

Package: tidypolars (via r-universe)

July 5, 2024

Type Package

Title Get the Power of Polars with the Syntax of the Tidyverse

Version 0.9.0

Description Polars is an amazing cross-language tool for manipulating very large data. However, one drawback is that the R implementation has a syntax that will look odd to many R users who are not used to Python syntax. The objective of tidypolars is to improve the ease-of-use of Polars in R by providing tidyverse syntax to polars.

License MIT + file LICENSE

Encoding UTF-8

LazyData true

RoxygenNote 7.3.1

Roxygen list(markdown = TRUE)

URL <https://www.tidypolars.etiennebacher.com/>,
<https://etiennebacher.r-universe.dev/tidypolars>

BugReports <https://github.com/etiennebacher/tidypolars/issues>

Depends R (>= 4.1.0)

Imports dplyr, glue, polars (>= 0.18.0), rlang, tidyr, tidysselect,
utils, vctrs

Suggests bench, data.table, knitr, lubridate, rmarkdown, rstudioapi,
tibble, tinytest

Remotes markvanderloo/tinytest/pkg

VignetteBuilder knitr

Repository <https://etiennebacher.r-universe.dev>

RemoteUrl <https://github.com/etiennebacher/tidypolars>

RemoteRef v0.9.0

RemoteSha b856d174d90010dde55a4348a1496ded797f40af

Contents

arrange.RPolarsDataFrame	3
as_tibble.tidypolars	3
bind_cols_polars	4
bind_rows_polars	5
complete.RPolarsDataFrame	6
compute.RPolarsLazyFrame	7
count.RPolarsDataFrame	9
cross_join.RPolarsDataFrame	10
describe	11
describe_plan	11
distinct.RPolarsDataFrame	12
drop_na.RPolarsDataFrame	13
fetch	14
fill.RPolarsDataFrame	15
filter.RPolarsDataFrame	16
group_by.RPolarsDataFrame	17
group_split.RPolarsDataFrame	18
group_vars.RPolarsDataFrame	19
left_join.RPolarsDataFrame	20
make_unique_id	24
mutate.RPolarsDataFrame	25
pivot_longer.RPolarsDataFrame	27
pivot_wider.RPolarsDataFrame	28
pull.RPolarsDataFrame	30
relocate.RPolarsDataFrame	31
rename.RPolarsDataFrame	32
replace_na.RPolarsDataFrame	33
rowwise.RPolarsDataFrame	34
select.RPolarsDataFrame	35
semi_join.RPolarsDataFrame	35
separate.RPolarsDataFrame	37
sink_csv	38
sink_parquet	40
slice_tail.RPolarsDataFrame	42
summarize.RPolarsDataFrame	43
tidypolars-options	44
uncount.RPolarsDataFrame	45
unite.RPolarsDataFrame	46

Index

48

```
arrange.RPolarsDataFrame
```

Order rows using column values

Description

Order rows using column values

Usage

```
## S3 method for class 'RPolarsDataFrame'
arrange(.data, ..., .by_group = FALSE)
```

Arguments

<code>.data</code>	A Polars Data/LazyFrame
<code>...</code>	Quoted or unquoted variable names. Select helpers cannot be used.
<code>.by_group</code>	If TRUE, will sort data within groups.

Examples

```
pl_test <- polars::pl$DataFrame(
  x1 = c("a", "a", "b", "a", "c"),
  x2 = c(2, 1, 5, 3, 1),
  value = sample(1:5)
)

arrange(pl_test, x1)
arrange(pl_test, x1, -x2)

# if the data is grouped, you need to specify `by_group = TRUE` to sort by
# the groups first
pl_test |>
  group_by(x1) |>
  arrange(-x2, .by_group = TRUE)
```

```
as_tibble.tidypolars
```

Convert a Polars DataFrame to an R data.frame or to a tibble

Description

This makes it easier to convert a polars [DataFrame](#) or [LazyFrame](#) to a [tibble](#) in a pipe workflow.

Usage

```
## S3 method for class 'tidypolars'
as_tibble(x, int64_conversion = polars::polars_options()$int64_conversion, ...)
```

Arguments

`x` A Polars Data/LazyFrame

`int64_conversion` How should Int64 values be handled when converting a polars object to R? See the documentation in [polars::as.data.frame.RPolarsDataFrame](#).

`...` Options passed to [polars::as.data.frame.RPolarsDataFrame](#).

About int64

Int64 is a format accepted in Polars but not natively in R (the package `bit64` helps with that).

Since `tidypolars` is simply a wrapper around `polars`, the behavior of `int64` values will depend on the options set in `polars`. Use `options(polars.int64_conversion =)` to specify how `int64` variables should be handled. See the [documentation in polars](#) for the possible options.

Examples

```
iris |>
  as_polars_df() |>
  filter(Sepal.Length > 6) |>
  as_tibble()
```

`bind_cols_polars` *Append multiple Data/LazyFrames next to each other*

Description

Append multiple Data/LazyFrames next to each other

Usage

```
bind_cols_polars(..., .name_repair = "unique")
```

Arguments

`...` Polars DataFrames or LazyFrames to combine. Each argument can either be a Data/LazyFrame, or a list of Data/LazyFrames. Columns are matched by name. All Data/LazyFrames must have the same number of rows and there mustn't be duplicated column names.

`.name_repair` Can be "unique", "universal", "check_unique", "minimal". See [vctrs::vec_as_names\(\)](#) for the explanations for each value.

Examples

```

p1 <- polars::pl$DataFrame(
  x = sample(letters, 20),
  y = sample(1:100, 20)
)
p2 <- polars::pl$DataFrame(
  z = sample(letters, 20),
  w = sample(1:100, 20)
)

bind_cols_polars(p1, p2)
bind_cols_polars(list(p1, p2))

```

bind_rows_polars	<i>Stack multiple Data/LazyFrames on top of each other</i>
------------------	--

Description

Stack multiple Data/LazyFrames on top of each other

Usage

```
bind_rows_polars(..., .id = NULL)
```

Arguments

...	Polars DataFrames or LazyFrames to combine. Each argument can either be a Data/LazyFrame, or a list of Data/LazyFrames. Columns are matched by name, and any missing columns will be filled with NA.
.id	The name of an optional identifier column. Provide a string to create an output column that identifies each input. If all elements in ... are named, the identifier will use their names. Otherwise, it will be a simple count.

Examples

```

library(polars)
p1 <- pl$DataFrame(
  x = c("a", "b"),
  y = 1:2
)
p2 <- pl$DataFrame(
  y = 3:4,
  z = c("c", "d")
)$with_columns(pl$col("y")$cast(pl$Int16))

bind_rows_polars(p1, p2)

# this is equivalent

```

```

bind_rows_polars(list(p1, p2))

# create an id colum
bind_rows_polars(p1, p2, .id = "id")

# create an id colum with named elements
bind_rows_polars(p1 = p1, p2 = p2, .id = "id")

```

```
complete.RPolarsDataFrame
```

Complete a data frame with missing combinations of data

Description

Turns implicit missing values into explicit missing values. This is useful for completing missing combinations of data.

Usage

```

## S3 method for class 'RPolarsDataFrame'
complete(data, ..., fill = list())

## S3 method for class 'RPolarsLazyFrame'
complete(data, ..., fill = list())

```

Arguments

data	A Polars Data/LazyFrame
...	Any expression accepted by <code>dplyr::select()</code> : variable names, column numbers, select helpers, etc.
fill	A named list that for each variable supplies a single value to use instead of NA for missing combinations.

Examples

```

df <- polars::pl$DataFrame(
  group = c(1:2, 1, 2),
  item_id = c(1:2, 2, 3),
  item_name = c("a", "a", "b", "b"),
  value1 = c(1, NA, 3, 4),
  value2 = 4:7
)
df

df |> complete(group, item_id, item_name)

df |>

```

```

complete(
  group, item_id, item_name,
  fill = list(value1 = 0, value2 = 99)
)

df |>
  group_by(group, maintain_order = TRUE) |>
  complete(item_id, item_name)

```

```
compute.RPolarsLazyFrame
```

Collect a LazyFrame

Description

compute() checks the query, optimizes it in the background, and runs it. The output is a [Polars DataFrame](#). collect() is similar to compute() but converts the output to an R [data.frame](#), which consumes more memory.

Until tidypolars 0.7.0, there was only collect() and it was used to collect a LazyFrame into a Polars DataFrame. This usage is still valid for now but will change in 0.8.0 to automatically convert a DataFrame to a data.frame. Use compute() to have a Polars DataFrame as output.

Usage

```

## S3 method for class 'RPolarsLazyFrame'
compute(
  x,
  ...,
  type_coercion = TRUE,
  predicate_pushdown = TRUE,
  projection_pushdown = TRUE,
  simplify_expression = TRUE,
  slice_pushdown = TRUE,
  comm_subplan_elim = TRUE,
  comm_subexpr_elim = TRUE,
  cluster_with_columns = TRUE,
  no_optimization = FALSE,
  streaming = FALSE,
  collect_in_background = FALSE
)

## S3 method for class 'RPolarsLazyFrame'
collect(
  x,
  ...,
  type_coercion = TRUE,

```

```

predicate_pushdown = TRUE,
projection_pushdown = TRUE,
simplify_expression = TRUE,
slice_pushdown = TRUE,
comm_subplan_elim = TRUE,
comm_subexpr_elim = TRUE,
cluster_with_columns = TRUE,
no_optimization = FALSE,
streaming = FALSE,
collect_in_background = FALSE
)

```

Arguments

x	A Polars LazyFrame
...	Not used.
type_coercion	Coerce types such that operations succeed and run on minimal required memory (default is TRUE).
predicate_pushdown	Applies filters as early as possible at scan level (default is TRUE).
projection_pushdown	Select only the columns that are needed at the scan level (default is TRUE).
simplify_expression	Various optimizations, such as constant folding and replacing expensive operations with faster alternatives (default is TRUE).
slice_pushdown	Only load the required slice from the scan. Don't materialize sliced outputs level. Don't materialize sliced outputs (default is TRUE).
comm_subplan_elim	Cache branching subplans that occur on self-joins or unions (default is TRUE).
comm_subexpr_elim	Cache common subexpressions (default is TRUE).
cluster_with_columns	Combine sequential independent calls to <code>\$with_columns()</code> .
no_optimization	Sets the following optimizations to FALSE: <code>predicate_pushdown</code> , <code>projection_pushdown</code> , <code>slice_pushdown</code> , <code>simplify_expression</code> . Default is FALSE.
streaming	Run parts of the query in a streaming fashion (this is in an alpha state). Default is FALSE.
collect_in_background	Detach this query from the R session. Computation will start in background. Get a handle which later can be converted into the resulting DataFrame. Useful in interactive mode to not lock R session (default is FALSE).

See Also

[fetch\(\)](#) for applying a lazy query on a subset of the data.

Examples

```

dat_lazy <- polars::pl$DataFrame(iris)$lazy()

compute(dat_lazy)

# you can build a query and add compute() as the last piece
dat_lazy |>
  select(starts_with("Sepal")) |>
  filter(between(Sepal.Length, 5, 6)) |>
  compute()

# call collect() instead to return a data.frame (note that this is more
# expensive than compute())
dat_lazy |>
  select(starts_with("Sepal")) |>
  filter(between(Sepal.Length, 5, 6)) |>
  collect()

```

count.RPolarsDataFrame

Count the observations in each group

Description

Count the observations in each group

Usage

```

## S3 method for class 'RPolarsDataFrame'
count(x, ..., sort = FALSE, name = "n")

## S3 method for class 'RPolarsLazyFrame'
count(x, ..., sort = FALSE, name = "n")

## S3 method for class 'RPolarsDataFrame'
add_count(x, ..., sort = FALSE, name = "n")

## S3 method for class 'RPolarsLazyFrame'
add_count(x, ..., sort = FALSE, name = "n")

```

Arguments

x	A Polars Data/LazyFrame
...	Any expression accepted by <code>dplyr::select()</code> : variable names, column numbers, select helpers, etc.
sort	If TRUE, will show the largest groups at the top.
name	Name of the new column.

Examples

```
test <- polars::pl$DataFrame(mtcars)
count(test, cyl)

count(test, cyl, am)

count(test, cyl, am, sort = TRUE, name = "count")

add_count(test, cyl, am, sort = TRUE, name = "count")
```

```
cross_join.RPolarsDataFrame
      Cross join
```

Description

Cross joins match each row in x to every row in y , resulting in a dataset with $nrow(x) * nrow(y)$ rows.

Usage

```
## S3 method for class 'RPolarsDataFrame'
cross_join(x, y, suffix = c(".x", ".y"), ...)

## S3 method for class 'RPolarsLazyFrame'
cross_join(x, y, suffix = c(".x", ".y"), ...)
```

Arguments

<code>x, y</code>	Two Polars Data/LazyFrames
<code>suffix</code>	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.
<code>...</code>	Not used.

Examples

```
test <- polars::pl$DataFrame(
  origin = c("ALG", "FRA", "GER"),
  year = c(2020, 2020, 2021)
)

test2 <- polars::pl$DataFrame(
  destination = c("USA", "JPN", "BRA"),
  language = c("english", "japanese", "portuguese")
)

test
```

```
test2
cross_join(test, test2)
```

describe	<i>Summary statistics for a Polars DataFrame</i>
----------	--

Description

Summary statistics for a Polars DataFrame

Usage

```
describe(.data, percentiles = c(0.25, 0.75))
```

Arguments

.data	A Polars Data/LazyFrame
percentiles	One or more percentiles to include in the summary statistics. All values must be between 0 and 1 (NULL are ignored).

Examples

```
mtcars |>
  as_polars_df() |>
  describe(percentiles = c(0.2, 0.4, 0.6, 0.8))
```

describe_plan	<i>Show the optimized and non-optimized query plans</i>
---------------	---

Description

These functions are available for LazyFrames only. `describe_plan()` shows the query plan as-is, without any optimization done by Polars. This is not the query performed. Before running the query, Polars applies a list of optimizations (such as filtering data before sorting it). The actual query plan ran by Polars is shown by `describe_optimized_plan()`. Note that the queries are read from bottom to top.

Usage

```
describe_plan(.data)

describe_optimized_plan(.data)
```

Arguments

`.data` A Polars LazyFrame

Examples

```
query <- mtcars |>
  as_polars_lf() |>
  arrange(drat) |>
  filter(cyl == 3) |>
  select(mpg)

describe_plan(query)

describe_optimized_plan(query)
```

distinct.RPolarsDataFrame

Remove or keep only duplicated rows in a Data/LazyFrame

Description

By default, duplicates are looked for in all variables. It is possible to specify a subset of variables where duplicates should be looked for. It is also possible to keep either the first occurrence, the last occurrence or remove all duplicates.

Usage

```
## S3 method for class 'RPolarsDataFrame'
distinct(.data, ..., keep = "first", maintain_order = TRUE)

## S3 method for class 'RPolarsLazyFrame'
distinct(.data, ..., keep = "first", maintain_order = TRUE)

duplicated_rows(.data, ...)
```

Arguments

`.data` A Polars Data/LazyFrame

`...` Any expression accepted by `dplyr::select()`: variable names, column numbers, select helpers, etc.

`keep` Either "first" (keep the first occurrence of the duplicated row), "last" (last occurrence) or "none" (remove all occurrences of duplicated rows).

`maintain_order` Maintain row order. This is the default but it can slow down the process with large datasets and it prevents the use of streaming.

Examples

```
pl_test <- polars::pl$DataFrame(
  iso_o = c(rep(c("AA", "AB"), each = 2), "AC", "DC"),
  iso_d = rep(c("BA", "BB", "BC"), each = 2),
  value = c(2, 2, 3, 4, 5, 6)
)

distinct(pl_test)
distinct(pl_test, iso_o)

duplicated_rows(pl_test)
duplicated_rows(pl_test, iso_o, iso_d)
```

```
drop_na.RPolarsDataFrame
```

Drop missing values

Description

By default, this will drop rows that contain any missing values. It is possible to specify a subset of variables so that only missing values in these variables will be considered.

Usage

```
## S3 method for class 'RPolarsDataFrame'
drop_na(data, ...)

## S3 method for class 'RPolarsLazyFrame'
drop_na(data, ...)
```

Arguments

data	A Polars Data/LazyFrame
...	Any expression accepted by <code>dplyr::select()</code> : variable names, column numbers, select helpers, etc.

Examples

```
tmp <- mtcars
tmp[1:3, "mpg"] <- NA
tmp[4, "hp"] <- NA
pl_tmp <- polars::pl$DataFrame(tmp)

drop_na(pl_tmp)
drop_na(pl_tmp, hp, mpg)
```

 fetch

Fetch n rows of a LazyFrame

Description

Fetch is a way to collect only the first `n` rows of a `LazyFrame`. It is mainly used to test that a query runs as expected on a subset of the data before using `collect()` on the full query. Note that fetching `n` rows doesn't mean that the output will actually contain `n` rows, see the section 'Details' for more information.

Usage

```
fetch(
  .data,
  n_rows = 500,
  type_coercion = TRUE,
  predicate_pushdown = TRUE,
  projection_pushdown = TRUE,
  simplify_expression = TRUE,
  slice_pushdown = TRUE,
  comm_subplan_elim = TRUE,
  comm_subexpr_elim = TRUE,
  cluster_with_columns = TRUE,
  no_optimization = FALSE,
  streaming = FALSE
)
```

Arguments

<code>.data</code>	A Polars <code>LazyFrame</code>
<code>n_rows</code>	Number of rows to fetch.
<code>type_coercion</code>	Coerce types such that operations succeed and run on minimal required memory (default is <code>TRUE</code>).
<code>predicate_pushdown</code>	Applies filters as early as possible at scan level (default is <code>TRUE</code>).
<code>projection_pushdown</code>	Select only the columns that are needed at the scan level (default is <code>TRUE</code>).
<code>simplify_expression</code>	Various optimizations, such as constant folding and replacing expensive operations with faster alternatives (default is <code>TRUE</code>).
<code>slice_pushdown</code>	Only load the required slice from the scan. Don't materialize sliced outputs level. Don't materialize sliced outputs (default is <code>TRUE</code>).
<code>comm_subplan_elim</code>	Cache branching subplans that occur on self-joins or unions (default is <code>TRUE</code>).

<code>comm_subexpr_elim</code>	Cache common subexpressions (default is TRUE).
<code>cluster_with_columns</code>	Combine sequential independent calls to <code>\$with_columns()</code> .
<code>no_optimization</code>	Sets the following optimizations to FALSE: <code>predicate_pushdown</code> , <code>projection_pushdown</code> , <code>slice_pushdown</code> , <code>simplify_expression</code> . Default is FALSE.
<code>streaming</code>	Run parts of the query in a streaming fashion (this is in an alpha state). Default is FALSE.

Details

The parameter `n_rows` indicates how many rows from the LazyFrame should be used at the beginning of the query, but it doesn't guarantee that `n_rows` will be returned. For example, if the query contains a filter or join operations with other datasets, then the final number of rows can be lower than `n_rows`. On the other hand, appending some rows during the query can lead to an output that has more rows than `n_rows`.

See Also

[collect\(\)](#) for applying a lazy query on the full data.

Examples

```
dat_lazy <- polars::pl$DataFrame(iris)$lazy()

# this will return 30 rows
fetch(dat_lazy, 30)

# this will return less than 30 rows because there are less than 30 matches
# for this filter in the whole dataset
dat_lazy |>
  filter(Sepal.Length > 7.0) |>
  fetch(30)
```

`fill.RPolarsDataFrame` *Fill in missing values with previous or next value*

Description

Fills missing values in selected columns using the next or previous entry. This is useful in the common output format where values are not repeated, and are only recorded when they change.

Usage

```
## S3 method for class 'RPolarsDataFrame'
fill(data, ..., .direction = c("down", "up", "downup", "updown"))
```

Arguments

data	A Polars Data/LazyFrame
...	Any expression accepted by <code>dplyr::select()</code> : variable names, column numbers, select helpers, etc.
.direction	Direction in which to fill missing values. Either "down" (the default), "up", "downup" (i.e. first down and then up) or "updown" (first up and then down).

Details

With grouped Data/LazyFrames, `fill()` will be applied within each group, meaning that it won't fill across group boundaries.

Examples

```
pl_test <- polars::pl$DataFrame(x = c(NA, 1), y = c(2, NA))

fill(pl_test, everything(), .direction = "down")
fill(pl_test, everything(), .direction = "up")

# with grouped data, it doesn't use values from other groups
pl_grouped <- polars::pl$DataFrame(
  grp = rep(c("A", "B"), each = 3),
  x = c(1, NA, NA, NA, 2, NA),
  y = c(3, NA, 4, NA, 3, 1)
) |>
  group_by(grp)

fill(pl_grouped, x, y, .direction = "down")
```

```
filter.RPolarsDataFrame
```

Keep rows that match a condition

Description

This 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 'RPolarsDataFrame'
filter(.data, ..., .by = NULL)

## S3 method for class 'RPolarsLazyFrame'
filter(.data, ..., .by = NULL)
```


Arguments

<code>.data</code>	A Polars Data/LazyFrame
<code>...</code>	Expressions that return a logical value, and are defined in terms of the variables in the data. If multiple expressions are included, they will be combined with the <code>&</code> operator. Only rows for which all conditions evaluate to TRUE are kept.
<code>.by</code>	Optionally, a selection of columns to group by for just this operation, functioning as an alternative to <code>group_by()</code> . The group order is not maintained, use <code>group_by()</code> if you want more control over it.

Examples

```
pl_iris <- polars::pl$DataFrame(iris)

filter(pl_iris, Sepal.Length < 5, Species == "setosa")

filter(pl_iris, Sepal.Length < Sepal.Width + Petal.Length)

filter(pl_iris, Species == "setosa" | is.na(Species))

iris2 <- iris
iris2$Species <- as.character(iris2$Species)
iris2 |>
  as_polars_df() |>
  filter(Species %in% c("setosa", "virginica"))

# filter by group
pl_iris |>
  group_by(Species) |>
  filter(Sepal.Length == max(Sepal.Length)) |>
  ungroup()

# an alternative syntax for grouping is to use `.by`
pl_iris |>
  filter(Sepal.Length == max(Sepal.Length), .by = Species)
```

group_by.RPolarsDataFrame

Group by one or more variables

Description

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

Usage

```
## S3 method for class 'RPolarsDataFrame'
group_by(.data, ..., maintain_order = FALSE)

## S3 method for class 'RPolarsDataFrame'
ungroup(x, ...)

## S3 method for class 'RPolarsLazyFrame'
group_by(.data, ..., maintain_order = FALSE)

## S3 method for class 'RPolarsLazyFrame'
ungroup(x, ...)
```

Arguments

.data	A Polars Data/LazyFrame
...	Variables to group by (used in group_by() only). Not used in ungroup().
maintain_order	Maintain row order. For performance reasons, this is FALSE by default). Setting it to TRUE can slow down the process with large datasets and prevents the use of streaming.
x	A Polars Data/LazyFrame

Examples

```
by_cyl <- mtcars |>
  as_polars_df() |>
  group_by(cyl)

by_cyl

by_cyl |> summarise(
  disp = mean(disp),
  hp = mean(hp)
)
by_cyl |> filter(disp == max(disp))
```

group_split.RPolarsDataFrame
Grouping metadata

Description

group_vars() returns a character vector with the names of the grouping variables. group_keys() returns a data frame with one row per group.

Usage

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

Arguments

.tbl	A Polars Data/LazyFrame
...	If .tbl is not grouped, variables to group by. If .tbl is already grouped, this is ignored.
.keep	Should the grouping columns be kept?

Examples

```
pl_g <- polars::as_polars_df(iris) |>
  group_by(Species)

group_split(pl_g)
```

```
group_vars.RPolarsDataFrame
      Grouping metadata
```

Description

group_vars() returns a character vector with the names of the grouping variables. group_keys() returns a data frame with one row per group.

Usage

```
## S3 method for class 'RPolarsDataFrame'
group_vars(x)

## S3 method for class 'RPolarsLazyFrame'
group_vars(x)

## S3 method for class 'RPolarsDataFrame'
group_keys(.tbl, ...)

## S3 method for class 'RPolarsLazyFrame'
group_keys(.tbl, ...)
```

Arguments

x, .tbl	A Polars Data/LazyFrame
...	Not used.

Examples

```
pl_g <- polars::as_polars_df(mtcars) |>
  group_by(cyl, am)

group_vars(pl_g)

group_keys(pl_g)
```

left_join.RPolarsDataFrame
Mutating joins

Description

Mutating joins add columns from y to x, matching observations based on the keys.

Usage

```
## S3 method for class 'RPolarsDataFrame'
left_join(
  x,
  y,
  by = NULL,
  copy = NULL,
  suffix = c(".x", ".y"),
  ...,
  keep = NULL,
  na_matches = "na",
  relationship = NULL
)

## S3 method for class 'RPolarsDataFrame'
right_join(
  x,
  y,
  by = NULL,
  copy = NULL,
  suffix = c(".x", ".y"),
  ...,
  keep = NULL,
  na_matches = "na",
  relationship = NULL
)

## S3 method for class 'RPolarsDataFrame'
full_join(
```

```
x,  
y,  
by = NULL,  
copy = NULL,  
suffix = c(".x", ".y"),  
...,  
keep = NULL,  
na_matches = "na",  
relationship = NULL  
)  
  
## S3 method for class 'RPolarsDataFrame'  
inner_join(  
  x,  
  y,  
  by = NULL,  
  copy = NULL,  
  suffix = c(".x", ".y"),  
  ...,  
  keep = NULL,  
  na_matches = "na",  
  relationship = NULL  
)  
  
## S3 method for class 'RPolarsLazyFrame'  
left_join(  
  x,  
  y,  
  by = NULL,  
  copy = NULL,  
  suffix = c(".x", ".y"),  
  ...,  
  keep = NULL,  
  na_matches = "na",  
  relationship = NULL  
)  
  
## S3 method for class 'RPolarsLazyFrame'  
right_join(  
  x,  
  y,  
  by = NULL,  
  copy = NULL,  
  suffix = c(".x", ".y"),  
  ...,  
  keep = NULL,  
  na_matches = "na",  
  relationship = NULL
```

```

)

## S3 method for class 'RPolarsLazyFrame'
full_join(
  x,
  y,
  by = NULL,
  copy = NULL,
  suffix = c(".x", ".y"),
  ...,
  keep = NULL,
  na_matches = "na",
  relationship = NULL
)

## S3 method for class 'RPolarsLazyFrame'
inner_join(
  x,
  y,
  by = NULL,
  copy = NULL,
  suffix = c(".x", ".y"),
  ...,
  keep = NULL,
  na_matches = "na",
  relationship = NULL
)

```

Arguments

x, y	Two Polars Data/LazyFrames
by	<p>Variables to join by. If NULL, the default, *_join() will perform a natural join, using all variables in common across x and y. A message lists the variables so that you can check they're correct; suppress the message by supplying by explicitly.</p> <p>by can take a character vector, like c("x", "y") if x and y are in both datasets. To join on variables that don't have the same name, use equalities in the character vector, like c("x1" = "x2", "y"). If you use a character vector, the join can only be done using strict equality.</p> <p>Finally, by can be a specification created by dplyr::join_by(). Contrary to the input as character vector shown above, join_by() uses unquoted column names, e.g join_by(x1 == x2, y). It also uses equality and inequality operators ==, > and similar. For now, only equality operators are supported.</p>
copy, keep	Not used.
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.
...	Not used.

na_matches	<p>Should two NA values match?</p> <ul style="list-style-type: none"> • "na", the default, treats two NA values as equal. • "never" treats two NA values as different and will never match them together or to any other values. <p>Note that when joining Polars Data/LazyFrames, NaN are always considered equal, no matter the value of na_matches. This differs from the original dplyr implementation.</p>
relationship	<p>Handling of the expected relationship between the keys of x and y. Must be one of the following:</p> <ul style="list-style-type: none"> • NULL, the default, is equivalent to "many-to-many". It doesn't expect any relationship between x and y. • "one-to-one" expects each row in x to match at most 1 row in y and each row in y to match at most 1 row in x. • "one-to-many" expects each row in y to match at most 1 row in x. • "many-to-one" expects each row in x matches at most 1 row in y.

Examples

```

test <- polars::pl$DataFrame(
  x = c(1, 2, 3),
  y1 = c(1, 2, 3),
  z = c(1, 2, 3)
)

test2 <- polars::pl$DataFrame(
  x = c(1, 2, 4),
  y2 = c(1, 2, 4),
  z2 = c(4, 5, 7)
)

test

test2

# default is to use common columns, here "x" only
left_join(test, test2)

# we can specify the columns on which to join with join_by()...
left_join(test, test2, by = join_by(x, y1 == y2))

# ... or with a character vector
left_join(test, test2, by = c("x", "y1" = "y2"))

# we can customize the suffix of common column names not used to join
test2 <- polars::pl$DataFrame(
  x = c(1, 2, 4),
  y1 = c(1, 2, 4),
  z = c(4, 5, 7)
)

```

```

left_join(test, test2, by = "x", suffix = c("_left", "_right"))

# the argument "relationship" ensures the join matches the expectation
country <- polars::pl$DataFrame(
  iso = c("FRA", "DEU"),
  value = 1:2
)
country

country_year <- polars::pl$DataFrame(
  iso = rep(c("FRA", "DEU"), each = 2),
  year = rep(2019:2020, 2),
  value2 = 3:6
)
country_year

# We expect that each row in "x" matches only one row in "y" but, it's not
# true as each row of "x" matches two rows of "y"
tryCatch(
  left_join(country, country_year, join_by(iso), relationship = "one-to-one"),
  error = function(e) e
)

# A correct expectation would be "one-to-many":
left_join(country, country_year, join_by(iso), relationship = "one-to-many")

```

make_unique_id

Create a column with unique id per row values

Description

Create a column with unique id per row values

Usage

```
make_unique_id(.data, ..., new_col = "hash")
```

Arguments

.data	A Polars Data/LazyFrame
...	Any expression accepted by <code>dplyr::select()</code> : variable names, column numbers, select helpers, etc.
new_col	Name of the new column

Examples

```
mtcars |>
  as_polars_df() |>
  make_unique_id(am, gear)
```

```
mutate.RPolarsDataFrame
```

Create, modify, and delete columns

Description

This 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 'RPolarsDataFrame'
mutate(.data, ..., .by = NULL, .keep = c("all", "used", "unused", "none"))
```

```
## S3 method for class 'RPolarsLazyFrame'
mutate(.data, ..., .by = NULL, .keep = c("all", "used", "unused", "none"))
```

Arguments

<code>.data</code>	A Polars Data/LazyFrame
<code>...</code>	Name-value pairs. The name gives the name of the column in the output. The value can be: <ul style="list-style-type: none"> • A vector the same length as the current group (or the whole data frame if ungrouped). • NULL, to remove the column. <p><code>across()</code> is mostly supported, except in a few cases. In particular, if the <code>.cols</code> argument is <code>where(...)</code>, it will <i>not</i> select variables that were created before <code>across()</code>. Other select helpers are supported. See the examples.</p>
<code>.by</code>	Optionally, a selection of columns to group by for just this operation, functioning as an alternative to <code>group_by()</code> . The group order is not maintained, use <code>group_by()</code> if you want more control over it.
<code>.keep</code>	Control which columns from <code>.data</code> are retained in the output. Grouping columns and columns created by <code>...</code> are always kept. <ul style="list-style-type: none"> • "all" retains all columns from <code>.data</code>. This is the default. • "used" retains only the columns used in <code>...</code> to create new columns. This is useful for checking your work, as it displays inputs and outputs side-by-side.

- "unused" retains only the columns not used in . . . to create new columns. This is useful if you generate new columns, but no longer need the columns used to generate them.
- "none" doesn't retain any extra columns from .data. Only the grouping variables and columns created by . . . are kept.

Details

A lot of functions available in base R (cos, mean, multiplying, etc.) or in other packages (dplyr::lag(), etc.) are implemented in an efficient way in Polars. These functions are automatically translated to Polars syntax under the hood so that you can continue using the classic R syntax and functions.

If a Polars built-in replacement doesn't exist (for example for custom functions), then tidypolars will throw an error. See the vignette on Polars expressions to know how to write custom functions that are accepted by tidypolars.

Examples

```
pl_iris <- polars::pl$DataFrame(iris)

# classic operation
mutate(pl_iris, x = Sepal.Width + Sepal.Length)

# logical operation
mutate(pl_iris, x = Sepal.Width > Sepal.Length & Petal.Width > Petal.Length)

# overwrite existing variable
mutate(pl_iris, Sepal.Width = Sepal.Width*2)

# grouped computation
pl_iris |>
  group_by(Species) |>
  mutate(foo = mean(Sepal.Length))

# an alternative syntax for grouping is to use `.by`
pl_iris |>
  mutate(foo = mean(Sepal.Length), .by = Species)

# across() is available
pl_iris |>
  mutate(
    across(.cols = contains("Sepal"), .fns = mean, .names = "{.fn}_of_{.col}")
  )
# It can receive several types of functions:
pl_iris |>
  mutate(
    across(
      .cols = contains("Sepal"),
      .fns = list(mean = mean, sd = ~ sd(.x)),
      .names = "{.fn}_of_{.col}"
    )
  )
```

```

# Be careful when using across(.cols = where(...), ...) as it will not include
# variables created in the same `...` (this is only the case for `where()`):
## Not run:
pl_iris |>
  mutate(
    foo = 1,
    across(
      .cols = where(is.numeric),
      \(x) x - 1000 # <<<<<<<<< this will not be applied on variable "foo"
    )
  )

## End(Not run)
# Warning message:
# In `across()`, the argument `.cols = where(is.numeric)` will not take into account
# variables created in the same `mutate()`/`summarize` call.

# Embracing an external variable works
some_value <- 1
mutate(pl_iris, x = {{ some_value }})

```

pivot_longer.RPolarsDataFrame

Pivot a Data/LazyFrame from wide to long

Description

Pivot a Data/LazyFrame from wide to long

Usage

```

## S3 method for class 'RPolarsDataFrame'
pivot_longer(
  data,
  cols,
  ...,
  names_to = "name",
  names_prefix = NULL,
  values_to = "value"
)

## S3 method for class 'RPolarsLazyFrame'
pivot_longer(
  data,
  cols,
  ...,
  names_to = "name",

```

```

  names_prefix = NULL,
  values_to = "value"
)

```

Arguments

data	A Polars Data/LazyFrame
cols	Columns to pivot into longer format. Can be anything accepted by <code>dplyr::select()</code> .
...	Not used.
names_to	The (quoted) name of the column that will contain the column names specified by cols.
names_prefix	A regular expression used to remove matching text from the start of each variable name.
values_to	A string specifying the name of the column to create from the data stored in cell values.

Examples

```

pl_relig_income <- polars::pl$DataFrame(tidyr::relig_income)
pl_relig_income

pl_relig_income |>
  pivot_longer(!religion, names_to = "income", values_to = "count")

pl_billboard <- polars::pl$DataFrame(tidyr::billboard)
pl_billboard

pl_billboard |>
  pivot_longer(
    cols = starts_with("wk"),
    names_to = "week",
    names_prefix = "wk",
    values_to = "rank",
  )

```

pivot_wider.RPolarsDataFrame

Pivot a DataFrame from long to wide

Description

Pivot a DataFrame from long to wide

Usage

```
## S3 method for class 'RPolarsDataFrame'
pivot_wider(
  data,
  ...,
  id_cols = NULL,
  names_from = name,
  values_from = value,
  names_prefix = "",
  names_sep = "_",
  names_glue = NULL,
  values_fill = NULL
)
```

Arguments

<code>data</code>	A Polars DataFrame (LazyFrames are not supported).
<code>...</code>	Not used.
<code>id_cols</code>	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> .
<code>names_from</code>	The (quoted or unquoted) column names whose values will be used for the names of the new columns.
<code>values_from</code>	The (quoted or unquoted) column names whose values will be used to fill the new columns.
<code>names_prefix</code>	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.
<code>names_sep</code>	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.
<code>names_glue</code>	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 to create custom column names.
<code>values_fill</code>	A scalar that will be used to replace missing values in the new columns. Note that the type of this value will be applied to new columns. For example, if you provide a character value to fill numeric columns, then all these columns will be converted to character.

Examples

```
pl_fish_encounters <- polars::pl$DataFrame(tidyr::fish_encounters)

pl_fish_encounters |>
  pivot_wider(names_from = station, values_from = seen)
```

```

pl_fish_encounters |>
  pivot_wider(names_from = station, values_from = seen, values_fill = 0)

# be careful about the type of the replacement value!
pl_fish_encounters |>
  pivot_wider(names_from = station, values_from = seen, values_fill = "a")

# using "names_glue" to specify the names of new columns
production <- expand_grid(
  product = c("A", "B"),
  country = c("AI", "EI"),
  year = 2000:2014
) |>
  filter((product == "A" & country == "AI") | product == "B") |>
  mutate(production = 1:45) |>
  as_polars_df()

production

production |>
  pivot_wider(
    names_from = c(product, country),
    values_from = production,
    names_glue = "prod_{product}_{country}"
  )

```

pull.RPolarsDataFrame *Extract a variable of a Data/LazyFrame*

Description

This returns an R vector and not a Polars Series.

Usage

```
## S3 method for class 'RPolarsDataFrame'
pull(.data, var, ...)
```

```
## S3 method for class 'RPolarsLazyFrame'
pull(.data, var, ...)
```

Arguments

.data	A Polars Data/LazyFrame
var	A quoted or unquoted variable name, or a variable index.
...	Not used.

Examples

```
pl_test <- as_polars_df(iris)
pull(pl_test, Sepal.Length)
pull(pl_test, "Sepal.Length")
```

```
relocate.RPolarsDataFrame
      Change column order
```

Description

Use `relocate()` to change column positions, using the same syntax as `select()` to make it easy to move blocks of columns at once.

Usage

```
## S3 method for class 'RPolarsDataFrame'
relocate(.data, ..., .before = NULL, .after = NULL)

## S3 method for class 'RPolarsLazyFrame'
relocate(.data, ..., .before = NULL, .after = NULL)
```

Arguments

<code>.data</code>	A Polars Data/LazyFrame
<code>...</code>	Any expression accepted by <code>dplyr::select()</code> : variable names, column numbers, select helpers, etc.
<code>.before</code> , <code>.after</code>	Column name (either quoted or unquoted) that indicates the destination of columns selected by <code>...</code> . Supplying neither will move columns to the left-hand side; specifying both is an error.

Examples

```
dat <- as_polars_df(mtcars)

dat |>
  relocate(hp, vs, .before = cyl)

# if .before and .after are not specified, selected columns are moved to the
# first positions
dat |>
  relocate(hp, vs)

# .before and .after can be quoted or unquoted
dat |>
  relocate(hp, vs, .after = "gear")
```

```
# select helpers are also available
dat |>
  relocate(contains("[aeiou]"))

dat |>
  relocate(hp, vs, .after = last_col())
```

```
rename.RPolarsDataFrame
```

```
Rename columns
```

Description

Rename columns

Usage

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

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

## S3 method for class 'RPolarsDataFrame'
rename_with(.data, .fn, .cols = tidyselect::everything(), ...)

## S3 method for class 'RPolarsLazyFrame'
rename_with(.data, .fn, .cols = tidyselect::everything(), ...)
```

Arguments

<code>.data</code>	A Polars Data/LazyFrame
<code>...</code>	For <code>rename()</code> , use <code>new_name = old_name</code> to rename selected variables. It is also possible to use quotation marks, e.g <code>"new_name" = "old_name"</code> . For <code>rename_with</code> , additional arguments passed to <code>fn</code> .
<code>.fn</code>	Function to apply on column names
<code>.cols</code>	Columns on which to apply <code>fn</code> . Can be anything accepted by <code>dplyr::select()</code> .

Examples

```
pl_test <- polars::pl$DataFrame(mtcars)

rename(pl_test, miles_per_gallon = mpg, horsepower = "hp")

rename(pl_test, `Miles per gallon` = "mpg", `Horse power` = "hp")
```



```

rename_with(pl_test, toupper, .cols = contains("p"))

pl_test_2 <- polars::pl$DataFrame(iris)

rename_with(pl_test_2, function(x) tolower(gsub(".", "_", x, fixed = TRUE)))

rename_with(pl_test_2, \(x) tolower(gsub(".", "_", x, fixed = TRUE)))

```

```
replace_na.RPolarsDataFrame
```

Replace NAs with specified values

Description

Replace NAs with specified values

Usage

```
## S3 method for class 'RPolarsDataFrame'
replace_na(data, replace, ...)
```

```
## S3 method for class 'RPolarsLazyFrame'
replace_na(data, replace, ...)
```

Arguments

data	A Polars Data/LazyFrame
replace	Either a scalar that will be used to replace NA in all columns, or a named list with the column name and the value that will be used to replace NA in it. The column type will be automatically converted to the type of the replacement value.
...	Not used.

Examples

```

pl_test <- polars::pl$DataFrame(x = c(NA, 1), y = c(2, NA))

# replace all NA with 0
replace_na(pl_test, 0)

# custom replacement per column
replace_na(pl_test, list(x = 0, y = 999))

# be careful to use the same type for the replacement and for the column!
replace_na(pl_test, list(x = "a", y = "unknown"))

```

`rowwise.RPolarsDataFrame`*Group input by rows*

Description

[EXPERIMENTAL]

`rowwise()` allows you to compute on a `Data/LazyFrame` a row-at-a-time. This is most useful when a vectorised function doesn't exist. `rowwise()` produces another type of grouped data, and therefore can be removed with `ungroup()`.

Usage

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

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

Arguments

<code>data</code>	A Polars Data/LazyFrame
<code>...</code>	Any expression accepted by <code>dplyr::select()</code> : variable names, column numbers, select helpers, etc.

Value

A Polars Data/LazyFrame.

Examples

```
df <- polars::pl$DataFrame(x = c(1, 3, 4), y = c(2, 1, 5), z = c(2, 3, 1))  
  
# Compute the mean of x, y, z in each row  
df |>  
  rowwise() |>  
  mutate(m = mean(c(x, y, z)))  
  
# Compute the min and max of x and y in each row  
df |>  
  rowwise() |>  
  mutate(min = min(c(x, y)), max = max(c(x, y)))
```

```
select.RPolarsDataFrame
```

Select columns from a Data/LazyFrame

Description

Select columns from a Data/LazyFrame

Usage

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

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

Arguments

<code>.data</code>	A Polars Data/LazyFrame
<code>...</code>	Any expression accepted by <code>dplyr::select()</code> : variable names, column numbers, select helpers, etc. Renaming is also possible.

Examples

```
pl_iris <- polars::pl$DataFrame(iris)

select(pl_iris, c("Sepal.Length", "Sepal.Width"))
select(pl_iris, Sepal.Length, Sepal.Width)
select(pl_iris, 1:3)
select(pl_iris, starts_with("Sepal"))
select(pl_iris, -ends_with("Length"))

# Renaming while selecting is also possible
select(pl_iris, foo1 = Sepal.Length, Sepal.Width)
```

```
semi_join.RPolarsDataFrame
```

Filtering joins

Description

Filtering joins filter rows from `x` based on the presence or absence of matches in `y`:

- `semi_join()` return all rows from `x` with a match in `y`.
- `anti_join()` return all rows from `x` without a match in `y`.

Usage

```
## S3 method for class 'RPolarsDataFrame'
semi_join(x, y, by = NULL, ..., na_matches = "na")

## S3 method for class 'RPolarsDataFrame'
anti_join(x, y, by = NULL, ..., na_matches = "na")

## S3 method for class 'RPolarsLazyFrame'
semi_join(x, y, by = NULL, ..., na_matches = "na")

## S3 method for class 'RPolarsLazyFrame'
anti_join(x, y, by = NULL, ..., na_matches = "na")
```

Arguments

x, y	Two Polars Data/LazyFrames
by	Variables to join by. If NULL, the default, *_join() will perform a natural join, using all variables in common across x and y. A message lists the variables so that you can check they're correct; suppress the message by supplying by explicitly. by can take a character vector, like c("x", "y") if x and y are in both datasets. To join on variables that don't have the same name, use equalities in the character vector, like c("x1" = "x2", "y"). If you use a character vector, the join can only be done using strict equality. Finally, by can be a specification created by dplyr::join_by(). Contrary to the input as character vector shown above, join_by() uses unquoted column names, e.g join_by(x1 == x2, y). It also uses equality and inequality operators ==, > and similar. For now, only equality operators are supported.
...	Not used.
na_matches	Should two NA values match? <ul style="list-style-type: none"> • "na", the default, treats two NA values as equal. • "never" treats two NA values as different and will never match them together or to any other values.

Note that when joining Polars Data/LazyFrames, NaN are always considered equal, no matter the value of na_matches. This differs from the original dplyr implementation.

Examples

```
test <- polars::pl$DataFrame(
  x = c(1, 2, 3),
  y = c(1, 2, 3),
  z = c(1, 2, 3)
)

test2 <- polars::pl$DataFrame(
  x = c(1, 2, 4),
```

```

  y = c(1, 2, 4),
  z2 = c(1, 2, 4)
)

test

test2

# only keep the rows of `test` that have matching keys in `test2`
semi_join(test, test2, by = c("x", "y"))

# only keep the rows of `test` that don't have matching keys in `test2`
anti_join(test, test2, by = c("x", "y"))

```

```
separate.RPolarsDataFrame
```

Separate a character column into multiple columns based on a substring

Description

Currently, splitting a column on a regular expression or position is not possible.

Usage

```
## S3 method for class 'RPolarsDataFrame'
separate(data, col, into, sep = " ", remove = TRUE, ...)
```

```
## S3 method for class 'RPolarsLazyFrame'
separate(data, col, into, sep = " ", remove = TRUE, ...)
```

Arguments

data	A Polars Data/LazyFrame
col	Column to split
into	Character vector containing the names of new variables to create. Use NA to omit the variable in the output.
sep	String that is used to split the column. Regular expressions are not supported yet.
remove	If TRUE, remove input column from output data frame.
...	Not used.

Examples

```
test <- polars::pl$DataFrame(
  x = c(NA, "x.y", "x.z", "y.z")
)
separate(test, x, into = c("foo", "foo2"), sep = ".")
```

sink_csv

Stream output to a CSV file

Description

This function allows to stream a LazyFrame that is larger than RAM directly to a .csv file without collecting it in the R session, thus preventing crashes because of too small memory.

Usage

```
sink_csv(
  .data,
  path,
  include_bom = FALSE,
  include_header = TRUE,
  separator = ",",
  line_terminator = "\n",
  quote = "\"",
  batch_size = 1024,
  datetime_format = NULL,
  date_format = NULL,
  time_format = NULL,
  float_precision = NULL,
  null_values = "",
  quote_style = "necessary",
  maintain_order = TRUE,
  type_coercion = TRUE,
  predicate_pushdown = TRUE,
  projection_pushdown = TRUE,
  simplify_expression = TRUE,
  slice_pushdown = TRUE,
  no_optimization = FALSE,
  inherit_optimization = FALSE
)
```

Arguments

.data	A Polars LazyFrame.
path	Output file (must be a .csv file).

include_bom	Whether to include UTF-8 BOM (byte order mark) in the CSV output.
include_header	Whether to include header in the CSV output.
separator	Separate CSV fields with this symbol.
line_terminator	String used to end each row.
quote	Byte to use as quoting character.
batch_size	Number of rows that will be processed per thread.
datetime_format, date_format, time_format	A format string used to format date and time values. See <code>strftime()</code> for accepted values. If no format specified, the default fractional-second precision is inferred from the maximum time unit found in the <code>Datetime</code> cols (if any).
float_precision	Number of decimal places to write, applied to both <code>Float32</code> and <code>Float64</code> datatypes.
null_values	A string representing null values (defaulting to the empty string).
quote_style	Determines the quoting strategy used: <ul style="list-style-type: none"> • "necessary" (default): This puts quotes around fields only when necessary. They are necessary when fields contain a quote, delimiter or record terminator. Quotes are also necessary when writing an empty record (which is indistinguishable from a record with one empty field). • "always": This puts quotes around every field. • "non_numeric": This puts quotes around all fields that are non-numeric. Namely, when writing a field that does not parse as a valid float or integer, then quotes will be used even if they aren't strictly necessary.
maintain_order	Whether maintain the order the data was processed (default is <code>TRUE</code>). Setting this to <code>FALSE</code> will be slightly faster.
type_coercion	Coerce types such that operations succeed and run on minimal required memory (default is <code>TRUE</code>).
predicate_pushdown	Applies filters as early as possible at scan level (default is <code>TRUE</code>).
projection_pushdown	Select only the columns that are needed at the scan level (default is <code>TRUE</code>).
simplify_expression	Various optimizations, such as constant folding and replacing expensive operations with faster alternatives (default is <code>TRUE</code>).
slice_pushdown	Only load the required slice from the scan. Don't materialize sliced outputs level. Don't materialize sliced outputs (default is <code>TRUE</code>).
no_optimization	Sets the following optimizations to <code>FALSE</code> : <code>predicate_pushdown</code> , <code>projection_pushdown</code> , <code>slice_pushdown</code> , <code>simplify_expression</code> . Default is <code>FALSE</code> .
inherit_optimization	Use existing optimization settings regardless of the settings specified in this function call. Default is <code>FALSE</code> .

Value

Writes a .csv file with the content of the LazyFrame.

Examples

```
## Not run:
# This is an example workflow where sink_csv() is not very useful because
# the data would fit in memory. It simply is an example of using it at the
# end of a piped workflow.

# Create files for the CSV input and output:
file_csv <- tempfile(fileext = ".csv")
file_csv2 <- tempfile(fileext = ".csv")

# Write some data in a CSV file
fake_data <- do.call("rbind", rep(list(mtcars), 1000))
write.csv(fake_data, file = file_csv)

# In a new R session, we could read this file as a LazyFrame, do some operations,
# and write it to another CSV file without ever collecting it in the R session:
polars::pl$scan_csv(file_csv) |>
  filter(cyl %in% c(4, 6), mpg > 22) |>
  mutate(
    hp_gear_ratio = hp / gear
  ) |>
  sink_csv(path = file_csv2)

## End(Not run)
```

sink_parquet

Stream output to a parquet file

Description

This function allows to stream a LazyFrame that is larger than RAM directly to a .parquet file without collecting it in the R session, thus preventing crashes because of too small memory.

Usage

```
sink_parquet(
  .data,
  path,
  compression = "zstd",
  compression_level = 3,
  statistics = FALSE,
  row_group_size = NULL,
  data_pagesize_limit = NULL,
```



```

maintain_order = TRUE,
type_coercion = TRUE,
predicate_pushdown = TRUE,
projection_pushdown = TRUE,
simplify_expression = TRUE,
slice_pushdown = TRUE,
no_optimization = FALSE,
inherit_optimization = FALSE
)

```

Arguments

<code>.data</code>	A Polars LazyFrame.
<code>path</code>	Output file (must be a <code>.parquet</code> file).
<code>compression</code>	The compression method. One of : <ul style="list-style-type: none"> • "uncompressed" • "zstd" (default): good compression performance • "lz4": fast compression / decompression • "snappy": more backwards compatibility guarantees when you deal with older parquet readers. • "gzip", "lzo", "brotli"
<code>compression_level</code>	The level of compression to use (default is 3). Only used if <code>compression</code> is one of "gzip", "brotli", or "zstd". Higher compression means smaller files on disk. <ul style="list-style-type: none"> • "gzip" : min-level = 0, max-level = 10 • "brotli" : min-level = 0, max-level = 11 • "zstd" : min-level = 1, max-level = 22.
<code>statistics</code>	Whether to compute and write column statistics (default is FALSE). This requires more computations.
<code>row_group_size</code>	Size of the row groups in number of rows. If NULL (default), the chunks of the DataFrame are used. Writing in smaller chunks may reduce memory pressure and improve writing speeds.
<code>data_pagesize_limit</code>	If NULL (default), the limit will be around 1MB.
<code>maintain_order</code>	Whether maintain the order the data was processed (default is TRUE). Setting this to FALSE will be slightly faster.
<code>type_coercion</code>	Coerce types such that operations succeed and run on minimal required memory (default is TRUE).
<code>predicate_pushdown</code>	Applies filters as early as possible at scan level (default is TRUE).
<code>projection_pushdown</code>	Select only the columns that are needed at the scan level (default is TRUE).
<code>simplify_expression</code>	Various optimizations, such as constant folding and replacing expensive operations with faster alternatives (default is TRUE).

`slice_pushdown` Only load the required slice from the scan. Don't materialize sliced outputs level. Don't materialize sliced outputs (default is TRUE).

`no_optimization` Sets the following optimizations to FALSE: `predicate_pushdown`, `projection_pushdown`, `slice_pushdown`, `simplify_expression`. Default is FALSE.

`inherit_optimization` Use existing optimization settings regardless of the settings specified in this function call. Default is FALSE.

Value

Writes a .parquet file with the content of the LazyFrame.

Examples

```
## Not run:
# This is an example workflow where sink_parquet() is not very useful because
# the data would fit in memory. It simply is an example of using it at the
# end of a piped workflow.

# Create files for the CSV input and the Parquet output:
file_csv <- tempfile(fileext = ".csv")
file_parquet <- tempfile(fileext = ".parquet")

# Write some data in a CSV file
fake_data <- do.call("rbind", rep(list(mtcars), 1000))
write.csv(fake_data, file = file_csv)

# In a new R session, we could read this file as a LazyFrame, do some operations,
# and write it to a parquet file without ever collecting it in the R session:
polars::pl$scan_csv(file_csv) |>
  filter(cyl %in% c(4, 6), mpg > 22) |>
  mutate(
    hp_gear_ratio = hp / gear
  ) |>
  sink_parquet(path = file_parquet)

## End(Not run)
```

slice_tail.RPolarsDataFrame

Subset rows of a Data/LazyFrame

Description

Subset rows of a Data/LazyFrame

Usage

```
## S3 method for class 'RPolarsDataFrame'
slice_tail(.data, ..., n, by = NULL)

## S3 method for class 'RPolarsLazyFrame'
slice_tail(.data, ..., n, by = NULL)

## S3 method for class 'RPolarsDataFrame'
slice_head(.data, ..., n, by = NULL)

## S3 method for class 'RPolarsLazyFrame'
slice_head(.data, ..., n, by = NULL)

## S3 method for class 'RPolarsDataFrame'
slice_sample(.data, ..., n = NULL, prop = NULL, replace = FALSE, by = NULL)
```

Arguments

.data	A Polars Data/LazyFrame
...	Not used.
n	The number of rows to select from the start or the end of the data. Cannot be used with prop.
by	Optionally, a selection of columns to group by for just this operation, functioning as an alternative to group_by(). The group order is not maintained, use group_by() if you want more control over it.
prop	Proportion of rows to select. Cannot be used with n.
replace	Perform the sampling with replacement (TRUE) or without (FALSE).

Examples

```
pl_test <- polars::pl$DataFrame(iris)
slice_head(pl_test, n = 3)
slice_tail(pl_test, n = 3)
slice_sample(pl_test, n = 5)
slice_sample(pl_test, prop = 0.1)
```

summarize.RPolarsDataFrame

Summarize each group down to one row

Description

summarize() returns one row for each combination of grouping variables (one difference with dplyr::summarize() is that summarize() only accepts grouped data). It will contain one column for each grouping variable and one column for each of the summary statistics that you have specified.

Usage

```
## S3 method for class 'RPolarsDataFrame'
summarize(.data, ..., .by = NULL)

## S3 method for class 'RPolarsDataFrame'
summarise(.data, ..., .by = NULL)

## S3 method for class 'RPolarsLazyFrame'
summarize(.data, ..., .by = NULL)

## S3 method for class 'RPolarsLazyFrame'
summarise(.data, ..., .by = NULL)
```

Arguments

<code>.data</code>	A Polars Data/LazyFrame
<code>...</code>	Name-value pairs. The name gives the name of the column in the output. The value can be: <ul style="list-style-type: none"> • A vector the same length as the current group (or the whole data frame if ungrouped). • <code>NULL</code>, to remove the column. <p><code>across()</code> is mostly supported, except in a few cases. In particular, if the <code>.cols</code> argument is <code>where(...)</code>, it will <i>not</i> select variables that were created before <code>across()</code>. Other select helpers are supported. See the examples.</p>
<code>.by</code>	Optionally, a selection of columns to group by for just this operation, functioning as an alternative to <code>group_by()</code> . The group order is not maintained, use <code>group_by()</code> if you want more control over it.

Examples

```
mtcars |>
  as_polars_df() |>
  group_by(cyl) |>
  summarize(m_gear = mean(gear), sd_gear = sd(gear))

# an alternative syntax is to use `by`
mtcars |>
  as_polars_df() |>
  summarize(m_gear = mean(gear), sd_gear = sd(gear), .by = cyl)
```

Description

There is currently one global option:

- `tidypolars_unknown_args` controls what happens when some arguments passed in an expression are unknown, e.g the argument `prob` in `sample()`. The default ("warn") only warns the user that some arguments are ignored by `tidypolars`. The only other accepted value is "error" to throw an error when this happens.

Examples

```
options(tidypolars_unknown_args = "warn")
test <- polars::pl$DataFrame(x = c(2, 1, 5, 3, 1))

# The default is to warn the user
mutate(test, x2 = sample(x, prob = 0.5))

# But one can make this stricter and throw an error when this happens
options(tidypolars_unknown_args = "error")
try(mutate(test, x2 = sample(x, prob = 0.5)))

options(tidypolars_unknown_args = "warn")
```

uncount.RPolarsDataFrame

Uncount a Data/LazyFrame

Description

This duplicates rows according to a weighting variable (or expression). This is the opposite of `count()`.

Usage

```
## S3 method for class 'RPolarsDataFrame'
uncount(data, weights, ..., .remove = TRUE, .id = NULL)

## S3 method for class 'RPolarsLazyFrame'
uncount(data, weights, ..., .remove = TRUE, .id = NULL)
```

Arguments

<code>data</code>	A Polars Data/LazyFrame
<code>weights</code>	A vector of weights. Evaluated in the context of data.
<code>...</code>	Not used.
<code>.remove</code>	If TRUE, and <code>weights</code> is the name of a column in <code>data</code> , then this column is removed.
<code>.id</code>	Supply a string to create a new variable which gives a unique identifier for each created row.

Examples

```

test <- polars::pl$DataFrame(x = c("a", "b"), y = 100:101, n = c(1, 2))
test

uncount(test, n)

uncount(test, n, .id = "id")

# using constants
uncount(test, 2)

# using expressions
uncount(test, 2 / n)

```

```
unite.RPolarsDataFrame
```

Unite multiple columns into one by pasting strings together

Description

Unite multiple columns into one by pasting strings together

Usage

```

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

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

```

Arguments

data	A Polars Data/LazyFrame
col	The name of the new column, as a string or symbol.
...	Any expression accepted by <code>dplyr::select()</code> : variable names, column numbers, select helpers, etc.
sep	Separator to use between values.
remove	If TRUE, remove input columns from the output Data/LazyFrame.
na.rm	If TRUE, missing values will be replaced with an empty string prior to uniting each value.

Examples

```
test <- polars::pl$DataFrame(  
  year = 2009:2011,  
  month = 10:12,  
  day = c(11L, 22L, 28L),  
  name_day = c("Monday", "Thursday", "Wednesday")  
)  
  
# By default, united columns are dropped  
unite(test, col = "full_date", year, month, day, sep = "-")  
unite(test, col = "full_date", year, month, day, sep = "-", remove = FALSE)  
  
test2 <- polars::pl$DataFrame(  
  name = c("John", "Jack", "Thomas"),  
  middlename = c("T.", NA, "F."),  
  surname = c("Smith", "Thompson", "Jones")  
)  
  
# By default, NA values are kept in the character output  
unite(test2, col = "full_name", everything(), sep = " ")  
unite(test2, col = "full_name", everything(), sep = " ", na.rm = TRUE)
```

Index

add_count.RPolarsDataFrame
(count.RPolarsDataFrame), 9

add_count.RPolarsLazyFrame
(count.RPolarsDataFrame), 9

anti_join.RPolarsDataFrame
(semi_join.RPolarsDataFrame),
35

anti_join.RPolarsLazyFrame
(semi_join.RPolarsDataFrame),
35

arrange.RPolarsDataFrame, 3

as_tibble.tidypolars, 3

bind_cols_polars, 4

bind_rows_polars, 5

collect(), 15

collect.RPolarsLazyFrame
(compute.RPolarsLazyFrame), 7

complete.RPolarsDataFrame, 6

complete.RPolarsLazyFrame
(complete.RPolarsDataFrame), 6

compute.RPolarsLazyFrame, 7

count.RPolarsDataFrame, 9

count.RPolarsLazyFrame
(count.RPolarsDataFrame), 9

cross_join.RPolarsDataFrame, 10

cross_join.RPolarsLazyFrame
(cross_join.RPolarsDataFrame),
10

data.frame, 7

DataFrame, 3

describe, 11

describe_optimized_plan
(describe_plan), 11

describe_plan, 11

distinct.RPolarsDataFrame, 12

distinct.RPolarsLazyFrame
(distinct.RPolarsDataFrame), 12

drop_na.RPolarsDataFrame, 13

drop_na.RPolarsLazyFrame
(drop_na.RPolarsDataFrame), 13

duplicated_rows
(distinct.RPolarsDataFrame), 12

fetch, 14

fetch(), 8

fill.RPolarsDataFrame, 15

filter.RPolarsDataFrame, 16

filter.RPolarsLazyFrame
(filter.RPolarsDataFrame), 16

full_join.RPolarsDataFrame
(left_join.RPolarsDataFrame),
20

full_join.RPolarsLazyFrame
(left_join.RPolarsDataFrame),
20

group_by.RPolarsDataFrame, 17

group_by.RPolarsLazyFrame
(group_by.RPolarsDataFrame), 17

group_keys.RPolarsDataFrame
(group_vars.RPolarsDataFrame),
19

group_keys.RPolarsLazyFrame
(group_vars.RPolarsDataFrame),
19

group_split.RPolarsDataFrame, 18

group_vars.RPolarsDataFrame, 19

group_vars.RPolarsLazyFrame
(group_vars.RPolarsDataFrame),
19

inner_join.RPolarsDataFrame
(left_join.RPolarsDataFrame),
20

inner_join.RPolarsLazyFrame
(left_join.RPolarsDataFrame),
20

LazyFrame, [3](#)
 left_join.RPolarsDataFrame, [20](#)
 left_join.RPolarsLazyFrame
 (left_join.RPolarsDataFrame),
 [20](#)

 make_unique_id, [24](#)
 mutate.RPolarsDataFrame, [25](#)
 mutate.RPolarsLazyFrame
 (mutate.RPolarsDataFrame), [25](#)

 pivot_longer.RPolarsDataFrame, [27](#)
 pivot_longer.RPolarsLazyFrame
 (pivot_longer.RPolarsDataFrame),
 [27](#)
 pivot_wider.RPolarsDataFrame, [28](#)
 Polars DataFrame, [7](#)
 polars::as.data.frame.RPolarsDataFrame,
 [4](#)
 pull.RPolarsDataFrame, [30](#)
 pull.RPolarsLazyFrame
 (pull.RPolarsDataFrame), [30](#)

 relocate.RPolarsDataFrame, [31](#)
 relocate.RPolarsLazyFrame
 (relocate.RPolarsDataFrame), [31](#)
 rename.RPolarsDataFrame, [32](#)
 rename.RPolarsLazyFrame
 (rename.RPolarsDataFrame), [32](#)
 rename_with.RPolarsDataFrame
 (rename.RPolarsDataFrame), [32](#)
 rename_with.RPolarsLazyFrame
 (rename.RPolarsDataFrame), [32](#)
 replace_na.RPolarsDataFrame, [33](#)
 replace_na.RPolarsLazyFrame
 (replace_na.RPolarsDataFrame),
 [33](#)
 right_join.RPolarsDataFrame
 (left_join.RPolarsDataFrame),
 [20](#)
 right_join.RPolarsLazyFrame
 (left_join.RPolarsDataFrame),
 [20](#)
 rowwise.RPolarsDataFrame, [34](#)
 rowwise.RPolarsLazyFrame
 (rowwise.RPolarsDataFrame), [34](#)

 select.RPolarsDataFrame, [35](#)
 select.RPolarsLazyFrame
 (select.RPolarsDataFrame), [35](#)

 semi_join.RPolarsDataFrame, [35](#)
 semi_join.RPolarsLazyFrame
 (semi_join.RPolarsDataFrame),
 [35](#)
 separate.RPolarsDataFrame, [37](#)
 separate.RPolarsLazyFrame
 (separate.RPolarsDataFrame), [37](#)
 sink_csv, [38](#)
 sink_parquet, [40](#)
 slice_head.RPolarsDataFrame
 (slice_tail.RPolarsDataFrame),
 [42](#)
 slice_head.RPolarsLazyFrame
 (slice_tail.RPolarsDataFrame),
 [42](#)
 slice_sample.RPolarsDataFrame
 (slice_tail.RPolarsDataFrame),
 [42](#)
 slice_tail.RPolarsDataFrame, [42](#)
 slice_tail.RPolarsLazyFrame
 (slice_tail.RPolarsDataFrame),
 [42](#)
 summarise.RPolarsDataFrame
 (summarize.RPolarsDataFrame),
 [43](#)
 summarise.RPolarsLazyFrame
 (summarize.RPolarsDataFrame),
 [43](#)
 summarize.RPolarsDataFrame, [43](#)
 summarize.RPolarsLazyFrame
 (summarize.RPolarsDataFrame),
 [43](#)

 tibble, [3](#)
 tidypolars-options, [44](#)

 uncount.RPolarsDataFrame, [45](#)
 uncount.RPolarsLazyFrame
 (uncount.RPolarsDataFrame), [45](#)
 ungroup.RPolarsDataFrame
 (group_by.RPolarsDataFrame), [17](#)
 ungroup.RPolarsLazyFrame
 (group_by.RPolarsDataFrame), [17](#)
 unite.RPolarsDataFrame, [46](#)
 unite.RPolarsLazyFrame
 (unite.RPolarsDataFrame), [46](#)

 vctrs::vec_as_names(), [4](#)