

Package ‘joyn’

March 30, 2024

Type Package

Title Tool for Diagnosis of Tables Joins and Complementary Join Features

Version 0.2.0

Description Tool for diagnosing table joins. It combines the speed of ``collapse`` and ``data.table``, the flexibility of ``dplyr``, and the diagnosis and features of the ``merge`` command in ``Stata``.

License MIT + file LICENSE

Encoding UTF-8

URL <https://github.com/randrescastaneda/joyn>,
<https://rossanatat.github.io/joyn/>

BugReports <https://github.com/randrescastaneda/joyn/issues>

Suggests badger, covr, knitr, rmarkdown, testthat (>= 3.0.0), withr, dplyr, tibble

Config/testthat/edition 3

Imports rlang, data.table, cli, utils, collapse (>= 2.0.9), lifecycle

Depends R (>= 2.10)

RoxygenNote 7.3.1

VignetteBuilder knitr

NeedsCompilation no

Author R.Andres Castaneda [aut, cre],
Zander Prinsloo [aut],
Rossana Tatulli [aut]

Maintainer R.Andres Castaneda <acastaneda@worldbank.org>

Repository CRAN

Date/Publication 2024-03-29 23:30:02 UTC

R topics documented:

freq_table	2
full_join	3
get_joyn_options	6
inner_join	7
is_balanced	11
is_id	12
joyn	13
joyn_msg	16
joyn_report	17
left_join	18
merge	21
possible_ids	24
rename_to_valid	25
right_join	25
set_joyn_options	29
Index	30

freq_table	<i>Tabulate simple frequencies</i>
------------	------------------------------------

Description

tabulate one variable frequencies

Usage

```
freq_table(x, byvar, digits = 1, na.rm = TRUE)
```

Arguments

x	data frame
byvar	character: name of variable to tabulate. Use Standard evaluation.
digits	numeric: number of decimal places to display. Default is 1.
na.rm	logical: if TRUE remove NAs from calculations. Default is TRUE

Value

data.table with frequencies.

Examples

```
library(data.table)
x4 = data.table(id1 = c(1, 1, 2, 3, 3),
                id2 = c(1, 1, 2, 3, 4),
                t   = c(1L, 2L, 1L, 2L, NA_integer_),
                x   = c(16, 12, NA, NA, 15))
freq_table(x4, "id1")
```

full_join

Full join two data frames

Description

This is a joyn wrapper that works in a similar fashion to [dplyr::full_join](#)

Usage

```
full_join(
  x,
  y,
  by = intersect(names(x), names(y)),
  copy = FALSE,
  suffix = c(".x", ".y"),
  keep = NULL,
  na_matches = c("na", "never"),
  multiple = "all",
  unmatched = "drop",
  relationship = "one-to-one",
  y_vars_to_keep = TRUE,
  update_values = FALSE,
  update_NAs = update_values,
  reportvar = getOption("joyn.reportvar"),
  reporttype = c("character", "numeric"),
  roll = NULL,
  keep_common_vars = FALSE,
  sort = TRUE,
  verbose = getOption("joyn.verbose"),
  ...
)
```

Arguments

x data frame: referred to as *left* in R terminology, or *master* in Stata terminology.
y data frame: referred to as *right* in R terminology, or *using* in Stata terminology.

by	a character vector of variables to join by. If NULL, the default, join will do a natural join, using all variables with common names across the two tables. A message lists the variables so that you can check they're correct (to suppress the message, simply explicitly list the variables that you want to join). To join by different variables on x and y use a vector of expressions. For example, <code>by = c("a = b", "z")</code> will use "a" in x, "b" in y, and "z" in both tables.
copy	If x and y are not from the same data source, and copy is TRUE, then y will be copied into the same src as x. This allows you to join tables across srcs, but it is a potentially expensive operation so you must opt into it.
suffix	If there are non-joined duplicate variables in x and y, these suffixes will be added to the output to disambiguate them. Should be a character vector of length 2.
keep	Should the join keys from both x and y be preserved in the output? <ul style="list-style-type: none"> • If NULL, the default, joins on equality retain only the keys from x, while joins on inequality retain the keys from both inputs. • If TRUE, all keys from both inputs are retained. • If FALSE, only keys from x are retained. For right and full joins, the data in key columns corresponding to rows that only exist in y are merged into the key columns from x. Can't be used when joining on inequality conditions.
na_matches	Should two NA or two NaN values match? <ul style="list-style-type: none"> • "na", the default, treats two NA or two NaN values as equal, like <code>%in%</code>, <code>match()</code>, and <code>merge()</code>. • "never" treats two NA or two NaN values as different, and will never match them together or to any other values. This is similar to joins for database sources and to <code>base::merge(incomparables = NA)</code>.
multiple	Handling of rows in x with multiple matches in y. For each row of x: <ul style="list-style-type: none"> • "all", the default, returns every match detected in y. This is the same behavior as SQL. • "any" returns one match detected in y, with no guarantees on which match will be returned. It is often faster than "first" and "last" if you just need to detect if there is at least one match. • "first" returns the first match detected in y. • "last" returns the last match detected in y.
unmatched	How should unmatched keys that would result in dropped rows be handled? <ul style="list-style-type: none"> • "drop" drops unmatched keys from the result. • "error" throws an error if unmatched keys are detected. <p>unmatched is intended to protect you from accidentally dropping rows during a join. It only checks for unmatched keys in the input that could potentially drop rows.</p> <ul style="list-style-type: none"> • For left joins, it checks y. • For right joins, it checks x. • For inner joins, it checks both x and y. In this case, unmatched is also allowed to be a character vector of length 2 to specify the behavior for x and y independently.

relationship	<p>Handling of the expected relationship between the keys of x and y. If the expectations chosen from the list below are invalidated, an error is thrown.</p> <ul style="list-style-type: none"> • NULL, the default, doesn't expect there to be any relationship between x and y. However, for equality joins it will check for a many-to-many relationship (which is typically unexpected) and will warn if one occurs, encouraging you to either take a closer look at your inputs or make this relationship explicit by specifying "many-to-many". See the <i>Many-to-many relationships</i> section for more details. • "one-to-one" expects: <ul style="list-style-type: none"> – Each row in x matches at most 1 row in y. – Each row in y matches at most 1 row in x. • "one-to-many" expects: <ul style="list-style-type: none"> – Each row in y matches at most 1 row in x. • "many-to-one" expects: <ul style="list-style-type: none"> – Each row in x matches at most 1 row in y. • "many-to-many" doesn't perform any relationship checks, but is provided to allow you to be explicit about this relationship if you know it exists. <p>relationship doesn't handle cases where there are zero matches. For that, see unmatched.</p>
y_vars_to_keep	<p>character: Vector of variable names in y that will be kept after the merge. If TRUE (the default), it keeps all the brings all the variables in y into x. If FALSE or NULL, it does not bring any variable into x, but a report will be generated.</p>
update_values	<p>logical: If TRUE, it will update all values of variables in x with the actual of variables in y with the same name as the ones in x. NAs from y won't be used to update actual values in x. Yet, by default, NAs in x will be updated with values in y. To avoid this, make sure to set update_NAs = FALSE</p>
update_NAs	<p>logical: If TRUE, it will update NA values of all variables in x with actual values of variables in y that have the same name as the ones in x. If FALSE, NA values won't be updated, even if update_values is TRUE</p>
reportvar	<p>character: Name of reporting variable. Default is ".joyn". This is the same as variable "_merge" in Stata after performing a merge. If FALSE or NULL, the reporting variable will be excluded from the final table, though a summary of the join will be display after concluding.</p>
reporttype	<p>character: One of "character" or "numeric". Default is "character". If "numeric", the reporting variable will contain numeric codes of the source and the contents of each observation in the joined table. See below for more information.</p>
roll	<p>double: <i>to be implemented</i></p>
keep_common_vars	<p>logical: If TRUE, it will keep the original variable from y when both tables have common variable names. Thus, the prefix "y." will be added to the original name to distinguish from the resulting variable in the joined table.</p>
sort	<p>logical: If TRUE, sort by key variables in by. Default is TRUE.</p>
verbose	<p>logical: if FALSE, it won't display any message (programmer's option). Default is TRUE.</p>

... Arguments passed on to `join`

`match_type` character: one of "m:m", "m:1", "1:m", "1:1". Default is "1:1" since this the most restrictive. However, following Stata's recommendation, it is better to be explicit and use any of the other three match types (See details in *match types sections*).

`allow.cartesian` logical: Check documentation in official [web site](#). Default is NULL, which implies that if the join is "1:1" it will be FALSE, but if the join has any "m" on it, it will be converted to TRUE. By specifying TRUE or FALSE you force the behavior of the join.

`suffixes` A character(2) specifying the suffixes to be used for making non-by column names unique. The suffix behaviour works in a similar fashion as the `base::merge` method does.

`yvars` **[Superseded]**: use now `y_vars_to_keep`

`keep_y_in_x` **[Superseded]**: use now `keep_common_vars`

`msg_type` character: type of messages to display by default

`na.last` logical. If TRUE, missing values in the data are placed last; if FALSE, they are placed first; if NA they are removed. `na.last=NA` is valid only for `x[order(., na.last)]` and its default is TRUE. `setorder` and `setorderv` only accept TRUE/FALSE with default FALSE.

Value

An data frame of the same class as `x`. The properties of the output are as close as possible to the ones returned by the dplyr alternative.

See Also

Other dplyr alternatives: `inner_join()`, `left_join()`, `right_join()`

Examples

```
# Simple full join
library(data.table)

x1 = data.table(id = c(1L, 1L, 2L, 3L, NA_integer_),
               t   = c(1L, 2L, 1L, 2L, NA_integer_),
               x   = 11:15)
y1 = data.table(id = c(1,2, 4),
               y   = c(11L, 15L, 16))
full_join(x1, y1, relationship = "many-to-one")
```

get_joy_n_options

Get joy_n options

Description

This function aims to display and store info on join options

Usage

```
get_joyn_options(env = .joynenv, display = TRUE, option = NULL)
```

Arguments

env	environment, which is joyn environment by default
display	logical, if TRUE displays (i.e., print) info on joyn options and corresponding default and current values
option	character or NULL. If character, name of a specific joyn option. If NULL, all joyn options

Value

joyn options and values invisibly as a list

See Also

JOYn options functions [set_joyn_options\(\)](#)

Examples

```
## Not run:

# display all joyn options, their default and current values
joyn::get_joyn_options()

# store list of option = value pairs AND do not display info
joyn_options <- joyn::get_joyn_options(display = FALSE)

# get info on one specific option and store it
joyn.verbose <- joyn::get_joyn_options(option = "joyn.verbose")

# get info on two specific option
joyn::get_joyn_options(option = c("joyn.verbose", "joyn.reportvar"))

## End(Not run)
```

inner_join

Inner join two data frames

Description

This is a joyn wrapper that works in a similar fashion to [dplyr::inner_join](#)

Usage

```
inner_join(
  x,
  y,
  by = intersect(names(x), names(y)),
  copy = FALSE,
  suffix = c(".x", ".y"),
  keep = NULL,
  na_matches = c("na", "never"),
  multiple = "all",
  unmatched = "drop",
  relationship = "one-to-one",
  y_vars_to_keep = TRUE,
  update_values = FALSE,
  update_NAs = update_values,
  reportvar = getOption("joyn.reportvar"),
  reporttype = c("character", "numeric"),
  roll = NULL,
  keep_common_vars = FALSE,
  sort = TRUE,
  verbose = getOption("joyn.verbose"),
  ...
)
```

Arguments

x	data frame: referred to as <i>left</i> in R terminology, or <i>master</i> in Stata terminology.
y	data frame: referred to as <i>right</i> in R terminology, or <i>using</i> in Stata terminology.
by	a character vector of variables to join by. If NULL, the default, joyn will do a natural join, using all variables with common names across the two tables. A message lists the variables so that you can check they're correct (to suppress the message, simply explicitly list the variables that you want to join). To join by different variables on x and y use a vector of expressions. For example, <code>by = c("a = b", "z")</code> will use "a" in x, "b" in y, and "z" in both tables.
copy	If x and y are not from the same data source, and copy is TRUE, then y will be copied into the same src as x. This allows you to join tables across srcs, but it is a potentially expensive operation so you must opt into it.
suffix	If there are non-joined duplicate variables in x and y, these suffixes will be added to the output to disambiguate them. Should be a character vector of length 2.
keep	Should the join keys from both x and y be preserved in the output? <ul style="list-style-type: none"> • If NULL, the default, joins on equality retain only the keys from x, while joins on inequality retain the keys from both inputs. • If TRUE, all keys from both inputs are retained. • If FALSE, only keys from x are retained. For right and full joins, the data in key columns corresponding to rows that only exist in y are merged into the key columns from x. Can't be used when joining on inequality conditions.

na_matches	<p>Should two NA or two NaN values match?</p> <ul style="list-style-type: none"> • "na", the default, treats two NA or two NaN values as equal, like <code>%in%</code>, <code>match()</code>, and <code>merge()</code>. • "never" treats two NA or two NaN values as different, and will never match them together or to any other values. This is similar to joins for database sources and to <code>base::merge(incomparables = NA)</code>.
multiple	<p>Handling of rows in <code>x</code> with multiple matches in <code>y</code>. For each row of <code>x</code>:</p> <ul style="list-style-type: none"> • "all", the default, returns every match detected in <code>y</code>. This is the same behavior as SQL. • "any" returns one match detected in <code>y</code>, with no guarantees on which match will be returned. It is often faster than "first" and "last" if you just need to detect if there is at least one match. • "first" returns the first match detected in <code>y</code>. • "last" returns the last match detected in <code>y</code>.
unmatched	<p>How should unmatched keys that would result in dropped rows be handled?</p> <ul style="list-style-type: none"> • "drop" drops unmatched keys from the result. • "error" throws an error if unmatched keys are detected. <p>unmatched is intended to protect you from accidentally dropping rows during a join. It only checks for unmatched keys in the input that could potentially drop rows.</p> <ul style="list-style-type: none"> • For left joins, it checks <code>y</code>. • For right joins, it checks <code>x</code>. • For inner joins, it checks both <code>x</code> and <code>y</code>. In this case, unmatched is also allowed to be a character vector of length 2 to specify the behavior for <code>x</code> and <code>y</code> independently.
relationship	<p>Handling of the expected relationship between the keys of <code>x</code> and <code>y</code>. If the expectations chosen from the list below are invalidated, an error is thrown.</p> <ul style="list-style-type: none"> • NULL, the default, doesn't expect there to be any relationship between <code>x</code> and <code>y</code>. However, for equality joins it will check for a many-to-many relationship (which is typically unexpected) and will warn if one occurs, encouraging you to either take a closer look at your inputs or make this relationship explicit by specifying "many-to-many". See the <i>Many-to-many relationships</i> section for more details. • "one-to-one" expects: <ul style="list-style-type: none"> – Each row in <code>x</code> matches at most 1 row in <code>y</code>. – Each row in <code>y</code> matches at most 1 row in <code>x</code>. • "one-to-many" expects: <ul style="list-style-type: none"> – Each row in <code>y</code> matches at most 1 row in <code>x</code>. • "many-to-one" expects: <ul style="list-style-type: none"> – Each row in <code>x</code> matches at most 1 row in <code>y</code>. • "many-to-many" doesn't perform any relationship checks, but is provided to allow you to be explicit about this relationship if you know it exists. <p>relationship doesn't handle cases where there are zero matches. For that, see unmatched.</p>

y_vars_to_keep	character: Vector of variable names in y that will be kept after the merge. If TRUE (the default), it keeps all the brings all the variables in y into x. If FALSE or NULL, it does not bring any variable into x, but a report will be generated.
update_values	logical: If TRUE, it will update all values of variables in x with the actual of variables in y with the same name as the ones in x. NAs from y won't be used to update actual values in x. Yet, by default, NAs in x will be updated with values in y. To avoid this, make sure to set update_NAs = FALSE
update_NAs	logical: If TRUE, it will update NA values of all variables in x with actual values of variables in y that have the same name as the ones in x. If FALSE, NA values won't be updated, even if update_values is TRUE
reportvar	character: Name of reporting variable. Default is ".joyn". This is the same as variable "_merge" in Stata after performing a merge. If FALSE or NULL, the reporting variable will be excluded from the final table, though a summary of the join will be display after concluding.
reporttype	character: One of "character" or "numeric". Default is "character". If "numeric", the reporting variable will contain numeric codes of the source and the contents of each observation in the joined table. See below for more information.
roll	double: <i>to be implemented</i>
keep_common_vars	logical: If TRUE, it will keep the original variable from y when both tables have common variable names. Thus, the prefix "y." will be added to the original name to distinguish from the resulting variable in the joined table.
sort	logical: If TRUE, sort by key variables in by. Default is TRUE.
verbose	logical: if FALSE, it won't display any message (programmer's option). Default is TRUE.
...	Arguments passed on to joyn
match_type	character: one of "m:m", "m:1", "1:m", "1:1". Default is "1:1" since this the most restrictive. However, following Stata's recommendation, it is better to be explicit and use any of the other three match types (See details in <i>match types sections</i>).
allow.cartesian	logical: Check documentation in official web site . Default is NULL, which implies that if the join is "1:1" it will be FALSE, but if the join has any "m" on it, it will be converted to TRUE. By specifying TRUE of FALSE you force the behavior of the join.
suffixes	A character(2) specifying the suffixes to be used for making non-by column names unique. The suffix behaviour works in a similar fashion as the base::merge method does.
yvars	[Superseded] : use now y_vars_to_keep
keep_y_in_x	[Superseded] : use now keep_common_vars
msg_type	character: type of messages to display by default
na.last	logical. If TRUE, missing values in the data are placed last; if FALSE, they are placed first; if NA they are removed. na.last=NA is valid only for x[order(., na.last)] and its default is TRUE. setorder and setorderv only accept TRUE/FALSE with default FALSE.

Value

An data frame of the same class as `x`. The properties of the output are as close as possible to the ones returned by the `dplyr` alternative.

See Also

Other `dplyr` alternatives: `full_join()`, `left_join()`, `right_join()`

Examples

```
# Simple full join
library(data.table)

x1 = data.table(id = c(1L, 1L, 2L, 3L, NA_integer_),
               t = c(1L, 2L, 1L, 2L, NA_integer_),
               x = 11:15)
y1 = data.table(id = c(1,2, 4),
               y = c(11L, 15L, 16))
inner_join(x1, y1, relationship = "many-to-one")
```

is_balanced

Is data frame balanced by group?

Description

Check if the data frame is balanced by group of columns, i.e., if it contains every combination of the elements in the specified variables

Usage

```
is_balanced(df, by, return = c("logic", "table"))
```

Arguments

<code>df</code>	data frame
<code>by</code>	character: variables used to check if <code>df</code> is balanced
<code>return</code>	character: either "logic" or "table". If "logic", returns TRUE or FALSE depending on whether data frame is balanced. If "table" returns the unbalanced observations - i.e. the combinations of elements in specified variables not found in input <code>df</code>

Value

logical, if `return == "logic"`, else returns data frame of unbalanced observations

Examples

```
x1 = data.frame(id = c(1L, 1L, 2L, 3L, NA_integer_),
               t = c(1L, 2L, 1L, 2L, NA_integer_),
               x = 11:15)
is_balanced(df = x1,
            by = c("id", "t"),
            return = "table") # returns combination of elements in "id" and "t" not present in df
is_balanced(df = x1,
            by = c("id", "t"),
            return = "logic") # FALSE
```

is_id	<i>Check if dt is uniquely identified by by variable</i>
-------	--

Description

report if dt is uniquely identified by by var or, if report = TRUE, the duplicates in by variable

Usage

```
is_id(dt, by, verbose = getOption("joyn.verbose"), return_report = FALSE)
```

Arguments

dt	either right or left table
by	variable to merge by
verbose	logical: if TRUE messages will be displayed
return_report	logical: if TRUE, returns data with summary of duplicates. If FALSE, returns logical value depending on whether dt is uniquely identified by by

Value

logical or data.frame, depending on the value of argument return_report

Examples

```
library(data.table)

# example with data frame not uniquely identified by `by` var
y <- data.table(id = c("c", "b", "c", "a"),
               y = c(11L, 15L, 18L, 20L))
is_id(y, by = "id")
is_id(y, by = "id", return_report = TRUE)

# example with data frame uniquely identified by `by` var
y1 <- data.table(id = c("1", "3", "2", "9"),
```

```

      y = c(11L, 15L, 18L, 20L))
is_id(y1, by = "id")

```

 joyn

Join two tables

Description

This is the primary function in the joyn package. It executes a full join, performs a number of checks, and filters to allow the user-specified join.

Usage

```

joyn(
  x,
  y,
  by = intersect(names(x), names(y)),
  match_type = c("1:1", "1:m", "m:1", "m:m"),
  keep = c("full", "left", "master", "right", "using", "inner"),
  y_vars_to_keep = TRUE,
  update_values = FALSE,
  update_NAs = update_values,
  reportvar = getOption("joyn.reportvar"),
  reporttype = c("character", "numeric"),
  roll = NULL,
  keep_common_vars = FALSE,
  sort = TRUE,
  verbose = getOption("joyn.verbose"),
  suffixes = getOption("joyn.suffixes"),
  allow.cartesian = deprecated(),
  yvars = deprecated(),
  keep_y_in_x = deprecated(),
  na.last = getOption("joyn.na.last"),
  msg_type = getOption("joyn.msg_type")
)

```

Arguments

x	data frame: referred to as <i>left</i> in R terminology, or <i>master</i> in Stata terminology.
y	data frame: referred to as <i>right</i> in R terminology, or <i>using</i> in Stata terminology.
by	a character vector of variables to join by. If NULL, the default, joyn will do a natural join, using all variables with common names across the two tables. A message lists the variables so that you can check they're correct (to suppress the message, simply explicitly list the variables that you want to join). To join by different variables on x and y use a vector of expressions. For example, by = c("a = b", "z") will use "a" in x, "b" in y, and "z" in both tables.

match_type	character: one of "m:m", "m:1", "1:m", "1:1". Default is "1:1" since this the most restrictive. However, following Stata's recommendation, it is better to be explicit and use any of the other three match types (See details in <i>match types sections</i>).
keep	atomic character vector of length 1: One of "full", "left", "master", "right", "using", "inner". Default is "full". Even though this is not the regular behavior of joins in R, the objective of joyn is to present a diagnosis of the join which requires a full join. That is why the default is a a full join. Yet, if "left" or "master", it keeps the observations that matched in both tables and the ones that did not match in x. The ones in y will be discarded. If "right" or "using", it keeps the observations that matched in both tables and the ones that did not match in y. The ones in x will be discarded. If "inner", it only keeps the observations that matched both tables. Note that if, for example, a keep = "left", the joyn() function still executes a full join under the hood and then filters so that only rows the output table is a left join. This behaviour, while inefficient, allows all the diagnostics and checks conducted by joyn.
y_vars_to_keep	character: Vector of variable names in y that will be kept after the merge. If TRUE (the default), it keeps all the brings all the variables in y into x. If FALSE or NULL, it does not bring any variable into x, but a report will be generated.
update_values	logical: If TRUE, it will update all values of variables in x with the actual of variables in y with the same name as the ones in x. NAs from y won't be used to update actual values in x. Yet, by default, NAs in x will be updated with values in y. To avoid this, make sure to set update_NAs = FALSE
update_NAs	logical: If TRUE, it will update NA values of all variables in x with actual values of variables in y that have the same name as the ones in x. If FALSE, NA values won't be updated, even if update_values is TRUE
reportvar	character: Name of reporting variable. Default is ".joyn". This is the same as variable "_merge" in Stata after performing a merge. If FALSE or NULL, the reporting variable will be excluded from the final table, though a summary of the join will be display after concluding.
reporttype	character: One of "character" or "numeric". Default is "character". If "numeric", the reporting variable will contain numeric codes of the source and the contents of each observation in the joined table. See below for more information.
roll	double: <i>to be implemented</i>
keep_common_vars	logical: If TRUE, it will keep the original variable from y when both tables have common variable names. Thus, the prefix "y." will be added to the original name to distinguish from the resulting variable in the joined table.
sort	logical: If TRUE, sort by key variables in by. Default is TRUE.
verbose	logical: if FALSE, it won't display any message (programmer's option). Default is TRUE.
suffixes	A character(2) specifying the suffixes to be used for making non-by column names unique. The suffix behaviour works in a similar fashion as the <code>base::merge</code> method does.

<code>allow.cartesian</code>	logical: Check documentation in official web site . Default is NULL, which implies that if the join is "1:1" it will be FALSE, but if the join has any "m" on it, it will be converted to TRUE. By specifying TRUE or FALSE you force the behavior of the join.
<code>yvars</code>	[Superseded] : use now <code>y_vars_to_keep</code>
<code>keep_y_in_x</code>	[Superseded] : use now <code>keep_common_vars</code>
<code>na.last</code>	logical. If TRUE, missing values in the data are placed last; if FALSE, they are placed first; if NA they are removed. <code>na.last=NA</code> is valid only for <code>x[order(. , na.last)]</code> and its default is TRUE. <code>setorder</code> and <code>setorderv</code> only accept TRUE/FALSE with default FALSE.
<code>msg_type</code>	character: type of messages to display by default

Value

a `data.table` joining `x` and `y`.

match types

Using the same wording of the [Stata manual](#)

1:1: specifies a one-to-one match merge. The variables specified in `by` uniquely identify single observations in both table.

1:m and m:1: specify *one-to-many* and *many-to-one* match merges, respectively. This means that in one of the tables the observations are uniquely identified by the variables in `by`, while in the other table many (two or more) of the observations are identified by the variables in `by`.

m:m refers to *many-to-many merge*. `variables in by` does not uniquely identify the observations in either table. Matching is performed by combining observations with equal values in `by`; within matching values, the first observation in the master (i.e. left or `x`) table is matched with the first matching observation in the using (i.e. right or `y`) table; the second, with the second; and so on. If there is an unequal number of observations within a group, then the last observation of the shorter group is used repeatedly to match with subsequent observations of the longer group.

reporttype

If `reporttype = "numeric"`, then the numeric values have the following meaning:

1: row comes from `x`, i.e. "x" 2: row comes from `y`, i.e. "y" 3: row from both `x` and `y`, i.e. "x & y" 4: row has NA in `x` that has been updated with `y`, i.e. "NA updated" 5: row has value in `x` that has been updated with `y`, i.e. "value updated" 6: row from `x` that has not been updated, i.e. "not updated"

NAs order

NAs are placed either at first or at last in the resulting `data.frame` depending on the value of `getOption("joyn.na.last")`. The Default is FALSE as it is the default value of `data.table::setorderv`.

Examples

```

# Simple join
library(data.table)
x1 = data.table(id = c(1L, 1L, 2L, 3L, NA_integer_),
  t = c(1L, 2L, 1L, 2L, NA_integer_),
  x = 11:15)

y1 = data.table(id = 1:2,
  y = c(11L, 15L))

x2 = data.table(id = c(1, 1, 2, 3, NA),
  t = c(1L, 2L, 1L, 2L, NA_integer_),
  x = c(16, 12, NA, NA, 15))

y2 = data.table(id = c(1, 2, 5, 6, 3),
  yd = c(1, 2, 5, 6, 3),
  y = c(11L, 15L, 20L, 13L, 10L),
  x = c(16:20))
joyn(x1, y1, match_type = "m:1")

# Bad merge for not specifying by argument or match_type
joyn(x2, y2)

# good merge, ignoring variable x from y
joyn(x2, y2, by = "id", match_type = "m:1")

# update NAs in x variable from x
joyn(x2, y2, by = "id", update_NAs = TRUE, match_type = "m:1")

# Update values in x with variables from y
joyn(x2, y2, by = "id", update_values = TRUE, match_type = "m:1")

```

joyn_msg

display type of joyn message

Description

display type of joyn message

Usage

```
joyn_msg(msg_type = getOption("joyn.msg_type"), msg = NULL)
```

Arguments

msg_type	character: one or more of the following: all, basic, info, note, warn, timing, or err
msg	character vector to be parsed to <code>cli::cli_abort()</code> . Default is NULL. It only works if "err" %in% msg_type. This is an internal argument.

Value

returns data frame with message invisibly. print message in console

See Also

Messages functions [clear_joynenv\(\)](#), [joyn_msgs_exist\(\)](#), [joyn_report\(\)](#), [msg_type_dt\(\)](#), [store_msg\(\)](#), [style\(\)](#), [type_choices\(\)](#)

Examples

```
library(data.table)
x1 = data.table(id = c(1L, 1L, 2L, 3L, NA_integer_),
               t = c(1L, 2L, 1L, 2L, NA_integer_),
               x = 11:15)

y1 = data.table(id = 1:2,
               y = c(11L, 15L))
df <- joyn(x1, y1, match_type = "m:1")
joyn_msg("basic")
joyn_msg("all")
```

joyn_report	<i>Print JOYn report table</i>
-------------	--------------------------------

Description

Print JOYn report table

Usage

```
joyn_report(verbose = getOption("joyn.verbose"))
```

Arguments

verbose logical: if FALSE, it won't display any message (programmer's option). Default is TRUE.

Value

invisible table of frequencies

See Also

Messages functions [clear_joynenv\(\)](#), [joyn_msg\(\)](#), [joyn_msgs_exist\(\)](#), [msg_type_dt\(\)](#), [store_msg\(\)](#), [style\(\)](#), [type_choices\(\)](#)

Examples

```
library(data.table)
x1 = data.table(id = c(1L, 1L, 2L, 3L, NA_integer_),
  t = c(1L, 2L, 1L, 2L, NA_integer_),
  x = 11:15)

y1 = data.table(id = 1:2,
  y = c(11L, 15L))

d <- joyn(x1, y1, match_type = "m:1")
joyn_report(verbose = TRUE)
```

left_join	<i>Left join two data frames</i>
-----------	----------------------------------

Description

This is a joyn wrapper that works in a similar fashion to [dplyr::left_join](#)

Usage

```
left_join(
  x,
  y,
  by = intersect(names(x), names(y)),
  copy = FALSE,
  suffix = c(".x", ".y"),
  keep = NULL,
  na_matches = c("na", "never"),
  multiple = "all",
  unmatched = "drop",
  relationship = NULL,
  y_vars_to_keep = TRUE,
  update_values = FALSE,
  update_NAs = update_values,
  reportvar = getOption("joyn.reportvar"),
  reporttype = c("character", "numeric"),
  roll = NULL,
  keep_common_vars = FALSE,
  sort = TRUE,
  verbose = getOption("joyn.verbose"),
  ...
)
```

Arguments

x data frame: referred to as *left* in R terminology, or *master* in Stata terminology.

y	data frame: referred to as <i>right</i> in R terminology, or <i>using</i> in Stata terminology.
by	a character vector of variables to join by. If NULL, the default, join will do a natural join, using all variables with common names across the two tables. A message lists the variables so that you can check they're correct (to suppress the message, simply explicitly list the variables that you want to join). To join by different variables on x and y use a vector of expressions. For example, by = c("a = b", "z") will use "a" in x, "b" in y, and "z" in both tables.
copy	If x and y are not from the same data source, and copy is TRUE, then y will be copied into the same src as x. This allows you to join tables across srcs, but it is a potentially expensive operation so you must opt into it.
suffix	If there are non-joined duplicate variables in x and y, these suffixes will be added to the output to disambiguate them. Should be a character vector of length 2.
keep	Should the join keys from both x and y be preserved in the output? <ul style="list-style-type: none"> • If NULL, the default, joins on equality retain only the keys from x, while joins on inequality retain the keys from both inputs. • If TRUE, all keys from both inputs are retained. • If FALSE, only keys from x are retained. For right and full joins, the data in key columns corresponding to rows that only exist in y are merged into the key columns from x. Can't be used when joining on inequality conditions.
na_matches	Should two NA or two NaN values match? <ul style="list-style-type: none"> • "na", the default, treats two NA or two NaN values as equal, like <code>%in%</code>, <code>match()</code>, and <code>merge()</code>. • "never" treats two NA or two NaN values as different, and will never match them together or to any other values. This is similar to joins for database sources and to <code>base::merge(incomparables = NA)</code>.
multiple	Handling of rows in x with multiple matches in y. For each row of x: <ul style="list-style-type: none"> • "all", the default, returns every match detected in y. This is the same behavior as SQL. • "any" returns one match detected in y, with no guarantees on which match will be returned. It is often faster than "first" and "last" if you just need to detect if there is at least one match. • "first" returns the first match detected in y. • "last" returns the last match detected in y.
unmatched	How should unmatched keys that would result in dropped rows be handled? <ul style="list-style-type: none"> • "drop" drops unmatched keys from the result. • "error" throws an error if unmatched keys are detected. <p>unmatched is intended to protect you from accidentally dropping rows during a join. It only checks for unmatched keys in the input that could potentially drop rows.</p> <ul style="list-style-type: none"> • For left joins, it checks y. • For right joins, it checks x. • For inner joins, it checks both x and y. In this case, unmatched is also allowed to be a character vector of length 2 to specify the behavior for x and y independently.

relationship	<p>Handling of the expected relationship between the keys of x and y. If the expectations chosen from the list below are invalidated, an error is thrown.</p> <ul style="list-style-type: none"> • NULL, the default, doesn't expect there to be any relationship between x and y. However, for equality joins it will check for a many-to-many relationship (which is typically unexpected) and will warn if one occurs, encouraging you to either take a closer look at your inputs or make this relationship explicit by specifying "many-to-many". See the <i>Many-to-many relationships</i> section for more details. • "one-to-one" expects: <ul style="list-style-type: none"> – Each row in x matches at most 1 row in y. – Each row in y matches at most 1 row in x. • "one-to-many" expects: <ul style="list-style-type: none"> – Each row in y matches at most 1 row in x. • "many-to-one" expects: <ul style="list-style-type: none"> – Each row in x matches at most 1 row in y. • "many-to-many" doesn't perform any relationship checks, but is provided to allow you to be explicit about this relationship if you know it exists. <p>relationship doesn't handle cases where there are zero matches. For that, see unmatched.</p>
y_vars_to_keep	<p>character: Vector of variable names in y that will be kept after the merge. If TRUE (the default), it keeps all the brings all the variables in y into x. If FALSE or NULL, it does not bring any variable into x, but a report will be generated.</p>
update_values	<p>logical: If TRUE, it will update all values of variables in x with the actual of variables in y with the same name as the ones in x. NAs from y won't be used to update actual values in x. Yet, by default, NAs in x will be updated with values in y. To avoid this, make sure to set update_NAs = FALSE</p>
update_NAs	<p>logical: If TRUE, it will update NA values of all variables in x with actual values of variables in y that have the same name as the ones in x. If FALSE, NA values won't be updated, even if update_values is TRUE</p>
reportvar	<p>character: Name of reporting variable. Default is ".joyn". This is the same as variable "_merge" in Stata after performing a merge. If FALSE or NULL, the reporting variable will be excluded from the final table, though a summary of the join will be display after concluding.</p>
reporttype	<p>character: One of "character" or "numeric". Default is "character". If "numeric", the reporting variable will contain numeric codes of the source and the contents of each observation in the joined table. See below for more information.</p>
roll	<p>double: <i>to be implemented</i></p>
keep_common_vars	<p>logical: If TRUE, it will keep the original variable from y when both tables have common variable names. Thus, the prefix "y." will be added to the original name to distinguish from the resulting variable in the joined table.</p>
sort	<p>logical: If TRUE, sort by key variables in by. Default is TRUE.</p>
verbose	<p>logical: if FALSE, it won't display any message (programmer's option). Default is TRUE.</p>

... Arguments passed on to [joyn](#)

`match_type` character: one of "m:m", "m:1", "1:m", "1:1". Default is "1:1" since this the most restrictive. However, following Stata's recommendation, it is better to be explicit and use any of the other three match types (See details in *match types sections*).

`allow.cartesian` logical: Check documentation in official [web site](#). Default is NULL, which implies that if the join is "1:1" it will be FALSE, but if the join has any "m" on it, it will be converted to TRUE. By specifying TRUE or FALSE you force the behavior of the join.

`suffixes` A character(2) specifying the suffixes to be used for making non-by column names unique. The suffix behaviour works in a similar fashion as the [base::merge](#) method does.

`yvars` [**Superseded**]: use now `y_vars_to_keep`

`keep_y_in_x` [**Superseded**]: use now `keep_common_vars`

`msg_type` character: type of messages to display by default

`na.last` logical. If TRUE, missing values in the data are placed last; if FALSE, they are placed first; if NA they are removed. `na.last=NA` is valid only for `x[order(., na.last)]` and its default is TRUE. `setorder` and `setorderv` only accept TRUE/FALSE with default FALSE.

Value

An data frame of the same class as `x`. The properties of the output are as close as possible to the ones returned by the `dplyr` alternative.

See Also

Other `dplyr` alternatives: [full_join\(\)](#), [inner_join\(\)](#), [right_join\(\)](#)

Examples

```
# Simple left join
library(data.table)

x1 = data.table(id = c(1L, 1L, 2L, 3L, NA_integer_),
               t   = c(1L, 2L, 1L, 2L, NA_integer_),
               x   = 11:15)
y1 = data.table(id = c(1,2, 4),
               y   = c(11L, 15L, 16))
left_join(x1, y1, relationship = "many-to-one")
```

merge

Merge two data frames

Description

This is a `joyn` wrapper that works in a similar fashion to [base::merge](#) and [data.table::merge](#), which is why [merge](#) masks the other two.

Usage

```
merge(
  x,
  y,
  by = NULL,
  by.x = NULL,
  by.y = NULL,
  all = FALSE,
  all.x = all,
  all.y = all,
  sort = TRUE,
  suffixes = c(".x", ".y"),
  no.dups = TRUE,
  allow.cartesian = getOption("datatable.allow.cartesian"),
  match_type = c("m:m", "m:1", "1:m", "1:1"),
  keep_common_vars = TRUE,
  ...
)
```

Arguments

<code>x, y</code>	data tables. <code>y</code> is coerced to a <code>data.table</code> if it isn't one already.
<code>by</code>	A vector of shared column names in <code>x</code> and <code>y</code> to merge on. This defaults to the shared key columns between the two tables. If <code>y</code> has no key columns, this defaults to the key of <code>x</code> .
<code>by.x, by.y</code>	Vectors of column names in <code>x</code> and <code>y</code> to merge on.
<code>all</code>	logical; <code>all = TRUE</code> is shorthand to save setting both <code>all.x = TRUE</code> and <code>all.y = TRUE</code> .
<code>all.x</code>	logical; if <code>TRUE</code> , rows from <code>x</code> which have no matching row in <code>y</code> are included. These rows will have 'NA's in the columns that are usually filled with values from <code>y</code> . The default is <code>FALSE</code> so that only rows with data from both <code>x</code> and <code>y</code> are included in the output.
<code>all.y</code>	logical; analogous to <code>all.x</code> above.
<code>sort</code>	logical. If <code>TRUE</code> (default), the rows of the merged <code>data.table</code> are sorted by setting the key to the <code>by / by.x</code> columns. If <code>FALSE</code> , unlike base R's <code>merge</code> for which row order is unspecified, the row order in <code>x</code> is retained (including retaining the position of missings when <code>all.x=TRUE</code>), followed by <code>y</code> rows that don't match <code>x</code> (when <code>all.y=TRUE</code>) retaining the order those appear in <code>y</code> .
<code>suffixes</code>	A character(2) specifying the suffixes to be used for making non- <code>by</code> column names unique. The suffix behaviour works in a similar fashion as the merge.data.frame method does.
<code>no.dups</code>	logical indicating that <code>suffixes</code> are also appended to non- <code>by.y</code> column names in <code>y</code> when they have the same column name as any <code>by.x</code> .
<code>allow.cartesian</code>	See <code>allow.cartesian</code> in [.data.table] .

match_type	character: one of "m:m", "m:1", "1:m", "1:1". Default is "1:1" since this the most restrictive. However, following Stata's recommendation, it is better to be explicit and use any of the other three match types (See details in <i>match types sections</i>).
keep_common_vars	logical: If TRUE, it will keep the original variable from y when both tables have common variable names. Thus, the prefix "y." will be added to the original name to distinguish from the resulting variable in the joined table.
...	Arguments passed on to <code>joyn</code>
y_vars_to_keep	character: Vector of variable names in y that will be kept after the merge. If TRUE (the default), it keeps all the brings all the variables in y into x. If FALSE or NULL, it does not bring any variable into x, but a report will be generated.
reportvar	character: Name of reporting variable. Default is ".joyn". This is the same as variable "_merge" in Stata after performing a merge. If FALSE or NULL, the reporting variable will be excluded from the final table, though a summary of the join will be display after concluding.
update_NAs	logical: If TRUE, it will update NA values of all variables in x with actual values of variables in y that have the same name as the ones in x. If FALSE, NA values won't be updated, even if update_values is TRUE
update_values	logical: If TRUE, it will update all values of variables in x with the actual of variables in y with the same name as the ones in x. NAs from y won't be used to update actual values in x. Yet, by default, NAs in x will be updated with values in y. To avoid this, make sure to set update_NAs = FALSE
verbose	logical: if FALSE, it won't display any message (programmer's option). Default is TRUE.

Value

data.table merging x and y

Examples

```
x1 = data.frame(id = c(1L, 1L, 2L, 3L, NA_integer_),
               t   = c(1L, 2L, 1L, 2L, NA_integer_),
               x   = 11:15)
y1 = data.frame(id = c(1,2, 4),
               y   = c(11L, 15L, 16))
joyn::merge(x1, y1, by = "id")
# example of using by.x and by.y
x2 = data.frame(id1 = c(1, 1, 2, 3, 3),
               id2 = c(1, 1, 2, 3, 4),
               t   = c(1L, 2L, 1L, 2L, NA_integer_),
               x   = c(16, 12, NA, NA, 15))
y2 = data.frame(id = c(1, 2, 5, 6, 3),
               id2 = c(1, 1, 2, 3, 4),
               y   = c(11L, 15L, 20L, 13L, 10L),
               x   = c(16:20))
```

```
jn <- joyn::merge(x2,
  y2,
  match_type = "m:m",
  all.x = TRUE,
  by.x = "id1",
  by.y = "id2")
# example with all = TRUE
jn <- joyn::merge(x2,
  y2,
  match_type = "m:m",
  by.x = "id1",
  by.y = "id2",
  all = TRUE)
```

possible_ids

Find possible unique identifies of data frame

Description

Identify possible variables uniquely identifying x

Usage

```
possible_ids(
  dt,
  exclude = NULL,
  include = NULL,
  verbose = getOption("possible_ids.verbose")
)
```

Arguments

dt	data frame
exclude	character: Exclude variables to be selected as identifiers. It could be either the name of the variables of one type of the variable prefixed by "_". For instance, "_numeric" or "_character".
include	character: Name of variable to be included, that might belong to the group excluded in the exclude
verbose	logical: If FALSE no message will be displayed. Default is TRUE

Value

list with possible identifiers

Examples

```
library(data.table)
x4 = data.table(id1 = c(1, 1, 2, 3, 3),
               id2 = c(1, 1, 2, 3, 4),
               t   = c(1L, 2L, 1L, 2L, NA_integer_),
               x   = c(16, 12, NA, NA, 15))
possible_ids(x4)
```

rename_to_valid	<i>Rename to syntactically valid names</i>
-----------------	--

Description

Rename to syntactically valid names

Usage

```
rename_to_valid(name, verbose = getOption("joyn.verbose"))
```

Arguments

name	character: name to be coerced to syntactically valid name
verbose	logical: if FALSE, it won't display any message (programmer's option). Default is TRUE.

Value

valid character name

Examples

```
joyn:::rename_to_valid("x y")
```

right_join	<i>Right join two data frames</i>
------------	-----------------------------------

Description

This is a joyn wrapper that works in a similar fashion to [dplyr::right_join](#)

Usage

```
right_join(
  x,
  y,
  by = intersect(names(x), names(y)),
  copy = FALSE,
  suffix = c(".x", ".y"),
  keep = NULL,
  na_matches = c("na", "never"),
  multiple = "all",
  unmatched = "drop",
  relationship = "one-to-one",
  y_vars_to_keep = TRUE,
  update_values = FALSE,
  update_NAs = update_values,
  reportvar = getOption("joyn.reportvar"),
  reporttype = c("character", "numeric"),
  roll = NULL,
  keep_common_vars = FALSE,
  sort = TRUE,
  verbose = getOption("joyn.verbose"),
  ...
)
```

Arguments

x	data frame: referred to as <i>left</i> in R terminology, or <i>master</i> in Stata terminology.
y	data frame: referred to as <i>right</i> in R terminology, or <i>using</i> in Stata terminology.
by	a character vector of variables to join by. If NULL, the default, joyn will do a natural join, using all variables with common names across the two tables. A message lists the variables so that you can check they're correct (to suppress the message, simply explicitly list the variables that you want to join). To join by different variables on x and y use a vector of expressions. For example, by = c("a = b", "z") will use "a" in x, "b" in y, and "z" in both tables.
copy	If x and y are not from the same data source, and copy is TRUE, then y will be copied into the same src as x. This allows you to join tables across srcs, but it is a potentially expensive operation so you must opt into it.
suffix	If there are non-joined duplicate variables in x and y, these suffixes will be added to the output to disambiguate them. Should be a character vector of length 2.
keep	Should the join keys from both x and y be preserved in the output? <ul style="list-style-type: none"> • If NULL, the default, joins on equality retain only the keys from x, while joins on inequality retain the keys from both inputs. • If TRUE, all keys from both inputs are retained. • If FALSE, only keys from x are retained. For right and full joins, the data in key columns corresponding to rows that only exist in y are merged into the key columns from x. Can't be used when joining on inequality conditions.

na_matches	<p>Should two NA or two NaN values match?</p> <ul style="list-style-type: none"> • "na", the default, treats two NA or two NaN values as equal, like <code>%in%</code>, <code>match()</code>, and <code>merge()</code>. • "never" treats two NA or two NaN values as different, and will never match them together or to any other values. This is similar to joins for database sources and to <code>base::merge(incomparables = NA)</code>.
multiple	<p>Handling of rows in <code>x</code> with multiple matches in <code>y</code>. For each row of <code>x</code>:</p> <ul style="list-style-type: none"> • "all", the default, returns every match detected in <code>y</code>. This is the same behavior as SQL. • "any" returns one match detected in <code>y</code>, with no guarantees on which match will be returned. It is often faster than "first" and "last" if you just need to detect if there is at least one match. • "first" returns the first match detected in <code>y</code>. • "last" returns the last match detected in <code>y</code>.
unmatched	<p>How should unmatched keys that would result in dropped rows be handled?</p> <ul style="list-style-type: none"> • "drop" drops unmatched keys from the result. • "error" throws an error if unmatched keys are detected. <p>unmatched is intended to protect you from accidentally dropping rows during a join. It only checks for unmatched keys in the input that could potentially drop rows.</p> <ul style="list-style-type: none"> • For left joins, it checks <code>y</code>. • For right joins, it checks <code>x</code>. • For inner joins, it checks both <code>x</code> and <code>y</code>. In this case, unmatched is also allowed to be a character vector of length 2 to specify the behavior for <code>x</code> and <code>y</code> independently.
relationship	<p>Handling of the expected relationship between the keys of <code>x</code> and <code>y</code>. If the expectations chosen from the list below are invalidated, an error is thrown.</p> <ul style="list-style-type: none"> • NULL, the default, doesn't expect there to be any relationship between <code>x</code> and <code>y</code>. However, for equality joins it will check for a many-to-many relationship (which is typically unexpected) and will warn if one occurs, encouraging you to either take a closer look at your inputs or make this relationship explicit by specifying "many-to-many". See the <i>Many-to-many relationships</i> section for more details. • "one-to-one" expects: <ul style="list-style-type: none"> – Each row in <code>x</code> matches at most 1 row in <code>y</code>. – Each row in <code>y</code> matches at most 1 row in <code>x</code>. • "one-to-many" expects: <ul style="list-style-type: none"> – Each row in <code>y</code> matches at most 1 row in <code>x</code>. • "many-to-one" expects: <ul style="list-style-type: none"> – Each row in <code>x</code> matches at most 1 row in <code>y</code>. • "many-to-many" doesn't perform any relationship checks, but is provided to allow you to be explicit about this relationship if you know it exists. <p>relationship doesn't handle cases where there are zero matches. For that, see unmatched.</p>

y_vars_to_keep	character: Vector of variable names in y that will be kept after the merge. If TRUE (the default), it keeps all the brings all the variables in y into x. If FALSE or NULL, it does not bring any variable into x, but a report will be generated.
update_values	logical: If TRUE, it will update all values of variables in x with the actual of variables in y with the same name as the ones in x. NAs from y won't be used to update actual values in x. Yet, by default, NAs in x will be updated with values in y. To avoid this, make sure to set update_NAs = FALSE
update_NAs	logical: If TRUE, it will update NA values of all variables in x with actual values of variables in y that have the same name as the ones in x. If FALSE, NA values won't be updated, even if update_values is TRUE
reportvar	character: Name of reporting variable. Default is ".joyn". This is the same as variable "_merge" in Stata after performing a merge. If FALSE or NULL, the reporting variable will be excluded from the final table, though a summary of the join will be display after concluding.
reporttype	character: One of "character" or "numeric". Default is "character". If "numeric", the reporting variable will contain numeric codes of the source and the contents of each observation in the joined table. See below for more information.
roll	double: <i>to be implemented</i>
keep_common_vars	logical: If TRUE, it will keep the original variable from y when both tables have common variable names. Thus, the prefix "y." will be added to the original name to distinguish from the resulting variable in the joined table.
sort	logical: If TRUE, sort by key variables in by. Default is TRUE.
verbose	logical: if FALSE, it won't display any message (programmer's option). Default is TRUE.
...	Arguments passed on to joyn
match_type	character: one of "m:m", "m:1", "1:m", "1:1". Default is "1:1" since this the most restrictive. However, following Stata's recommendation, it is better to be explicit and use any of the other three match types (See details in <i>match types sections</i>).
allow.cartesian	logical: Check documentation in official web site . Default is NULL, which implies that if the join is "1:1" it will be FALSE, but if the join has any "m" on it, it will be converted to TRUE. By specifying TRUE of FALSE you force the behavior of the join.
suffixes	A character(2) specifying the suffixes to be used for making non-by column names unique. The suffix behaviour works in a similar fashion as the base::merge method does.
yvars	[Superseded] : use now y_vars_to_keep
keep_y_in_x	[Superseded] : use now keep_common_vars
msg_type	character: type of messages to display by default
na.last	logical. If TRUE, missing values in the data are placed last; if FALSE, they are placed first; if NA they are removed. na.last=NA is valid only for x[order(., na.last)] and its default is TRUE. setorder and setorderv only accept TRUE/FALSE with default FALSE.

Value

An data frame of the same class as *x*. The properties of the output are as close as possible to the ones returned by the dplyr alternative.

See Also

Other dplyr alternatives: [full_join\(\)](#), [inner_join\(\)](#), [left_join\(\)](#)

Examples

```
# Simple right join
library(data.table)

x1 = data.table(id = c(1L, 1L, 2L, 3L, NA_integer_),
               t   = c(1L, 2L, 1L, 2L, NA_integer_),
               x   = 11:15)
y1 = data.table(id = c(1,2, 4),
               y   = c(11L, 15L, 16))
right_join(x1, y1, relationship = "many-to-one")
```

set_joy_n_options	<i>Set joy_n options</i>
-------------------	--------------------------

Description

This function is used to change the value of one or more joy_n options

Usage

```
set_joy_n_options(..., env = .joy_nenv)
```

Arguments

...	pairs of option = value
env	environment, which is joy_n environment by default

Value

joy_n new options and values invisibly as a list

See Also

JOYn options functions [get_joy_n_options\(\)](#)

Examples

```
joy_n::set_joy_n_options(joy_n.verbose = FALSE, joy_n.reportvar = "joy_n_status")
joy_n::set_joy_n_options() # return to default options
```

Index

* dplyr alternatives

- full_join, 3
- inner_join, 7
- left_join, 18
- right_join, 25

* messages

- joyn_msg, 16
- joyn_report, 17

* options

- get_joyn_options, 6
- set_joyn_options, 29

[.data.table, 22

base::merge, 6, 10, 14, 21, 28

clear_joynenv, 17

cli::cli_abort(), 16

data.table::merge, 21

data.table::setorderv, 15

dplyr::full_join, 3

dplyr::inner_join, 7

dplyr::left_join, 18

dplyr::right_join, 25

freq_table, 2

full_join, 3, 11, 21, 29

get_joyn_options, 6, 29

inner_join, 6, 7, 21, 29

is_balanced, 11

is_id, 12

joyn, 6, 10, 13, 21, 23, 28

joyn_msg, 16, 17

joyn_msgs_exist, 17

joyn_report, 17, 17

left_join, 6, 11, 18, 29

match(), 4, 9, 19, 27

merge, 21, 21

merge(), 4, 9, 19, 27

merge.data.frame, 22

msg_type_dt, 17

possible_ids, 24

rename_to_valid, 25

right_join, 6, 11, 21, 25

set_joyn_options, 7, 29

store_msg, 17

style, 17

type_choices, 17