>join_by()</code> specification with multiple expressions. For example, <code>join_by(a == b, c == d)</code> will match <code>x\$a</code> to <code>y\$b</code> and <code>x\$c</code> to <code>y\$d</code>. If the column names are the same between <code>x</code> and <code>y</code>, you can shorten this by listing only the variable names, like <code>join_by(a, c)</code>. <code>join_by()</code> can also be used to perform inequality, rolling, and overlap joins. See the documentation at ?join_by for details on these types of joins. For simple equality joins, you can alternatively specify a character vector of variable names to join by. For example, <code>by = c("a", "b")</code> joins <code>x\$a</code> to <code>y\$a</code> and <code>x\$b</code> to <code>y\$b</code>. If variable names differ between <code>x</code> and <code>y</code>, use a named character vector like <code>by = c("x_a" = "y_a", "x_b" = "y_b")</code>. To perform a cross-join, generating all combinations of <code>x</code> and <code>y</code>, see <code>cross_join()</code>.
<code>copy</code>	If <code>x</code> and <code>y</code> are not from the same data source, and <code>copy</code> is <code>TRUE</code> , then <code>y</code> will be copied into the same <code>src</code> as <code>x</code> . This allows you to join tables across <code>srcs</code> , but it is a potentially expensive operation so you must opt into it.
<code>suffix</code>	If there are non-joined duplicate variables in <code>x</code> and <code>y</code> , these suffixes will be added to the output to disambiguate them. Should be a character vector of length 2.
<code>...</code>	Other parameters passed onto methods.

Value

An object of the same type as `x` (including the same groups). The order of the rows and columns of `x` is preserved as much as possible. The output has the following properties:

- The rows are affected by the join type.
 - `inner_join()` returns matched `x` rows.
 - `left_join()` returns all `x` rows.
 - `right_join()` returns matched of `x` rows, followed by unmatched `y` rows.
 - `full_join()` returns all `x` rows, followed by unmatched `y` rows.
- Output columns include all columns from `x` and all non-key columns from `y`. If `keep = TRUE`, the key columns from `y` are included as well.
- If non-key columns in `x` and `y` have the same name, suffixes are added to disambiguate. If `keep = TRUE` and key columns in `x` and `y` have the same name, suffixes are added to disambiguate these as well.
- If `keep = FALSE`, output columns included in `by` are coerced to their common type between `x` and `y`.

Many-to-many relationships

By default, dplyr guards against many-to-many relationships in equality joins by throwing a warning. These occur when both of the following are true:

- A row in *x* matches multiple rows in *y*.
- A row in *y* matches multiple rows in *x*.

This is typically surprising, as most joins involve a relationship of one-to-one, one-to-many, or many-to-one, and is often the result of an improperly specified join. Many-to-many relationships are particularly problematic because they can result in a Cartesian explosion of the number of rows returned from the join.

If a many-to-many relationship is expected, silence this warning by explicitly setting `relationship = "many-to-many"`.

In production code, it is best to preemptively set `relationship` to whatever relationship you expect to exist between the keys of *x* and *y*, as this forces an error to occur immediately if the data doesn't align with your expectations.

Inequality joins typically result in many-to-many relationships by nature, so they don't warn on them by default, but you should still take extra care when specifying an inequality join, because they also have the capability to return a large number of rows.

Rolling joins don't warn on many-to-many relationships either, but many rolling joins follow a many-to-one relationship, so it is often useful to set `relationship = "many-to-one"` to enforce this.

Note that in SQL, most database providers won't let you specify a many-to-many relationship between two tables, instead requiring that you create a third *junction table* that results in two one-to-many relationships instead.

Methods

These functions are **generics**, which means that packages can provide implementations (methods) for other classes. See the documentation of individual methods for extra arguments and differences in behaviour.

Methods available in currently loaded packages:

- `inner_join()`: no methods found.
- `left_join()`: no methods found.
- `right_join()`: no methods found.
- `full_join()`: no methods found.

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. *Nat Methods* 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., François, R., Henry, L., Müller, K., Vaughan, D. (2023). *dplyr: A Grammar of Data Manipulation*. R package version 2.1.4, <https://CRAN.R-project.org/package=dplyr>

See Also

Other joins: [cross_join\(\)](#), [filter-joins](#), [nest_join\(\)](#)

Examples

```
data(pasilla)

tt <- pasilla
tt |> full_join(tibble::tibble(condition="treated", dose=10))
```

ggplot

Create a new ggplot from a SummarizedExperiment

Description

ggplot() initializes a ggplot object. It can be used to declare the input data frame for a graphic and to specify the set of plot aesthetics intended to be common throughout all subsequent layers unless specifically overridden.

Usage

```
## S3 method for class 'SummarizedExperiment'
ggplot(data = NULL, mapping = aes(), ..., environment = parent.frame())
```

Arguments

data	Default dataset to use for plot. If not already a data.frame, will be converted to one by fortify() . If not specified, must be supplied in each layer added to the plot.
mapping	Default list of aesthetic mappings to use for plot. If not specified, must be supplied in each layer added to the plot.
...	Other arguments passed on to methods. Not currently used.
environment	[Deprecated] Used prior to tidy evaluation.

Details

ggplot() is used to construct the initial plot object, and is almost always followed by a plus sign (+) to add components to the plot.

There are three common patterns used to invoke ggplot():

- ggplot(data = df, mapping = aes(x, y, other aesthetics))
- ggplot(data = df)
- ggplot()

The first pattern is recommended if all layers use the same data and the same set of aesthetics, although this method can also be used when adding a layer using data from another data frame.

The second pattern specifies the default data frame to use for the plot, but no aesthetics are defined up front. This is useful when one data frame is used predominantly for the plot, but the aesthetics vary from one layer to another.

The third pattern initializes a skeleton ggplot object, which is fleshed out as layers are added. This is useful when multiple data frames are used to produce different layers, as is often the case in complex graphics.

The data = and mapping = specifications in the arguments are optional (and are often omitted in practice), so long as the data and the mapping values are passed into the function in the right order. In the examples below, however, they are left in place for clarity.

Value

ggplot

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. *Nat Methods* 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H. (2016). *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York. ISBN 978-3-319-24277-4, <https://ggplot2.tidyverse.org>

See Also

The [first steps chapter](#) of the online ggplot2 book.

Examples

```
library(ggplot2)
data(pasilla)
pasilla %>%
  ggplot(aes(.sample, counts)) +
  geom_boxplot()
```

group_by	<i>Group by one or more variables</i>
----------	---------------------------------------

Description

Most data operations are done on groups defined by variables. `group_by()` takes an existing `tbl` and converts it into a grouped `tbl` where operations are performed "by group". `ungroup()` removes grouping.

Usage

```
## S3 method for class 'SummarizedExperiment'
group_by(.data, ..., .add = FALSE, .drop = group_by_drop_default(.data))
```

Arguments

<code>.data</code>	A data frame, data frame extension (e.g. a tibble), or a lazy data frame (e.g. from <code>dbplyr</code> or <code>dtplyr</code>). See <i>Methods</i> , below, for more details.
<code>...</code>	In <code>group_by()</code> , variables or computations to group by. Computations are always done on the ungrouped data frame. To perform computations on the grouped data, you need to use a separate <code>mutate()</code> step before the <code>group_by()</code> . Computations are not allowed in <code>nest_by()</code> . In <code>ungroup()</code> , variables to remove from the grouping.
<code>.add</code>	When <code>FALSE</code> , the default, <code>group_by()</code> will override existing groups. To add to the existing groups, use <code>.add = TRUE</code> . This argument was previously called <code>add</code> , but that prevented creating a new grouping variable called <code>add</code> , and conflicts with our naming conventions.

`.drop` Drop groups formed by factor levels that don't appear in the data? The default is TRUE except when `.data` has been previously grouped with `.drop = FALSE`. See [group_by_drop_default\(\)](#) for details.

Value

A grouped data frame with class `grouped_df`, unless the combination of `...` and `add` yields a empty set of grouping columns, in which case a tibble will be returned.

Methods

These function are **generics**, which means that packages can provide implementations (methods) for other classes. See the documentation of individual methods for extra arguments and differences in behaviour.

Methods available in currently loaded packages:

- `group_by()`: no methods found.
- `ungroup()`: no methods found.

Ordering

Currently, `group_by()` internally orders the groups in ascending order. This results in ordered output from functions that aggregate groups, such as [summarise\(\)](#).

When used as grouping columns, character vectors are ordered in the C locale for performance and reproducibility across R sessions. If the resulting ordering of your grouped operation matters and is dependent on the locale, you should follow up the grouped operation with an explicit call to [arrange\(\)](#) and set the `.locale` argument. For example:

```
data %>%
  group_by(chr) %>%
  summarise(avg = mean(x)) %>%
  arrange(chr, .locale = "en")
```

This is often useful as a preliminary step before generating content intended for humans, such as an HTML table.

Legacy behavior:

Prior to dplyr 1.1.0, character vector grouping columns were ordered in the system locale. If you need to temporarily revert to this behavior, you can set the global option `dplyr.legacy_locale` to TRUE, but this should be used sparingly and you should expect this option to be removed in a future version of dplyr. It is better to update existing code to explicitly call `arrange(.locale = )` instead. Note that setting `dplyr.legacy_locale` will also force calls to [arrange\(\)](#) to use the system locale.

References

- Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. *Nat Methods* 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>
- Wickham, H., François, R., Henry, L., Müller, K., Vaughan, D. (2023). dplyr: A Grammar of Data Manipulation. R package version 2.1.4, <https://CRAN.R-project.org/package=dplyr>

See Also

Other grouping functions: [group_map\(\)](#), [group_nest\(\)](#), [group_split\(\)](#), [group_trim\(\)](#)

Examples

```
data(pasilla)
pasilla |> group_by(.sample)
```

group_split	<i>Split data frame by groups</i>
-------------	-----------------------------------

Description**[Experimental]**

[group_split\(\)](#) works like [base::split\(\)](#) but:

- It uses the grouping structure from [group_by\(\)](#) and therefore is subject to the data mask
- It does not name the elements of the list based on the grouping as this only works well for a single character grouping variable. Instead, use [group_keys\(\)](#) to access a data frame that defines the groups.

[group_split\(\)](#) is primarily designed to work with grouped data frames. You can pass `...` to group and split an ungrouped data frame, but this is generally not very useful as you want have easy access to the group metadata.

Usage

```
## S3 method for class 'SummarizedExperiment'
group_split(.tbl, ..., .keep = TRUE)
```

Arguments

<code>.tbl</code>	A tbl.
<code>...</code>	If <code>.tbl</code> is an ungrouped data frame, a grouping specification, forwarded to group_by() .
<code>.keep</code>	Should the grouping columns be kept?

Value

A list of tibbles. Each tibble contains the rows of `.tbl` for the associated group and all the columns, including the grouping variables. Note that this returns a [list_of](#) which is slightly stricter than a simple list but is useful for representing lists where every element has the same type.

Lifecycle

[group_split\(\)](#) is not stable because you can achieve very similar results by manipulating the nested column returned from [tidyr::nest\(.by =\)](#). That also retains the group keys all within a single data structure. [group_split\(\)](#) may be deprecated in the future.

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. *Nat Methods* 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., François, R., Henry, L., Müller, K., Vaughan, D. (2023). *dplyr: A Grammar of Data Manipulation*. R package version 2.1.4, <https://CRAN.R-project.org/package=dplyr>

See Also

Other grouping functions: `group_by()`, `group_map()`, `group_nest()`, `group_trim()`

Examples

```
data(pasilla, package = "tidySummarizedExperiment")
pasilla |> group_split(condition)
pasilla |> group_split(counts > 0)
pasilla |> group_split(condition, counts > 0)
```

inner_join

Mutating joins

Description

Mutating joins add columns from *y* to *x*, matching observations based on the keys. There are four mutating joins: the inner join, and the three outer joins.

Inner join:

An `inner_join()` only keeps observations from *x* that have a matching key in *y*.

The most important property of an inner join is that unmatched rows in either input are not included in the result. This means that generally inner joins are not appropriate in most analyses, because it is too easy to lose observations.

Outer joins:

The three outer joins keep observations that appear in at least one of the data frames:

- A `left_join()` keeps all observations in *x*.
- A `right_join()` keeps all observations in *y*.
- A `full_join()` keeps all observations in *x* and *y*.

Usage

```
## S3 method for class 'SummarizedExperiment'
inner_join(x, y, by = NULL, copy = FALSE, suffix = c(".x", ".y"), ...)
```

Arguments

<code>x, y</code>	A pair of data frames, data frame extensions (e.g. a tibble), or lazy data frames (e.g. from <code>dbplyr</code> or <code>dtplyr</code>). See <i>Methods</i> , below, for more details.
<code>by</code>	A join specification created with <code>join_by()</code>, or a character vector of variables to join by. If <code>NULL</code>, the default, <code>*_join()</code> will perform a natural join, using all variables in common across <code>x</code> and <code>y</code>. A message lists the variables so that you can check they're correct; suppress the message by supplying <code>by</code> explicitly. To join on different variables between <code>x</code> and <code>y</code>, use a <code>join_by()</code> specification. For example, <code>join_by(a == b)</code> will match <code>x\$a</code> to <code>y\$b</code>. To join by multiple variables, use a <code>join_by()</code> specification with multiple expressions. For example, <code>join_by(a == b, c == d)</code> will match <code>x\$a</code> to <code>y\$b</code> and <code>x\$c</code> to <code>y\$d</code>. If the column names are the same between <code>x</code> and <code>y</code>, you can shorten this by listing only the variable names, like <code>join_by(a, c)</code>. <code>join_by()</code> can also be used to perform inequality, rolling, and overlap joins. See the documentation at ?join_by for details on these types of joins. For simple equality joins, you can alternatively specify a character vector of variable names to join by. For example, <code>by = c("a", "b")</code> joins <code>x\$a</code> to <code>y\$a</code> and <code>x\$b</code> to <code>y\$b</code>. If variable names differ between <code>x</code> and <code>y</code>, use a named character vector like <code>by = c("x_a" = "y_a", "x_b" = "y_b")</code>. To perform a cross-join, generating all combinations of <code>x</code> and <code>y</code>, see <code>cross_join()</code>.
<code>copy</code>	If <code>x</code> and <code>y</code> are not from the same data source, and <code>copy</code> is <code>TRUE</code> , then <code>y</code> will be copied into the same <code>src</code> as <code>x</code> . This allows you to join tables across <code>srcs</code> , but it is a potentially expensive operation so you must opt into it.
<code>suffix</code>	If there are non-joined duplicate variables in <code>x</code> and <code>y</code> , these suffixes will be added to the output to disambiguate them. Should be a character vector of length 2.
<code>...</code>	Other parameters passed onto methods.

Value

An object of the same type as `x` (including the same groups). The order of the rows and columns of `x` is preserved as much as possible. The output has the following properties:

- The rows are affected by the join type.
 - `inner_join()` returns matched `x` rows.
 - `left_join()` returns all `x` rows.
 - `right_join()` returns matched of `x` rows, followed by unmatched `y` rows.
 - `full_join()` returns all `x` rows, followed by unmatched `y` rows.
- Output columns include all columns from `x` and all non-key columns from `y`. If `keep = TRUE`, the key columns from `y` are included as well.
- If non-key columns in `x` and `y` have the same name, suffixes are added to disambiguate. If `keep = TRUE` and key columns in `x` and `y` have the same name, suffixes are added to disambiguate these as well.
- If `keep = FALSE`, output columns included in `by` are coerced to their common type between `x` and `y`.

Many-to-many relationships

By default, dplyr guards against many-to-many relationships in equality joins by throwing a warning. These occur when both of the following are true:

- A row in *x* matches multiple rows in *y*.
- A row in *y* matches multiple rows in *x*.

This is typically surprising, as most joins involve a relationship of one-to-one, one-to-many, or many-to-one, and is often the result of an improperly specified join. Many-to-many relationships are particularly problematic because they can result in a Cartesian explosion of the number of rows returned from the join.

If a many-to-many relationship is expected, silence this warning by explicitly setting `relationship = "many-to-many"`.

In production code, it is best to preemptively set `relationship` to whatever relationship you expect to exist between the keys of *x* and *y*, as this forces an error to occur immediately if the data doesn't align with your expectations.

Inequality joins typically result in many-to-many relationships by nature, so they don't warn on them by default, but you should still take extra care when specifying an inequality join, because they also have the capability to return a large number of rows.

Rolling joins don't warn on many-to-many relationships either, but many rolling joins follow a many-to-one relationship, so it is often useful to set `relationship = "many-to-one"` to enforce this.

Note that in SQL, most database providers won't let you specify a many-to-many relationship between two tables, instead requiring that you create a third *junction table* that results in two one-to-many relationships instead.

Methods

These functions are **generics**, which means that packages can provide implementations (methods) for other classes. See the documentation of individual methods for extra arguments and differences in behaviour.

Methods available in currently loaded packages:

- `inner_join()`: no methods found.
- `left_join()`: no methods found.
- `right_join()`: no methods found.
- `full_join()`: no methods found.

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. *Nat Methods* 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., François, R., Henry, L., Müller, K., Vaughan, D. (2023). *dplyr: A Grammar of Data Manipulation*. R package version 2.1.4, <https://CRAN.R-project.org/package=dplyr>

See Also

Other joins: [cross_join\(\)](#), [filter-joins](#), [nest_join\(\)](#)

Examples

```
data(pasilla)

tt <- pasilla
tt |> inner_join(tt |>
  distinct(condition) |>
  mutate(new_column=1:2) |>
  slice(1))
```

left_join

Mutating joins

Description

Mutating joins add columns from *y* to *x*, matching observations based on the keys. There are four mutating joins: the inner join, and the three outer joins.

Inner join:

An `inner_join()` only keeps observations from *x* that have a matching key in *y*.

The most important property of an inner join is that unmatched rows in either input are not included in the result. This means that generally inner joins are not appropriate in most analyses, because it is too easy to lose observations.

Outer joins:

The three outer joins keep observations that appear in at least one of the data frames:

- A `left_join()` keeps all observations in *x*.
- A `right_join()` keeps all observations in *y*.
- A `full_join()` keeps all observations in *x* and *y*.

Usage

```
## S3 method for class 'SummarizedExperiment'
left_join(x, y, by = NULL, copy = FALSE, suffix = c(".x", ".y"), ...)
```

Arguments

- | | |
|------|--|
| x, y | A pair of data frames, data frame extensions (e.g. a tibble), or lazy data frames (e.g. from dbplyr or dtplyr). See <i>Methods</i> , below, for more details. |
| by | <p>A join specification created with <code>join_by()</code>, or a character vector of variables to join by.</p> <p>If NULL, the default, <code>*_join()</code> will perform a natural join, using all variables in common across <i>x</i> and <i>y</i>. A message lists the variables so that you can check they're correct; suppress the message by supplying <code>by</code> explicitly.</p> <p>To join on different variables between <i>x</i> and <i>y</i>, use a <code>join_by()</code> specification. For example, <code>join_by(a == b)</code> will match <i>x</i>\$a to <i>y</i>\$b.</p> <p>To join by multiple variables, use a <code>join_by()</code> specification with multiple expressions. For example, <code>join_by(a == b, c == d)</code> will match <i>x</i>\$a to <i>y</i>\$b and <i>x</i>\$c to <i>y</i>\$d. If the column names are the same between <i>x</i> and <i>y</i>, you can shorten this by listing only the variable names, like <code>join_by(a, c)</code>.</p> |

`join_by()` can also be used to perform inequality, rolling, and overlap joins. See the documentation at [?join_by](#) for details on these types of joins.

For simple equality joins, you can alternatively specify a character vector of variable names to join by. For example, `by = c("a", "b")` joins `x$a` to `y$a` and `x$b` to `y$b`. If variable names differ between `x` and `y`, use a named character vector like `by = c("x_a" = "y_a", "x_b" = "y_b")`.

To perform a cross-join, generating all combinations of `x` and `y`, see [cross_join\(\)](#).

copy	If <code>x</code> and <code>y</code> are not from the same data source, and <code>copy</code> is <code>TRUE</code> , then <code>y</code> will be copied into the same <code>src</code> as <code>x</code> . This allows you to join tables across <code>srcs</code> , but it is a potentially expensive operation so you must opt into it.
suffix	If there are non-joined duplicate variables in <code>x</code> and <code>y</code> , these suffixes will be added to the output to disambiguate them. Should be a character vector of length 2.
...	Other parameters passed onto methods.

Value

An object of the same type as `x` (including the same groups). The order of the rows and columns of `x` is preserved as much as possible. The output has the following properties:

- The rows are affected by the join type.
 - `inner_join()` returns matched `x` rows.
 - `left_join()` returns all `x` rows.
 - `right_join()` returns matched of `x` rows, followed by unmatched `y` rows.
 - `full_join()` returns all `x` rows, followed by unmatched `y` rows.
- Output columns include all columns from `x` and all non-key columns from `y`. If `keep = TRUE`, the key columns from `y` are included as well.
- If non-key columns in `x` and `y` have the same name, suffixes are added to disambiguate. If `keep = TRUE` and key columns in `x` and `y` have the same name, suffixes are added to disambiguate these as well.
- If `keep = FALSE`, output columns included in `by` are coerced to their common type between `x` and `y`.

Many-to-many relationships

By default, `dplyr` guards against many-to-many relationships in equality joins by throwing a warning. These occur when both of the following are true:

- A row in `x` matches multiple rows in `y`.
- A row in `y` matches multiple rows in `x`.

This is typically surprising, as most joins involve a relationship of one-to-one, one-to-many, or many-to-one, and is often the result of an improperly specified join. Many-to-many relationships are particularly problematic because they can result in a Cartesian explosion of the number of rows returned from the join.

If a many-to-many relationship is expected, silence this warning by explicitly setting `relationship = "many-to-many"`.

In production code, it is best to preemptively set `relationship` to whatever relationship you expect to exist between the keys of `x` and `y`, as this forces an error to occur immediately if the data doesn't align with your expectations.

Inequality joins typically result in many-to-many relationships by nature, so they don't warn on them by default, but you should still take extra care when specifying an inequality join, because they also have the capability to return a large number of rows.

Rolling joins don't warn on many-to-many relationships either, but many rolling joins follow a many-to-one relationship, so it is often useful to set `relationship = "many-to-one"` to enforce this.

Note that in SQL, most database providers won't let you specify a many-to-many relationship between two tables, instead requiring that you create a third *junction table* that results in two one-to-many relationships instead.

Methods

These functions are **generics**, which means that packages can provide implementations (methods) for other classes. See the documentation of individual methods for extra arguments and differences in behaviour.

Methods available in currently loaded packages:

- `inner_join()`: no methods found.
- `left_join()`: no methods found.
- `right_join()`: no methods found.
- `full_join()`: no methods found.

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. *Nat Methods* 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., François, R., Henry, L., Müller, K., Vaughan, D. (2023). *dplyr: A Grammar of Data Manipulation*. R package version 2.1.4, <https://CRAN.R-project.org/package=dplyr>

See Also

Other joins: [cross_join\(\)](#), [filter-joins](#), [nest_join\(\)](#)

Examples

```
data(pasilla)

tt <- pasilla
tt |> left_join(tt |>
  distinct(condition) |>
  mutate(new_column=1:2) |>
  slice(1))
```

mutate

*Create, modify, and delete columns***Description**

`mutate()` creates new columns that are functions of existing variables. It can also modify (if the name is the same as an existing column) and delete columns (by setting their value to `NULL`).

Usage

```
## S3 method for class 'SummarizedExperiment'
mutate(.data, ...)
```

Arguments

`.data` A data frame, data frame extension (e.g. a tibble), or a lazy data frame (e.g. from `dbplyr` or `dtplyr`). See *Methods*, below, for more details.

`...` [<data-masking>](#) Name-value pairs. The name gives the name of the column in the output.
The value can be:

- A vector of length 1, which will be recycled to the correct length.
- A vector the same length as the current group (or the whole data frame if ungrouped).
- `NULL`, to remove the column.
- A data frame or tibble, to create multiple columns in the output.

Value

An object of the same type as `.data`. The output has the following properties:

- Columns from `.data` will be preserved according to the `.keep` argument.
- Existing columns that are modified by `...` will always be returned in their original location.
- New columns created through `...` will be placed according to the `.before` and `.after` arguments.
- The number of rows is not affected.
- Columns given the value `NULL` will be removed.
- Groups will be recomputed if a grouping variable is mutated.
- Data frame attributes are preserved.

Useful mutate functions

- `+`, `-`, `log()`, etc., for their usual mathematical meanings
- `lead()`, `lag()`
- `dense_rank()`, `min_rank()`, `percent_rank()`, `row_number()`, `cume_dist()`, `ntile()`
- `cumsum()`, `cummean()`, `cummin()`, `cummax()`, `cumany()`, `cumall()`
- `na_if()`, `coalesce()`
- `if_else()`, `recode()`, `case_when()`

Grouped tibbles

Because mutating expressions are computed within groups, they may yield different results on grouped tibbles. This will be the case as soon as an aggregating, lagging, or ranking function is involved. Compare this ungrouped mutate:

```
starwars %>%
  select(name, mass, species) %>%
  mutate(mass_norm = mass / mean(mass, na.rm = TRUE))
```

With the grouped equivalent:

```
starwars %>%
  select(name, mass, species) %>%
  group_by(species) %>%
  mutate(mass_norm = mass / mean(mass, na.rm = TRUE))
```

The former normalises mass by the global average whereas the latter normalises by the averages within species levels.

Methods

This function is a **generic**, which means that packages can provide implementations (methods) for other classes. See the documentation of individual methods for extra arguments and differences in behaviour.

Methods available in currently loaded packages: no methods found.

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. Nat Methods 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., François, R., Henry, L., Müller, K., Vaughan, D. (2023). dplyr: A Grammar of Data Manipulation. R package version 2.1.4, <https://CRAN.R-project.org/package=dplyr>

See Also

Other single table verbs: [rename\(\)](#), [slice\(\)](#), [summarise\(\)](#)

Examples

```
data(pasilla)
pasilla |> mutate(logcounts=log2(counts))
```

nest

*Nest rows into a list-column of data frames***Description**

Nesting creates a list-column of data frames; unnesting flattens it back out into regular columns. Nesting is implicitly a summarising operation: you get one row for each group defined by the non-nested columns. This is useful in conjunction with other summaries that work with whole datasets, most notably models.

Learn more in `vignette("nest")`.

Usage

```
## S3 method for class 'SummarizedExperiment'
nest(.data, ..., .names_sep = NULL)
```

Arguments

<code>.data</code>	A data frame.
<code>...</code>	<code><tidy-select></code> Columns to nest; these will appear in the inner data frames. Specified using name-variable pairs of the form <code>new_col = c(col1, col2, col3)</code>. The right hand side can be any valid tidyselect expression. If not supplied, then <code>...</code> is derived as all columns <i>not</i> selected by <code>.by</code>, and will use the column name from <code>.key</code>. [Deprecated]: previously you could write <code>df %>% nest(x, y, z)</code>. Convert to <code>df %>% nest(data = c(x, y, z))</code>.
<code>.names_sep</code>	If <code>NULL</code>, the default, the inner names will come from the former outer names. If a string, the new inner names will use the outer names with <code>names_sep</code> automatically stripped. This makes <code>names_sep</code> roughly symmetric between nesting and unnesting.

Details

If neither `...` nor `.by` are supplied, `nest()` will nest all variables, and will use the column name supplied through `.key`.

Value

`tidySummarizedExperiment_nested`

New syntax

tidyr 1.0.0 introduced a new syntax for `nest()` and `unnest()` that's designed to be more similar to other functions. Converting to the new syntax should be straightforward (guided by the message you'll receive) but if you just need to run an old analysis, you can easily revert to the previous behaviour using `nest_legacy()` and `unnest_legacy()` as follows:

```
library(tidyr)
nest <- nest_legacy
unnest <- unnest_legacy
```

Grouped data frames

`df %>% nest(data = c(x, y))` specifies the columns to be nested; i.e. the columns that will appear in the inner data frame. `df %>% nest(.by = c(x, y))` specifies the columns to nest *by*; i.e. the columns that will remain in the outer data frame. An alternative way to achieve the latter is to `nest()` a grouped data frame created by `dplyr::group_by()`. The grouping variables remain in the outer data frame and the others are nested. The result preserves the grouping of the input.

Variables supplied to `nest()` will override grouping variables so that `df %>% group_by(x, y) %>% nest(data = !z)` will be equivalent to `df %>% nest(data = !z)`.

You can't supply `.by` with a grouped data frame, as the groups already represent what you are nesting by.

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. *Nat Methods* 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., Vaughan, D. (2023). *tidyr: Tidy Messy Data*. R package version 2.0.0, <https://CRAN.R-project.org/package=tidyr>

Examples

```
tidySummarizedExperiment::pasilla |>
  nest(data=-condition)
```

pasilla	<i>Read counts of RNA-seq samples of Pasilla knock-down by Brooks et al.</i>
---------	--

Description

A SummarizedExperiment dataset containing the transcriptome information for *Drosophila Melanogaster*.

Usage

```
data(pasilla)
```

Format

containing 14599 features and 7 biological replicates.

Source

<https://bioconductor.org/packages/release/data/experiment/html/pasilla.html>

pivot_longer	<i>Pivot data from wide to long</i>
--------------	-------------------------------------

Description

`pivot_longer()` "lengthens" data, increasing the number of rows and decreasing the number of columns. The inverse transformation is [pivot_wider\(\)](#)

Learn more in `vignette("pivot")`.

Usage

```
## S3 method for class 'SummarizedExperiment'
pivot_longer(
  data,
  cols,
  ...,
  cols_vary = "fastest",
  names_to = "name",
  names_prefix = NULL,
  names_sep = NULL,
  names_pattern = NULL,
  names_ptypes = NULL,
  names_transform = NULL,
  names_repair = "check_unique",
  values_to = "value",
  values_drop_na = FALSE,
  values_ptypes = NULL,
  values_transform = NULL
)
```

Arguments

<code>data</code>	A data frame to pivot.
<code>cols</code>	<tidy-select> Columns to pivot into longer format.
<code>...</code>	Additional arguments passed on to methods.
<code>cols_vary</code>	When pivoting <code>cols</code> into longer format, how should the output rows be arranged relative to their original row number? • "fastest", the default, keeps individual rows from <code>cols</code> close together in the output. This often produces intuitively ordered output when you have at least one key column from data that is not involved in the pivoting process. • "slowest" keeps individual columns from <code>cols</code> close together in the output. This often produces intuitively ordered output when you utilize all of the columns from data in the pivoting process.
<code>names_to</code>	A character vector specifying the new column or columns to create from the information stored in the column names of data specified by <code>cols</code> . • If length 0, or if <code>NULL</code> is supplied, no columns will be created. • If length 1, a single column will be created which will contain the column names specified by <code>cols</code>.

- If length > 1, multiple columns will be created. In this case, one of `names_sep` or `names_pattern` must be supplied to specify how the column names should be split. There are also two additional character values you can take advantage of:
 - NA will discard the corresponding component of the column name.
 - `".value"` indicates that the corresponding component of the column name defines the name of the output column containing the cell values, overriding `values_to` entirely.

`names_prefix` A regular expression used to remove matching text from the start of each variable name.

`names_sep, names_pattern`

If `names_to` contains multiple values, these arguments control how the column name is broken up.

`names_sep` takes the same specification as `separate()`, and can either be a numeric vector (specifying positions to break on), or a single string (specifying a regular expression to split on).

`names_pattern` takes the same specification as `extract()`, a regular expression containing matching groups (`()`).

If these arguments do not give you enough control, use `pivot_longer_spec()` to create a spec object and process manually as needed.

`names_ptypes, values_ptypes`

Optionally, a list of column name-prototype pairs. Alternatively, a single empty prototype can be supplied, which will be applied to all columns. A prototype (or ptype for short) is a zero-length vector (like `integer()` or `numeric()`) that defines the type, class, and attributes of a vector. Use these arguments if you want to confirm that the created columns are the types that you expect. Note that if you want to change (instead of confirm) the types of specific columns, you should use `names_transform` or `values_transform` instead.

`names_transform, values_transform`

Optionally, a list of column name-function pairs. Alternatively, a single function can be supplied, which will be applied to all columns. Use these arguments if you need to change the types of specific columns. For example, `names_transform = list(week = as.integer)` would convert a character variable called `week` to an integer.

If not specified, the type of the columns generated from `names_to` will be character, and the type of the variables generated from `values_to` will be the common type of the input columns used to generate them.

`names_repair` What happens if the output has invalid column names? The default, `"check_unique"` is to error if the columns are duplicated. Use `"minimal"` to allow duplicates in the output, or `"unique"` to de-duplicated by adding numeric suffixes. See `vctrs::vec_as_names()` for more options.

`values_to` A string specifying the name of the column to create from the data stored in cell values. If `names_to` is a character containing the special `.value` sentinel, this value will be ignored, and the name of the value column will be derived from part of the existing column names.

`values_drop_na` If TRUE, will drop rows that contain only NAs in the `values_to` column. This effectively converts explicit missing values to implicit missing values, and should generally be used only when missing values in data were created by its structure.

Details

`pivot_longer()` is an updated approach to `gather()`, designed to be both simpler to use and to handle more use cases. We recommend you use `pivot_longer()` for new code; `gather()` isn't going away but is no longer under active development.

Value

`tidySummarizedExperiment`

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. *Nat Methods* 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., Vaughan, D. (2023). *tidyr: Tidy Messy Data*. R package version 2.0.0, <https://CRAN.R-project.org/package=tidyr>

Examples

```
# See vignette("pivot") for examples and explanation
library(dplyr)
tidySummarizedExperiment::pasilla %>%
  pivot_longer(c(condition, type),
    names_to="name", values_to="value")
```

<code>pivot_wider</code>	<i>Pivot data from long to wide</i>
--------------------------	-------------------------------------

Description

`pivot_wider()` "widens" data, increasing the number of columns and decreasing the number of rows. The inverse transformation is `pivot_longer()`.
Learn more in `vignette("pivot")`.

Usage

```
## S3 method for class 'SummarizedExperiment'
pivot_wider(
  data,
  ...,
  id_cols = NULL,
  id_expand = FALSE,
  names_from = name,
  names_prefix = "",
  names_sep = "_",
  names_glue = NULL,
  names_sort = FALSE,
  names_vary = "fastest",
  names_expand = FALSE,
  names_repair = "check_unique",
```

```

    values_from = value,
    values_fill = NULL,
    values_fn = NULL,
    unused_fn = NULL
  )

```

Arguments

data	A data frame to pivot.
...	Additional arguments passed on to methods.
id_cols	<tidy-select> A set of columns that uniquely identify each observation. Typically used when you have redundant variables, i.e. variables whose values are perfectly correlated with existing variables. Defaults to all columns in data except for the columns specified through <code>names_from</code> and <code>values_from</code>. If a tidyselect expression is supplied, it will be evaluated on data after removing the columns specified through <code>names_from</code> and <code>values_from</code>.
id_expand	Should the values in the <code>id_cols</code> columns be expanded by <code>expand()</code> before pivoting? This results in more rows, the output will contain a complete expansion of all possible values in <code>id_cols</code> . Implicit factor levels that aren't represented in the data will become explicit. Additionally, the row values corresponding to the expanded <code>id_cols</code> will be sorted.
names_from, values_from	<tidy-select> A pair of arguments describing which column (or columns) to get the name of the output column (<code>names_from</code>), and which column (or columns) to get the cell values from (<code>values_from</code>). If <code>values_from</code> contains multiple values, the value will be added to the front of the output column.
names_prefix	String added to the start of every variable name. This is particularly useful if <code>names_from</code> is a numeric vector and you want to create syntactic variable names.
names_sep	If <code>names_from</code> or <code>values_from</code> contains multiple variables, this will be used to join their values together into a single string to use as a column name.
names_glue	Instead of <code>names_sep</code> and <code>names_prefix</code> , you can supply a glue specification that uses the <code>names_from</code> columns (and special <code>.value</code>) to create custom column names.
names_sort	Should the column names be sorted? If <code>FALSE</code> , the default, column names are ordered by first appearance.
names_vary	When <code>names_from</code> identifies a column (or columns) with multiple unique values, and multiple <code>values_from</code> columns are provided, in what order should the resulting column names be combined? • "fastest" varies <code>names_from</code> values fastest, resulting in a column naming scheme of the form: <code>value1_name1</code>, <code>value1_name2</code>, <code>value2_name1</code>, <code>value2_name2</code>. This is the default. • "slowest" varies <code>names_from</code> values slowest, resulting in a column naming scheme of the form: <code>value1_name1</code>, <code>value2_name1</code>, <code>value1_name2</code>, <code>value2_name2</code>.
names_expand	Should the values in the <code>names_from</code> columns be expanded by <code>expand()</code> before pivoting? This results in more columns, the output will contain column names corresponding to a complete expansion of all possible values in <code>names_from</code> . Implicit factor levels that aren't represented in the data will become explicit.

	Additionally, the column names will be sorted, identical to what <code>names_sort</code> would produce.
<code>names_repair</code>	What happens if the output has invalid column names? The default, "check_unique" is to error if the columns are duplicated. Use "minimal" to allow duplicates in the output, or "unique" to de-duplicated by adding numeric suffixes. See vctrs::vec_as_names() for more options.
<code>values_fill</code>	Optionally, a (scalar) value that specifies what each value should be filled in with when missing. This can be a named list if you want to apply different fill values to different value columns.
<code>values_fn</code>	Optionally, a function applied to the value in each cell in the output. You will typically use this when the combination of <code>id_cols</code> and <code>names_from</code> columns does not uniquely identify an observation. This can be a named list if you want to apply different aggregations to different <code>values_from</code> columns.
<code>unused_fn</code>	Optionally, a function applied to summarize the values from the unused columns (i.e. columns not identified by <code>id_cols</code> , <code>names_from</code> , or <code>values_from</code>). The default drops all unused columns from the result. This can be a named list if you want to apply different aggregations to different unused columns. <code>id_cols</code> must be supplied for <code>unused_fn</code> to be useful, since otherwise all unspecified columns will be considered <code>id_cols</code> . This is similar to grouping by the <code>id_cols</code> then summarizing the unused columns using <code>unused_fn</code> .

Details

`pivot_wider()` is an updated approach to [spread\(\)](#), designed to be both simpler to use and to handle more use cases. We recommend you use `pivot_wider()` for new code; `spread()` isn't going away but is no longer under active development.

Value

`tidySummarizedExperiment`

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. *Nat Methods* 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., Vaughan, D. (2023). *tidyr: Tidy Messy Data*. R package version 2.0.0, <https://CRAN.R-project.org/package=tidyr>

See Also

[pivot_wider_spec\(\)](#) to pivot "by hand" with a data frame that defines a pivoting specification.

Examples

```
# See vignette("pivot") for examples and explanation
library(dplyr)
tidySummarizedExperiment::pasilla %>%
  pivot_wider(names_from=feature, values_from=counts)
```

plot_ly

Initiate a plotly visualization

Description

This function maps R objects to [plotly.js](#), an (MIT licensed) web-based interactive charting library. It provides abstractions for doing common things (e.g. mapping data values to fill colors (via `color`) or creating [animations](#) (via `frame`)) and sets some different defaults to make the interface feel more 'R-like' (i.e., closer to `plot()` and `ggplot2::qplot()`).

Usage

```
## S3 method for class 'tbl_df'
plot_ly(
  data = data.frame(),
  ...,
  type = NULL,
  name = NULL,
  color = NULL,
  colors = NULL,
  alpha = NULL,
  stroke = NULL,
  strokes = NULL,
  alpha_stroke = 1,
  size = NULL,
  sizes = c(10, 100),
  span = NULL,
  spans = c(1, 20),
  symbol = NULL,
  symbols = NULL,
  linetype = NULL,
  linetypes = NULL,
  split = NULL,
  frame = NULL,
  width = NULL,
  height = NULL,
  source = "A"
)

## S3 method for class 'SummarizedExperiment'
plot_ly(
  data = data.frame(),
  ...,
  type = NULL,
```

```

name = NULL,
color = NULL,
colors = NULL,
alpha = NULL,
stroke = NULL,
strokes = NULL,
alpha_stroke = 1,
size = NULL,
sizes = c(10, 100),
span = NULL,
spans = c(1, 20),
symbol = NULL,
symbols = NULL,
linetype = NULL,
linetypes = NULL,
split = NULL,
frame = NULL,
width = NULL,
height = NULL,
source = "A"
)

```

Arguments

data	A data frame (optional) or crosstalk::SharedData object.
...	Arguments (i.e., attributes) passed along to the trace type. See schema() for a list of acceptable attributes for a given trace type (by going to traces -> type -> attributes). Note that attributes provided at this level may override other arguments (e.g. <code>plot_ly(x = 1:10, y = 1:10, color = I("red"), marker = list(color = "blue"))</code>).
type	A character string specifying the trace type (e.g. "scatter", "bar", "box", etc). If specified, it <i>always</i> creates a trace, otherwise
name	Values mapped to the trace's name attribute. Since a trace can only have one name, this argument acts very much like <code>split</code> in that it creates one trace for every unique value.
color	Values mapped to relevant 'fill-color' attribute(s) (e.g. fillcolor , marker.color , textfont.color , etc.). The mapping from data values to color codes may be controlled using <code>colors</code> and <code>alpha</code> , or avoided altogether via I() (e.g., <code>color = I("red")</code>). Any color understood by grDevices::col2rgb() may be used in this way.
colors	Either a colorbrewer2.org palette name (e.g. "YlOrRd" or "Blues"), or a vector of colors to interpolate in hexadecimal "#RRGGBB" format, or a color interpolation function like <code>colorRamp()</code> .
alpha	A number between 0 and 1 specifying the alpha channel applied to color. Defaults to 0.5 when mapping to fillcolor and 1 otherwise.
stroke	Similar to <code>color</code> , but values are mapped to relevant 'stroke-color' attribute(s) (e.g., marker.line.color and line.color for filled polygons). If not specified, stroke inherits from <code>color</code> .
strokes	Similar to <code>colors</code> , but controls the stroke mapping.
alpha_stroke	Similar to <code>alpha</code> , but applied to stroke.

size	(Numeric) values mapped to relevant 'fill-size' attribute(s) (e.g., <code>marker.size</code> , <code>textfont.size</code> , and <code>error_x.width</code>). The mapping from data values to symbols may be controlled using <code>sizes</code> , or avoided altogether via <code>I()</code> (e.g., <code>size = I(30)</code>).
sizes	A numeric vector of length 2 used to scale <code>size</code> to pixels.
span	(Numeric) values mapped to relevant 'stroke-size' attribute(s) (e.g., <code>marker.line.width</code> , <code>line.width</code> for filled polygons, and <code>error_x.thickness</code>). The mapping from data values to symbols may be controlled using <code>spans</code> , or avoided altogether via <code>I()</code> (e.g., <code>span = I(30)</code>).
spans	A numeric vector of length 2 used to scale <code>span</code> to pixels.
symbol	(Discrete) values mapped to <code>marker.symbol</code> . The mapping from data values to symbols may be controlled using <code>symbols</code> , or avoided altogether via <code>I()</code> (e.g., <code>symbol = I("pentagon")</code>). Any <code>pch</code> value or <code>symbol name</code> may be used in this way.
symbols	A character vector of <code>pch</code> values or <code>symbol names</code> .
linetype	(Discrete) values mapped to <code>line.dash</code> . The mapping from data values to symbols may be controlled using <code>linetypes</code> , or avoided altogether via <code>I()</code> (e.g., <code>linetype = I("dash")</code>). Any <code>lty</code> (see <code>par</code>) value or <code>dash name</code> may be used in this way.
linetypes	A character vector of <code>lty</code> values or <code>dash names</code> .
split	(Discrete) values used to create multiple traces (one trace per value).
frame	(Discrete) values used to create animation frames.
width	Width in pixels (optional, defaults to automatic sizing).
height	Height in pixels (optional, defaults to automatic sizing).
source	a character string of length 1. Match the value of this string with the source argument in <code>event_data()</code> to retrieve the event data corresponding to a specific plot (shiny apps can have multiple plots).

Details

Unless `type` is specified, this function just initiates a plotly object with 'global' attributes that are passed onto downstream uses of `add_trace()` (or similar). A `formula` must always be used when referencing column name(s) in data (e.g. `plot_ly(mtcars, x = ~wt)`). Formulas are optional when supplying values directly, but they do help inform default axis/scale titles (e.g., `plot_ly(x = mtcars$wt)` vs `plot_ly(x = ~mtcars$wt)`)

Value

plotly
plotly

Author(s)

Carson Sievert

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. *Nat Methods* 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. Nat Methods 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

See Also

- For initializing a plotly-geo object: [plot_geo\(\)](#)
- For initializing a plotly-mapbox object: [plot_mapbox\(\)](#)
- For translating a ggplot2 object to a plotly object: [ggplotly\(\)](#)
- For modifying any plotly object: [layout\(\)](#), [add_trace\(\)](#), [style\(\)](#)
- For linked brushing: [highlight\(\)](#)
- For arranging multiple plots: [subplot\(\)](#), [crosstalk::bscols\(\)](#)
- For inspecting plotly objects: [plotly_json\(\)](#)
- For quick, accurate, and searchable plotly.js reference: [schema\(\)](#)

Examples

```
data(se)
se |>
  plot_ly(x = ~counts)

data(se)
se |>
  plot_ly(x = ~counts)
```

pull	<i>Extract a single column</i>
------	--------------------------------

Description

`pull()` is similar to `$`. It’s mostly useful because it looks a little nicer in pipes, it also works with remote data frames, and it can optionally name the output.

Usage

```
## S3 method for class 'SummarizedExperiment'
pull(.data, var = -1, name = NULL, ...)
```

Arguments

<code>.data</code>	A data frame, data frame extension (e.g. a tibble), or a lazy data frame (e.g. from <code>dbplyr</code> or <code>dtplyr</code>). See <i>Methods</i> , below, for more details.
<code>var</code>	A variable specified as: • a literal variable name • a positive integer, giving the position counting from the left • a negative integer, giving the position counting from the right.

	The default returns the last column (on the assumption that's the column you've created most recently).
	This argument is taken by expression and supports quasiquotation (you can unquote column names and column locations).
name	An optional parameter that specifies the column to be used as names for a named vector. Specified in a similar manner as var.
...	For use by methods.

Value

A vector the same size as .data.

Methods

This function is a **generic**, which means that packages can provide implementations (methods) for other classes. See the documentation of individual methods for extra arguments and differences in behaviour.

The following methods are currently available in loaded packages: no methods found.

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. Nat Methods 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., François, R., Henry, L., Müller, K., Vaughan, D. (2023). dplyr: A Grammar of Data Manipulation. R package version 2.1.4, <https://CRAN.R-project.org/package=dplyr>

Examples

```
data(pasilla)
pasilla |> pull(feature)
```

rename	<i>Rename columns</i>
--------	-----------------------

Description

rename() changes the names of individual variables using new_name = old_name syntax; rename_with() renames columns using a function.

Usage

```
## S3 method for class 'SummarizedExperiment'
rename(.data, ...)
```

Arguments

.data	A data frame, data frame extension (e.g. a tibble), or a lazy data frame (e.g. from dbplyr or dtplyr). See <i>Methods</i> , below, for more details.
...	For rename(): <tidy-select> Use new_name = old_name to rename selected variables. For rename_with(): additional arguments passed onto .fn.

Value

An object of the same type as `.data`. The output has the following properties:

- Rows are not affected.
- Column names are changed; column order is preserved.
- Data frame attributes are preserved.
- Groups are updated to reflect new names.

Methods

This function is a **generic**, which means that packages can provide implementations (methods) for other classes. See the documentation of individual methods for extra arguments and differences in behaviour.

The following methods are currently available in loaded packages: no methods found.

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. Nat Methods 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., François, R., Henry, L., Müller, K., Vaughan, D. (2023). dplyr: A Grammar of Data Manipulation. R package version 2.1.4, <https://CRAN.R-project.org/package=dplyr>

See Also

Other single table verbs: `mutate()`, `slice()`, `summarise()`

Examples

```
data(pasilla)
pasilla |> rename(cond=condition)
```

right_join	<i>Mutating joins</i>
------------	-----------------------

Description

Mutating joins add columns from `y` to `x`, matching observations based on the keys. There are four mutating joins: the inner join, and the three outer joins.

Inner join:

An `inner_join()` only keeps observations from `x` that have a matching key in `y`. The most important property of an inner join is that unmatched rows in either input are not included in the result. This means that generally inner joins are not appropriate in most analyses, because it is too easy to lose observations.

Outer joins:

The three outer joins keep observations that appear in at least one of the data frames:

- A `left_join()` keeps all observations in `x`.
- A `right_join()` keeps all observations in `y`.
- A `full_join()` keeps all observations in `x` and `y`.

Usage

```
## S3 method for class 'SummarizedExperiment'
right_join(x, y, by = NULL, copy = FALSE, suffix = c(".x", ".y"), ...)
```

Arguments

<code>x, y</code>	A pair of data frames, data frame extensions (e.g. a tibble), or lazy data frames (e.g. from dbplyr or dtplyr). See <i>Methods</i> , below, for more details.
<code>by</code>	A join specification created with <code>join_by()</code>, or a character vector of variables to join by. If <code>NULL</code>, the default, <code>*_join()</code> will perform a natural join, using all variables in common across <code>x</code> and <code>y</code>. A message lists the variables so that you can check they're correct; suppress the message by supplying <code>by</code> explicitly. To join on different variables between <code>x</code> and <code>y</code>, use a <code>join_by()</code> specification. For example, <code>join_by(a == b)</code> will match <code>x\$a</code> to <code>y\$b</code>. To join by multiple variables, use a <code>join_by()</code> specification with multiple expressions. For example, <code>join_by(a == b, c == d)</code> will match <code>x\$a</code> to <code>y\$b</code> and <code>x\$c</code> to <code>y\$d</code>. If the column names are the same between <code>x</code> and <code>y</code>, you can shorten this by listing only the variable names, like <code>join_by(a, c)</code>. <code>join_by()</code> can also be used to perform inequality, rolling, and overlap joins. See the documentation at <code>?join_by</code> for details on these types of joins. For simple equality joins, you can alternatively specify a character vector of variable names to join by. For example, <code>by = c("a", "b")</code> joins <code>x\$a</code> to <code>y\$a</code> and <code>x\$b</code> to <code>y\$b</code>. If variable names differ between <code>x</code> and <code>y</code>, use a named character vector like <code>by = c("x_a" = "y_a", "x_b" = "y_b")</code>. To perform a cross-join, generating all combinations of <code>x</code> and <code>y</code>, see <code>cross_join()</code>.
<code>copy</code>	If <code>x</code> and <code>y</code> are not from the same data source, and <code>copy</code> is <code>TRUE</code> , then <code>y</code> will be copied into the same src as <code>x</code> . This allows you to join tables across srcs, but it is a potentially expensive operation so you must opt into it.
<code>suffix</code>	If there are non-joined duplicate variables in <code>x</code> and <code>y</code> , these suffixes will be added to the output to disambiguate them. Should be a character vector of length 2.
<code>...</code>	Other parameters passed onto methods.

Value

An object of the same type as `x` (including the same groups). The order of the rows and columns of `x` is preserved as much as possible. The output has the following properties:

- The rows are affected by the join type.
 - `inner_join()` returns matched `x` rows.
 - `left_join()` returns all `x` rows.
 - `right_join()` returns matched of `x` rows, followed by unmatched `y` rows.
 - `full_join()` returns all `x` rows, followed by unmatched `y` rows.
- Output columns include all columns from `x` and all non-key columns from `y`. If `keep = TRUE`, the key columns from `y` are included as well.
- If non-key columns in `x` and `y` have the same name, suffixes are added to disambiguate. If `keep = TRUE` and key columns in `x` and `y` have the same name, suffixes are added to disambiguate these as well.
- If `keep = FALSE`, output columns included in `by` are coerced to their common type between `x` and `y`.

Many-to-many relationships

By default, dplyr guards against many-to-many relationships in equality joins by throwing a warning. These occur when both of the following are true:

- A row in *x* matches multiple rows in *y*.
- A row in *y* matches multiple rows in *x*.

This is typically surprising, as most joins involve a relationship of one-to-one, one-to-many, or many-to-one, and is often the result of an improperly specified join. Many-to-many relationships are particularly problematic because they can result in a Cartesian explosion of the number of rows returned from the join.

If a many-to-many relationship is expected, silence this warning by explicitly setting `relationship = "many-to-many"`.

In production code, it is best to preemptively set `relationship` to whatever relationship you expect to exist between the keys of *x* and *y*, as this forces an error to occur immediately if the data doesn't align with your expectations.

Inequality joins typically result in many-to-many relationships by nature, so they don't warn on them by default, but you should still take extra care when specifying an inequality join, because they also have the capability to return a large number of rows.

Rolling joins don't warn on many-to-many relationships either, but many rolling joins follow a many-to-one relationship, so it is often useful to set `relationship = "many-to-one"` to enforce this.

Note that in SQL, most database providers won't let you specify a many-to-many relationship between two tables, instead requiring that you create a third *junction table* that results in two one-to-many relationships instead.

Methods

These functions are **generics**, which means that packages can provide implementations (methods) for other classes. See the documentation of individual methods for extra arguments and differences in behaviour.

Methods available in currently loaded packages:

- `inner_join()`: no methods found.
- `left_join()`: no methods found.
- `right_join()`: no methods found.
- `full_join()`: no methods found.

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. *Nat Methods* 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., François, R., Henry, L., Müller, K., Vaughan, D. (2023). dplyr: A Grammar of Data Manipulation. R package version 2.1.4, <https://CRAN.R-project.org/package=dplyr>

See Also

Other joins: [cross_join\(\)](#), [filter-joins](#), [nest_join\(\)](#)

Examples

```
data(pasilla)

tt <- pasilla
tt |> right_join(tt |>
  distinct(condition) |>
  mutate(new_column=1:2) |>
  slice(1))
```

rowwise

Group input by rows

Description

`rowwise()` allows you to compute on a data frame a row-at-a-time. This is most useful when a vectorised function doesn't exist.

Most dplyr verbs preserve row-wise grouping. The exception is `summarise()`, which return a `grouped_df`. You can explicitly ungroup with `ungroup()` or `as_tibble()`, or convert to a `grouped_df` with `group_by()`.

Usage

```
## S3 method for class 'SummarizedExperiment'
rowwise(data, ...)
```

Arguments

<code>data</code>	Input data frame.
<code>...</code>	<code><tidy-select></code> Variables to be preserved when calling <code>summarise()</code> . This is typically a set of variables whose combination uniquely identify each row. NB: unlike <code>group_by()</code> you can not create new variables here but instead you can select multiple variables with (e.g.) <code>everything()</code> .

Value

A row-wise data frame with class `rowwise_df`. Note that a `rowwise_df` is implicitly grouped by row, but is not a `grouped_df`.

List-columns

Because a `rowwise` has exactly one row per group it offers a small convenience for working with list-columns. Normally, `summarise()` and `mutate()` extract a groups worth of data with `[]`. But when you index a list in this way, you get back another list. When you're working with a `rowwise` tibble, then dplyr will use `[[` instead of `[]` to make your life a little easier.

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. *Nat Methods* 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., François, R., Henry, L., Müller, K., Vaughan, D. (2023). dplyr: A Grammar of Data Manipulation. R package version 2.1.4, <https://CRAN.R-project.org/package=dplyr>

See Also

[nest_by\(\)](#) for a convenient way of creating rowwise data frames with nested data.

Examples

```
# TODO
```

sample_n	<i>Sample n rows from a table</i>
----------	-----------------------------------

Description

[Superseded] `sample_n()` and `sample_frac()` have been superseded in favour of [slice_sample\(\)](#). While they will not be deprecated in the near future, retirement means that we will only perform critical bug fixes, so we recommend moving to the newer alternative.

These functions were superseded because we realised it was more convenient to have two mutually exclusive arguments to one function, rather than two separate functions. This also made it to clean up a few other smaller design issues with `sample_n()/sample_frac`:

- The connection to `slice()` was not obvious.
- The name of the first argument, `tbl`, is inconsistent with other single table verbs which use `.data`.
- The size argument uses tidy evaluation, which is surprising and undocumented.
- It was easier to remove the deprecated `.env` argument.
- ... was in a suboptimal position.

Usage

```
## S3 method for class 'SummarizedExperiment'
sample_n(tbl, size, replace = FALSE, weight = NULL, .env = NULL, ...)

## S3 method for class 'SummarizedExperiment'
sample_frac(tbl, size = 1, replace = FALSE, weight = NULL, .env = NULL, ...)
```

Arguments

<code>tbl</code>	A data.frame.
<code>size</code>	<tidy-select> For <code>sample_n()</code> , the number of rows to select. For <code>sample_frac()</code> , the fraction of rows to select. If <code>tbl</code> is grouped, <code>size</code> applies to each group.
<code>replace</code>	Sample with or without replacement?
<code>weight</code>	<tidy-select> Sampling weights. This must evaluate to a vector of non-negative numbers the same length as the input. Weights are automatically standardised to sum to 1.
<code>.env</code>	DEPRECATED.
<code>...</code>	ignored

Value

tidySummarizedExperiment

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. Nat Methods 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., François, R., Henry, L., Müller, K., Vaughan, D. (2023). dplyr: A Grammar of Data Manipulation. R package version 2.1.4, <https://CRAN.R-project.org/package=dplyr>

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. Nat Methods 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., François, R., Henry, L., Müller, K., Vaughan, D. (2023). dplyr: A Grammar of Data Manipulation. R package version 2.1.4, <https://CRAN.R-project.org/package=dplyr>

Examples

```
data(pasilla)
pasilla |> sample_n(50)
pasilla |> sample_frac(0.1)
```

se

Read counts of RNA-seq samples derived from Pasilla knock-down by Brooks et al.

Description

A SummarizedExperiment dataset containing the transcriptome information for *Drosophila Melanogaster*.

Usage

```
data(se)
```

Format

containing 14599 features and 7 biological replicates.

Source

<https://bioconductor.org/packages/release/data/experiment/html/pasilla.html>

select

*Keep or drop columns using their names and types***Description**

Select (and optionally rename) variables in a data frame, using a concise mini-language that makes it easy to refer to variables based on their name (e.g. `a:f` selects all columns from `a` on the left to `f` on the right) or type (e.g. `where(is.numeric)` selects all numeric columns).

Overview of selection features:

Tidyverse selections implement a dialect of R where operators make it easy to select variables:

- `:` for selecting a range of consecutive variables.
- `!` for taking the complement of a set of variables.
- `&` and `|` for selecting the intersection or the union of two sets of variables.
- `c()` for combining selections.

In addition, you can use **selection helpers**. Some helpers select specific columns:

- `everything()`: Matches all variables.
- `last_col()`: Select last variable, possibly with an offset.
- `group_cols()`: Select all grouping columns.

Other helpers select variables by matching patterns in their names:

- `starts_with()`: Starts with a prefix.
- `ends_with()`: Ends with a suffix.
- `contains()`: Contains a literal string.
- `matches()`: Matches a regular expression.
- `num_range()`: Matches a numerical range like `x01`, `x02`, `x03`.

Or from variables stored in a character vector:

- `all_of()`: Matches variable names in a character vector. All names must be present, otherwise an out-of-bounds error is thrown.
- `any_of()`: Same as `all_of()`, except that no error is thrown for names that don't exist.

Or using a predicate function:

- `where()`: Applies a function to all variables and selects those for which the function returns `TRUE`.

Usage

```
## S3 method for class 'SummarizedExperiment'
select(.data, ...)
```

Arguments

<code>.data</code>	A data frame, data frame extension (e.g. a tibble), or a lazy data frame (e.g. from <code>dbplyr</code> or <code>dtplyr</code>). See <i>Methods</i> , below, for more details.
<code>...</code>	<tidy-select> One or more unquoted expressions separated by commas. Variable names can be used as if they were positions in the data frame, so expressions like <code>x:y</code> can be used to select a range of variables.

Value

An object of the same type as `.data`. The output has the following properties:

- Rows are not affected.
- Output columns are a subset of input columns, potentially with a different order. Columns will be renamed if `new_name = old_name` form is used.
- Data frame attributes are preserved.
- Groups are maintained; you can't select off grouping variables.

Methods

This function is a **generic**, which means that packages can provide implementations (methods) for other classes. See the documentation of individual methods for extra arguments and differences in behaviour.

The following methods are currently available in loaded packages: no methods found.

Examples

Here we show the usage for the basic selection operators. See the specific help pages to learn about helpers like `starts_with()`.

The selection language can be used in functions like `dplyr::select()` or `tidyr::pivot_longer()`. Let's first attach the tidyverse:

```
library(tidyverse)
```

```
# For better printing
iris <- as_tibble(iris)
```

Select variables by name:

```
starwars %>% select(height)
#> # A tibble: 87 x 1
#>   height
#>   <int>
#> 1    172
#> 2    167
#> 3     96
#> 4    202
#> # i 83 more rows
```

```
iris %>% pivot_longer(Sepal.Length)
#> # A tibble: 150 x 6
#>   Sepal.Width Petal.Length Petal.Width Species name      value
#>   <dbl>         <dbl>         <dbl> <fct>   <chr>    <dbl>
#> 1      3.5         1.4           0.2 setosa Sepal.Length  5.1
#> 2      3          1.4           0.2 setosa Sepal.Length  4.9
#> 3      3.2         1.3           0.2 setosa Sepal.Length  4.7
#> 4      3.1         1.5           0.2 setosa Sepal.Length  4.6
#> # i 146 more rows
```

Select multiple variables by separating them with commas. Note how the order of columns is determined by the order of inputs:

```
starwars %>% select(homeworld, height, mass)
#> # A tibble: 87 x 3
#>   homeworld height  mass
#>   <chr>      <int> <dbl>
#> 1 Tatooine    172    77
#> 2 Tatooine    167    75
#> 3 Naboo       96     32
#> 4 Tatooine    202   136
#> # i 83 more rows
```

Functions like `tidyr::pivot_longer()` don't take variables with dots. In this case use `c()` to select multiple variables:

```
iris %>% pivot_longer(c(Sepal.Length, Petal.Length))
#> # A tibble: 300 x 5
#>   Sepal.Width Petal.Width Species name          value
#>   <dbl>      <dbl> <fct> <chr>      <dbl>
#> 1         3.5         0.2 setosa Sepal.Length  5.1
#> 2         3.5         0.2 setosa Petal.Length  1.4
#> 3          3         0.2 setosa Sepal.Length  4.9
#> 4          3         0.2 setosa Petal.Length  1.4
#> # i 296 more rows
```

Operators::

The `:` operator selects a range of consecutive variables:

```
starwars %>% select(name:mass)
#> # A tibble: 87 x 3
#>   name          height  mass
#>   <chr>      <int> <dbl>
#> 1 Luke Skywalker    172    77
#> 2 C-3PO             167    75
#> 3 R2-D2              96    32
#> 4 Darth Vader      202   136
#> # i 83 more rows
```

The `!` operator negates a selection:

```
starwars %>% select(!(name:mass))
#> # A tibble: 87 x 11
#>   hair_color skin_color eye_color birth_year sex  gender homeworld species
#>   <chr>      <chr>      <chr>      <dbl> <chr> <chr>   <chr>   <chr>
#> 1 blond     fair        blue         19  male  masculine Tatooine Human
#> 2 <NA>      gold        yellow       112  none  masculine Tatooine Droid
#> 3 <NA>      white, blue red         33  none  masculine Naboo   Droid
#> 4 none      white        yellow      41.9  male  masculine Tatooine Human
#> # i 83 more rows
#> # i 3 more variables: films <list>, vehicles <list>, starships <list>
```

```
iris %>% select(!c(Sepal.Length, Petal.Length))
#> # A tibble: 150 x 3
#>   Sepal.Width Petal.Width Species
#>   <dbl>      <dbl> <fct>
#> 1         3.5         0.2 setosa
```

```
#> 2          3          0.2 setosa
#> 3          3.2          0.2 setosa
#> 4          3.1          0.2 setosa
#> # i 146 more rows
```

```
iris %>% select(!ends_with("Width"))
#> # A tibble: 150 x 3
#>   Sepal.Length Petal.Length Species
#>   <dbl>         <dbl> <fct>
#> 1         5.1         1.4 setosa
#> 2         4.9         1.4 setosa
#> 3         4.7         1.3 setosa
#> 4         4.6         1.5 setosa
#> # i 146 more rows
```

& and | take the intersection or the union of two selections:

```
iris %>% select(starts_with("Petal") & ends_with("Width"))
#> # A tibble: 150 x 1
#>   Petal.Width
#>   <dbl>
#> 1     0.2
#> 2     0.2
#> 3     0.2
#> 4     0.2
#> # i 146 more rows
```

```
iris %>% select(starts_with("Petal") | ends_with("Width"))
#> # A tibble: 150 x 3
#>   Petal.Length Petal.Width Sepal.Width
#>   <dbl>         <dbl>         <dbl>
#> 1         1.4         0.2         3.5
#> 2         1.4         0.2         3
#> 3         1.3         0.2         3.2
#> 4         1.5         0.2         3.1
#> # i 146 more rows
```

To take the difference between two selections, combine the & and ! operators:

```
iris %>% select(starts_with("Petal") & !ends_with("Width"))
#> # A tibble: 150 x 1
#>   Petal.Length
#>   <dbl>
#> 1         1.4
#> 2         1.4
#> 3         1.3
#> 4         1.5
#> # i 146 more rows
```

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. *Nat Methods* 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., François, R., Henry, L., Müller, K., Vaughan, D. (2023). dplyr: A Grammar of Data Manipulation. R package version 2.1.4, <https://CRAN.R-project.org/package=dplyr>

See Also

Other single table verbs: [arrange\(\)](#), [filter\(\)](#), [mutate\(\)](#), [reframe\(\)](#), [rename\(\)](#), [slice\(\)](#), [summarise\(\)](#)

Examples

```
data(pasilla)
pasilla |> select(.sample, .feature, counts)
```

separate	<i>Separate a character column into multiple columns with a regular expression or numeric locations</i>
----------	---

Description

[Superseded]

`separate()` has been superseded in favour of [separate_wider_position\(\)](#) and [separate_wider_delim\(\)](#) because the two functions make the two uses more obvious, the API is more polished, and the handling of problems is better. Superseded functions will not go away, but will only receive critical bug fixes.

Given either a regular expression or a vector of character positions, `separate()` turns a single character column into multiple columns.

Usage

```
## S3 method for class 'SummarizedExperiment'
separate(
  data,
  col,
  into,
  sep = "[^[:alnum:]]+",
  remove = TRUE,
  convert = FALSE,
  extra = "warn",
  fill = "warn",
  ...
)
```

Arguments

<code>data</code>	A data frame.
<code>col</code>	<tidy-select> Column to expand.
<code>into</code>	Names of new variables to create as character vector. Use NA to omit the variable in the output.

sep	Separator between columns. If character, sep is interpreted as a regular expression. The default value is a regular expression that matches any sequence of non-alphanumeric values. If numeric, sep is interpreted as character positions to split at. Positive values start at 1 at the far-left of the string; negative value start at -1 at the far-right of the string. The length of sep should be one less than into.
remove	If TRUE, remove input column from output data frame.
convert	If TRUE, will run <code>type.convert()</code> with <code>as.is = TRUE</code> on new columns. This is useful if the component columns are integer, numeric or logical. NB: this will cause string "NA"s to be converted to NAs.
extra	If sep is a character vector, this controls what happens when there are too many pieces. There are three valid options: • "warn" (the default): emit a warning and drop extra values. • "drop": drop any extra values without a warning. • "merge": only splits at most length(into) times
fill	If sep is a character vector, this controls what happens when there are not enough pieces. There are three valid options: • "warn" (the default): emit a warning and fill from the right • "right": fill with missing values on the right • "left": fill with missing values on the left
...	Additional arguments passed on to methods.

Value

tidySummarizedExperiment

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. *Nat Methods* 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., Vaughan, D. (2023). *tidyr: Tidy Messy Data*. R package version 2.0.0, <https://CRAN.R-project.org/package=tidyr>

See Also

`unite()`, the complement, `extract()` which uses regular expression capturing groups.

Examples

```
un <- tidySummarizedExperiment::pasilla |>
  unite("group", c(condition, type))
un |> separate(col=group, into=c("condition", "type"))
```

 slice

Subset rows using their positions

Description

`slice()` lets you index rows by their (integer) locations. It allows you to select, remove, and duplicate rows. It is accompanied by a number of helpers for common use cases:

- `slice_head()` and `slice_tail()` select the first or last rows.
- `slice_sample()` randomly selects rows.
- `slice_min()` and `slice_max()` select rows with the smallest or largest values of a variable.

If `.data` is a [grouped_df](#), the operation will be performed on each group, so that (e.g.) `slice_head(df, n = 5)` will select the first five rows in each group.

Usage

```
## S3 method for class 'SummarizedExperiment'
slice(.data, ..., .by = NULL, .preserve = FALSE)
```

Arguments

<code>.data</code>	A data frame, data frame extension (e.g. a tibble), or a lazy data frame (e.g. from <code>dbplyr</code> or <code>dtplyr</code>). See <i>Methods</i> , below, for more details.
<code>...</code>	For <code>slice()</code> : <data-masking> Integer row values. Provide either positive values to keep, or negative values to drop. The values provided must be either all positive or all negative. Indices beyond the number of rows in the input are silently ignored. For <code>slice_*()</code> , these arguments are passed on to methods.
<code>.by</code>	[Experimental] <tidy-select> Optionally, a selection of columns to group by for just this operation, functioning as an alternative to <code>group_by()</code> . For details and examples, see ?dplyr_by .
<code>.preserve</code>	Relevant when the <code>.data</code> input is grouped. If <code>.preserve = FALSE</code> (the default), the grouping structure is recalculated based on the resulting data, otherwise the grouping is kept as is.

Details

Slice does not work with relational databases because they have no intrinsic notion of row order. If you want to perform the equivalent operation, use [filter\(\)](#) and [row_number\(\)](#).

Value

An object of the same type as `.data`. The output has the following properties:

- Each row may appear 0, 1, or many times in the output.
- Columns are not modified.
- Groups are not modified.
- Data frame attributes are preserved.

Methods

These function are **generics**, which means that packages can provide implementations (methods) for other classes. See the documentation of individual methods for extra arguments and differences in behaviour.

Methods available in currently loaded packages:

- `slice()`: no methods found.
- `slice_head()`: no methods found.
- `slice_tail()`: no methods found.
- `slice_min()`: no methods found.
- `slice_max()`: no methods found.
- `slice_sample()`: no methods found.

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. Nat Methods 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., François, R., Henry, L., Müller, K., Vaughan, D. (2023). dplyr: A Grammar of Data Manipulation. R package version 2.1.4, <https://CRAN.R-project.org/package=dplyr>

See Also

Other single table verbs: [mutate\(\)](#), [rename\(\)](#), [summarise\(\)](#)

Examples

```
data(pasilla)
pasilla |> slice(1)
```

summarise	<i>Summarise each group down to one row</i>
-----------	---

Description

`summarise()` creates a new data frame. It returns one row for each combination of grouping variables; if there are no grouping variables, the output will have a single row summarising all observations in the input. It will contain one column for each grouping variable and one column for each of the summary statistics that you have specified.

`summarise()` and `summarize()` are synonyms.

Usage

```
## S3 method for class 'SummarizedExperiment'
summarise(.data, ...)

## S3 method for class 'SummarizedExperiment'
summarize(.data, ...)
```

Arguments

- `.data` A data frame, data frame extension (e.g. a tibble), or a lazy data frame (e.g. from `dbplyr` or `dtplyr`). See *Methods*, below, for more details.
- `...` `<data-masking>` Name-value pairs of summary functions. The name will be the name of the variable in the result.
The value can be:
- A vector of length 1, e.g. `min(x)`, `n()`, or `sum(is.na(y))`.
 - A data frame, to add multiple columns from a single expression.
- [Deprecated]** Returning values with size 0 or >1 was deprecated as of 1.1.0. Please use `reframe()` for this instead.

Value

An object *usually* of the same type as `.data`.

- The rows come from the underlying `group_keys()`.
- The columns are a combination of the grouping keys and the summary expressions that you provide.
- The grouping structure is controlled by the `.groups=` argument, the output may be another `grouped_df`, a `tibble` or a `rowwise` data frame.
- Data frame attributes are **not** preserved, because `summarise()` fundamentally creates a new data frame.

Useful functions

- Center: `mean()`, `median()`
- Spread: `sd()`, `IQR()`, `mad()`
- Range: `min()`, `max()`,
- Position: `first()`, `last()`, `nth()`,
- Count: `n()`, `n_distinct()`
- Logical: `any()`, `all()`

Backend variations

The data frame backend supports creating a variable and using it in the same summary. This means that previously created summary variables can be further transformed or combined within the summary, as in `mutate()`. However, it also means that summary variables with the same names as previous variables overwrite them, making those variables unavailable to later summary variables.

This behaviour may not be supported in other backends. To avoid unexpected results, consider using new names for your summary variables, especially when creating multiple summaries.

Methods

This function is a **generic**, which means that packages can provide implementations (methods) for other classes. See the documentation of individual methods for extra arguments and differences in behaviour.

The following methods are currently available in loaded packages: no methods found.

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. Nat Methods 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., François, R., Henry, L., Müller, K., Vaughan, D. (2023). dplyr: A Grammar of Data Manipulation. R package version 2.1.4, <https://CRAN.R-project.org/package=dplyr>

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. Nat Methods 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., François, R., Henry, L., Müller, K., Vaughan, D. (2023). dplyr: A Grammar of Data Manipulation. R package version 2.1.4, <https://CRAN.R-project.org/package=dplyr>

See Also

Other single table verbs: [mutate\(\)](#), [rename\(\)](#), [slice\(\)](#)

Examples

```
data(pasilla)
pasilla |> summarise(mean(counts))
```

tidy	<i>tidy for Seurat</i>
------	------------------------

Description

tidy for Seurat

Usage

```
tidy_(x, ...)

## S3 method for class 'SummarizedExperiment'
tidy(x, ...)

## S3 method for class 'RangedSummarizedExperiment'
tidy(x, ...)
```

Arguments

- x A SummarizedExperiment object
- ... Additional arguments passed to methods

Value

A tidyseurat object.

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. Nat Methods 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. Nat Methods 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. Nat Methods 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Examples

```
data(pasilla)
pasilla %>% tidy()
```

unite	<i>Unite multiple columns into one by pasting strings together</i>
-------	--

Description

Convenience function to paste together multiple columns into one.

Usage

```
## S3 method for class 'SummarizedExperiment'
unite(data, col, ..., sep = "_", remove = TRUE, na.rm = FALSE)
```

Arguments

data	A data frame.
col	The name of the new column, as a string or symbol. This argument is passed by expression and supports quasiquotation (you can unquote strings and symbols). The name is captured from the expression with rlang::ensym() (note that this kind of interface where symbols do not represent actual objects is now discouraged in the tidyverse; we support it here for backward compatibility).
...	<tidy-select> Columns to unite
sep	Separator to use between values.
remove	If TRUE, remove input columns from output data frame.
na.rm	If TRUE, missing values will be removed prior to uniting each value.

Value

tidySummarizedExperiment

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. *Nat Methods* 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., Vaughan, D. (2023). *tidyr: Tidy Messy Data*. R package version 2.0.0, <https://CRAN.R-project.org/package=tidyr>

See Also

[separate\(\)](#), the complement.

Examples

```
tidySummarizedExperiment::pasilla |>
  unite("group", c(condition, type))
```

unnest

Unnest a list-column of data frames into rows and columns

Description

Unnest expands a list-column containing data frames into rows and columns.

Usage

```
## S3 method for class 'tidySummarizedExperiment_nested'
unnest(
  data,
  cols,
  ...,
  keep_empty = FALSE,
  ptype = NULL,
  names_sep = NULL,
  names_repair = "check_unique",
  .drop,
  .id,
  .sep,
  .preserve
)

unnest_summarized_experiment(
  data,
  cols,
  ...,
  keep_empty = FALSE,
  ptype = NULL,
  names_sep = NULL,
  names_repair = "check_unique",
  .drop,
  .id,
```

```
.sep,
.preserve
)
```

Arguments

data	A data frame.
cols	<code><tidy-select></code> List-columns to unnest. When selecting multiple columns, values from the same row will be recycled to their common size.
...	[Deprecated]: previously you could write <code>df %>% unnest(x, y, z)</code> . Convert to <code>df %>% unnest(c(x, y, z))</code> . If you previously created a new variable in <code>unnest()</code> you'll now need to do it explicitly with <code>mutate()</code> . Convert <code>df %>% unnest(y = fun(x, y, z))</code> to <code>df %>% mutate(y = fun(x, y, z)) %>% unnest(y)</code> .
keep_empty	By default, you get one row of output for each element of the list that you are unchopping/unnesting. This means that if there's a size-0 element (like NULL or an empty data frame or vector), then that entire row will be dropped from the output. If you want to preserve all rows, use <code>keep_empty = TRUE</code> to replace size-0 elements with a single row of missing values.
ptype	Optionally, a named list of column name-prototype pairs to coerce <code>cols</code> to, overriding the default that will be guessed from combining the individual values. Alternatively, a single empty <code>ptype</code> can be supplied, which will be applied to all <code>cols</code> .
names_sep	If NULL, the default, the outer names will come from the inner names. If a string, the outer names will be formed by pasting together the outer and the inner column names, separated by <code>names_sep</code> .
names_repair	Used to check that output data frame has valid names. Must be one of the following options: "minimal": no name repair or checks, beyond basic existence, "unique": make sure names are unique and not empty, "check_unique": (the default), no name repair, but check they are unique, "universal": make the names unique and syntactic - a function: apply custom name repair. <code>tidyr_legacy</code>: use the name repair from tidyr 0.8. - a formula: a purrr-style anonymous function (see <code>rlang::as_function()</code>) See <code>vctrs::vec_as_names()</code> for more details on these terms and the strategies used to enforce them.
.drop, .preserve	[Deprecated]: all list-columns are now preserved; If there are any that you don't want in the output use <code>select()</code> to remove them prior to unnesting.
.id	[Deprecated]: convert <code>df %>% unnest(x, .id = "id")</code> to <code>df %>% mutate(id = names(x)) %>% unnest(x)</code>
.sep	[Deprecated]: use <code>names_sep</code> instead.

Value

tidySummarizedExperiment

New syntax

tidyr 1.0.0 introduced a new syntax for `nest()` and `unnest()` that's designed to be more similar to other functions. Converting to the new syntax should be straightforward (guided by the message you'll receive) but if you just need to run an old analysis, you can easily revert to the previous behaviour using `nest_legacy()` and `unnest_legacy()` as follows:

```
library(tidyr)
nest <- nest_legacy
unnest <- unnest_legacy
```

References

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. *Nat Methods* 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., Vaughan, D. (2023). tidyr: Tidy Messy Data. R package version 2.0.0, <https://CRAN.R-project.org/package=tidyr>

Hutchison, W.J., Keyes, T.J., The tidyomics Consortium. et al. The tidyomics ecosystem: enhancing omic data analyses. *Nat Methods* 21, 1166–1170 (2024). <https://doi.org/10.1038/s41592-024-02299-2>

Wickham, H., Vaughan, D. (2023). tidyr: Tidy Messy Data. R package version 2.0.0, <https://CRAN.R-project.org/package=tidyr>

See Also

Other rectangling: `hoist()`, `unnest_longer()`, `unnest_wider()`

Examples

```
tidySummarizedExperiment::pasilla |>
  nest(data=-condition) |>
  unnest(data)

tidySummarizedExperiment::pasilla |>
  nest(data=-condition) |>
  unnest_summarized_experiment(data)
```

%>%

Pipe operator

Description

See `magrittr::%>%` for details.

Usage

```
lhs %>% rhs
```

Arguments

lhs	A value or the magrittr placeholder.
rhs	A function call using the magrittr semantics.

Value

The result of calling rhs(lhs).

Examples

```
library(magrittr)
1 %>% sum(2)
```

Index

- * **datasets**
 - pasilla, [29](#)
 - se, [45](#)
- * **internal**
 - %>%, [59](#)
 - tidySummarizedExperiment-package, [3](#)
- * **single table verbs**
 - mutate, [26](#)
 - rename, [39](#)
 - slice, [52](#)
 - summarise, [53](#)
- +, [26](#)
- .onLoad(), [5](#)
- ==, [12](#)
- >, [12](#)
- >=, [12](#)
- ?dplyr_by, [52](#)
- ?join_by, [14](#), [21](#), [24](#), [41](#)
- &, [12](#)
- %>%, [59](#), [59](#)
- add_trace(), [37](#), [38](#)
- all(), [54](#)
- all_of(), [46](#)
- animation, [35](#)
- any(), [54](#)
- any_of(), [46](#)
- append_samples, [4](#)
- arrange, [3](#), [13](#), [50](#)
- arrange(), [18](#)
- as_tibble, [3](#), [4](#)
- as_tibble(), [43](#)
- base::as.data.frame(), [4](#)
- base::data.frame(), [4](#)
- base::split(), [19](#)
- between(), [12](#)
- bind_cols(bind_rows), [6](#)
- bind_rows, [6](#)
- case_when(), [26](#)
- coalesce(), [26](#)
- contains(), [46](#)
- count, [8](#)
- cross_join, [15](#), [22](#), [25](#), [42](#)
- cross_join(), [14](#), [21](#), [24](#), [41](#)
- crosstalk::bscols(), [38](#)
- crosstalk::SharedData, [36](#)
- cumall(), [26](#)
- cumany(), [26](#)
- cume_dist(), [26](#)
- cummax(), [26](#)
- cummean(), [26](#)
- cummin(), [26](#)
- cumsum(), [26](#)
- data.frame, [4](#)
- dense_rank(), [26](#)
- distinct, [9](#)
- dplyr::group_by(), [29](#)
- ends_with(), [46](#)
- enframe(), [6](#)
- event_data(), [37](#)
- everything(), [46](#)
- expand(), [33](#)
- extract, [10](#)
- extract(), [31](#), [51](#)
- filter, [3](#), [11](#), [50](#)
- filter(), [52](#)
- first(), [54](#)
- formula, [37](#)
- fortify(), [16](#)
- full_join, [3](#), [13](#)
- gather(), [32](#)
- ggplot, [3](#), [16](#)
- ggplot2::qplot(), [35](#)
- ggplotly(), [38](#)
- grDevices::col2rgb(), [36](#)
- group_by, [17](#), [20](#)
- group_by(), [8](#), [12](#), [19](#), [43](#), [52](#)
- group_by_drop_default(), [18](#)
- group_cols(), [46](#)
- group_keys(), [19](#), [54](#)
- group_map, [19](#), [20](#)

group_nest, [19, 20](#)
 group_split, [19, 19](#)
 group_split(), [19](#)
 group_trim, [19, 20](#)
 grouped_df, [18, 43, 52, 54](#)

 highlight(), [38](#)
 hoist, [59](#)

 I(), [36, 37](#)
 if_else(), [26](#)
 inner_join, [3, 20](#)
 IQR(), [54](#)
 is.na(), [12](#)

 join_by(), [14, 21, 23, 24, 41](#)

 lag(), [26](#)
 last(), [54](#)
 last_col(), [46](#)
 layout(), [38](#)
 lead(), [26](#)
 left_join, [3, 23](#)
 list_of, [19](#)
 log(), [26](#)

 mad(), [54](#)
 matches(), [46](#)
 matrix, [4](#)
 max(), [54](#)
 mean(), [54](#)
 median(), [54](#)
 min(), [54](#)
 min_rank(), [26](#)
 mutate, [3, 13, 26, 40, 50, 53, 55](#)
 mutate(), [54](#)

 n(), [54](#)
 n_distinct(), [54](#)
 na_if(), [26](#)
 near(), [12](#)
 nest, [3, 28](#)
 nest_by(), [44](#)
 nest_join, [15, 22, 25, 42](#)
 nest_legacy(), [28, 59](#)
 nth(), [54](#)
 ntile(), [26](#)
 num_range(), [46](#)

 par, [37](#)
 pasilla, [29](#)
 pch, [37](#)
 percent_rank(), [26](#)
 pivot_longer, [3, 30](#)
 pivot_longer(), [32](#)
 pivot_wider, [3, 32](#)
 pivot_wider(), [30](#)
 pivot_wider_spec(), [34](#)
 plot(), [35](#)
 plot_geo(), [38](#)
 plot_ly, [35](#)
 plot_mapbox(), [38](#)
 plotly_json(), [38](#)
 poly, [4](#)
 pull, [38](#)

 quasiquotation, [39, 56](#)

 recode(), [26](#)
 reframe, [13, 50](#)
 reframe(), [54](#)
 rename, [13, 27, 39, 50, 53, 55](#)
 right_join, [3, 40](#)
 rlang::as_function(), [5, 58](#)
 rlang::ensym(), [56](#)
 row_number(), [26, 52](#)
 rownames, [4, 5](#)
 rowwise, [43, 54](#)

 sample_frac(sample_n), [44](#)
 sample_n, [44](#)
 schema(), [36, 38](#)
 sd(), [54](#)
 se, [45](#)
 select, [3, 13, 46](#)
 separate, [50](#)
 separate(), [11, 31, 57](#)
 separate_wider_delim(), [50](#)
 separate_wider_position(), [50](#)
 separate_wider_regex(), [10](#)
 slice, [13, 27, 40, 50, 52, 55](#)
 slice_head(slice), [52](#)
 slice_max(slice), [52](#)
 slice_min(slice), [52](#)
 slice_sample(slice), [52](#)
 slice_sample(), [44](#)
 slice_tail(slice), [52](#)
 spread(), [34](#)
 starts_with(), [46, 47](#)
 style(), [38](#)
 subplot(), [38](#)
 summarise, [13, 27, 40, 50, 53, 53](#)
 summarise(), [18, 43](#)
 summarize(summarise), [53](#)

 table, [4](#)
 tbl_df, [4](#)

tibble, [54](#)
tibble(), [4](#), [6](#)
tidy, [3](#), [55](#)
tidy_(tidy), [55](#)
tidyr_legacy, [58](#)
tidySummarizedExperiment
 (tidySummarizedExperiment-package),
 [3](#)
tidySummarizedExperiment-package, [3](#)
ts, [4](#)
type.convert(), [11](#), [51](#)

ungroup(), [12](#), [43](#)
unique.data.frame(), [9](#)
unite, [56](#)
unite(), [51](#)
unnest, [3](#), [57](#)
unnest_legacy(), [28](#), [59](#)
unnest_longer, [59](#)
unnest_summarized_experiment (unnest),
 [57](#)
unnest_wider, [59](#)

vctrs::vec_as_names(), [5](#), [6](#), [31](#), [34](#), [58](#)

where(), [46](#)

xor(), [12](#)