

Package: arrow (via r-universe)

October 12, 2026

Title Integration to 'Apache' 'Arrow'

Version 25.0.1.9000

Description 'Apache' 'Arrow' <<https://arrow.apache.org/>> is a cross-language development platform for in-memory data. It specifies a standardized language-independent columnar memory format for flat and hierarchical data, organized for efficient analytic operations on modern hardware. This package provides an interface to the 'Arrow C++' library.

Depends R (>= 4.2)

License Apache License (>= 2.0)

URL <https://github.com/apache/arrow/>, <https://arrow.apache.org/docs/r/>

BugReports <https://github.com/apache/arrow/issues>

Encoding UTF-8

Language en-US

SystemRequirements C++20; for AWS S3 support on Linux, libcurl and openssl, and libxml2 for Azure (optional); cmake >= 3.26 (build-time only, and only for full source build)

Biarch true

Imports assertthat, bit64 (>= 0.9-7), glue, methods, purrr, R6, rlang (>= 1.0.0), stats, tidyselect (>= 1.0.0), utils, vctrs

Roxygen list(markdown = TRUE, r6 = FALSE, load = ``source"')

Config/testthat/edition 3

Config/build/bootstrap TRUE

Suggests blob, curl, cli, DBI, dbplyr, decor, distro, dplyr, duckdb (>= 0.2.8), hms, jsonlite, knitr, lubridate, pillar, pkgload, reticulate, rmarkdown, stringi, stringr, sys, testthat (>= 3.3.0), tibble, tzdb, withr

LinkingTo cpp11 (>= 0.4.2)

Collate 'arrowExports.R' 'enums.R' 'arrow-object.R' 'type.R'
 'array-data.R' 'arrow-datum.R' 'array.R' 'arrow-info.R'
 'arrow-package.R' 'arrow-tabular.R' 'buffer.R'
 'chunked-array.R' 'io.R' 'compression.R' 'scalar.R' 'compute.R'
 'config.R' 'csv.R' 'dataset.R' 'dataset-factory.R'
 'dataset-format.R' 'dataset-partition.R' 'dataset-scan.R'
 'dataset-write.R' 'dictionary.R' 'dplyr-across.R'
 'dplyr-arrange.R' 'dplyr-by.R' 'dplyr-collect.R'
 'dplyr-count.R' 'dplyr-datetime-helpers.R' 'dplyr-distinct.R'
 'dplyr-eval.R' 'dplyr-filter.R' 'dplyr-funcs-agg.R'
 'dplyr-funcs-augmented.R' 'dplyr-funcs-conditional.R'
 'dplyr-funcs-datetime.R' 'dplyr-funcs-doc.R'
 'dplyr-funcs-math.R' 'dplyr-funcs-simple.R'
 'dplyr-funcs-string.R' 'dplyr-funcs-type.R' 'expression.R'
 'dplyr-funcs.R' 'dplyr-glimpse.R' 'dplyr-group-by.R'
 'dplyr-join.R' 'dplyr-mutate.R' 'dplyr-select.R'
 'dplyr-slice.R' 'dplyr-summarize.R' 'dplyr-union.R'
 'record-batch.R' 'table.R' 'dplyr.R' 'duckdb.R' 'extension.R'
 'feather.R' 'field.R' 'filesystem.R' 'flight.R'
 'install-arrow.R' 'ipc-stream.R' 'json.R' 'memory-pool.R'
 'message.R' 'metadata.R' 'parquet.R' 'python.R'
 'query-engine.R' 'record-batch-reader.R'
 'record-batch-writer.R' 'reexports-tidysselect.R' 'schema.R'
 'udf.R' 'util.R'

Config/roxygen2/version 8.1.0

Config/pak/sysreqs cmake libxml2-dev libssl-dev

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

Date/Publication 2026-10-12 10:01:13 UTC

RemoteUrl <https://github.com/apache/arrow>

RemoteRef HEAD

RemoteSha 9b34c4a6a8856212b885f6a1e934795af04f3192

RemoteSubdir r

Contents

acero	5
Array	14
ArrayData	15
arrow_array	16
arrow_info	17
arrow_table	17
as_arrow_array	18
as_arrow_table	19
as_chunked_array	20
as_data_type	21

as_record_batch	22
as_record_batch_reader	23
as_schema	24
az_container	25
buffer	25
Buffer	26
call_function	27
chunked_array	28
ChunkedArray	28
Codec	30
codec_is_available	30
compression	31
concat_arrays	31
concat_tables	32
copy_files	32
cpu_count	33
create_package_with_all_dependencies	34
csv_convert_options	35
csv_parse_options	36
csv_read_options	37
csv_write_options	39
CsvFileFormat	39
CsvReadOptions	40
CsvTableReader	42
data-type	43
Dataset	47
dataset_factory	48
DataType	50
dictionary	51
DictionaryType	51
Expression	52
ExtensionArray	52
ExtensionType	53
FeatherReader	53
field	54
Field	55
FileFormat	55
FileInfo	56
FileSelector	57
FileSystem	57
FileWriteOptions	60
flight_connect	61
flight_disconnect	61
flight_get	62
flight_put	62
FragmentScanOptions	63
gs_bucket	63
hive_partition	64

infer_schema	65
infer_type	65
InputStream	66
install_arrow	66
install_pyarrow	68
io_thread_count	68
JsonFileFormat	69
list_compute_functions	69
list_flights	70
load_flight_server	71
map_batches	71
match_arrow	72
Message	73
MessageReader	74
mmap_create	74
mmap_open	75
new_extension_type	75
open_dataset	78
open_delim_dataset	82
OutputStream	87
ParquetArrowReaderProperties	87
ParquetFileReader	88
ParquetFileWriter	89
ParquetReaderProperties	90
ParquetWriterProperties	90
Partitioning	91
read_delim_arrow	92
read_feather	97
read_ipc_file	98
read_ipc_stream	99
read_json_arrow	99
read_message	101
read_parquet	101
read_schema	102
record_batch	103
RecordBatch	103
RecordBatchReader	104
RecordBatchWriter	106
register_scalar_function	107
s3_bucket	109
scalar	110
Scalar	110
Scanner	111
schema	112
Schema	113
show_exec_plan	114
Table	115
to_arrow	116

to_duckdb	117
unify_schemas	118
value_counts	119
vctrs_extension_array	119
write_csv_arrow	120
write_dataset	121
write_delim_dataset	124
write_feather	127
write_ipc_file	128
write_ipc_stream	129
write_parquet	130
write_to_raw	132

Index	133
--------------	------------

acero

Functions available in Arrow dplyr queries

Description

The arrow package contains methods for 38 dplyr table functions, many of which are "verbs" that do transformations to one or more tables. The package also has mappings of 229 R functions to the corresponding functions in the Arrow compute library. These allow you to write code inside of dplyr methods that call R functions, including many in packages like stringr and lubridate, and they will get translated to Arrow and run on the Arrow query engine (Acero). This document lists all of the mapped functions.

dplyr verbs

Most verb functions return an arrow_dplyr_query object, similar in spirit to a dbplyr::tbl_lazy. This means that the verbs do not eagerly evaluate the query on the data. To run the query, call either compute(), which returns an arrow Table, or collect(), which pulls the resulting Table into an R tibble.

- `anti_join()`: the copy argument is ignored
- `arrange()`
- `collapse()`
- `collect()`
- `compute()`
- `count()`
- `distinct()`: `.keep_all = TRUE` returns a non-missing value if present, only returning missing values if all are missing.
- `explain()`
- `filter()`
- `filter_out()`

- `full_join()`: the copy argument is ignored
- `glimpse()`
- `group_by()`
- `group_by_drop_default()`
- `group_vars()`
- `groups()`
- `inner_join()`: the copy argument is ignored
- `left_join()`: the copy argument is ignored
- `mutate()`
- `pull()`: the name argument is not supported; returns an R vector by default. Provide `as_vector = FALSE` to return an Arrow `ChunkedArray` instead, or set `options(arrow.pull_as_vector = FALSE)` globally.
- `relocate()`
- `rename()`
- `rename_with()`
- `right_join()`: the copy argument is ignored
- `select()`
- `semi_join()`: the copy argument is ignored
- `show_query()`
- `slice_head()`: slicing within groups not supported; Arrow datasets do not have row order, so head is non-deterministic; prop only supported on queries where `nrow()` is knowable without evaluating
- `slice_max()`: slicing within groups not supported; `with_ties = TRUE` (dplyr default) is not supported; prop only supported on queries where `nrow()` is knowable without evaluating
- `slice_min()`: slicing within groups not supported; `with_ties = TRUE` (dplyr default) is not supported; prop only supported on queries where `nrow()` is knowable without evaluating
- `slice_sample()`: slicing within groups not supported; `replace = TRUE` and the `weight_by` argument not supported; `n` only supported on queries where `nrow()` is knowable without evaluating
- `slice_tail()`: slicing within groups not supported; Arrow datasets do not have row order, so tail is non-deterministic; prop only supported on queries where `nrow()` is knowable without evaluating
- `summarise()`: window functions not currently supported; arguments `.drop = FALSE` and `.groups = "rowwise"` not supported
- `tally()`
- `transmute()`
- `ungroup()`
- `union()`
- `union_all()`

Function mappings

In the list below, any differences in behavior or support between Acero and the R function are listed. If no notes follow the function name, then you can assume that the function works in Acero just as it does in R.

Functions can be called either as `pkg::fun()` or just `fun()`, i.e. both `str_sub()` and `stringr::str_sub()` work.

In addition to these functions, you can call any of Arrow's 283 compute functions directly. Arrow has many functions that don't map to an existing R function. In other cases where there is an R function mapping, you can still call the Arrow function directly if you don't want the adaptations that the R mapping has that make Acero behave like R. These functions are listed in the [C++ documentation](#), and in the function registry in R, they are named with an `arrow_` prefix, such as `arrow_ascii_is_decimal`.

arrow:

- `add_filename()`
- `cast()`

base:

- `!`
- `!=`
- `%%`
- `%%/`
- `%in%`
- `&`
- `*`
- `+`
- `-`
- `/`
- `<`
- `<=`
- `==`
- `>`
- `>=`
- `ISOdate()`
- `ISOdatetime()`
- `^`
- `abs()`
- `acos()`
- `acosh()`
- `all()`
- `any()`
- `as.Date()`: Multiple `tryFormats` not supported in Arrow. Consider using the `lubridate` specialised parsing functions `ymd()`, `ymd()`, etc.

- `as.character()`
- `as.difftime()`: only supports `units = "secs"` (the default)
- `as.double()`
- `as.integer()`
- `as.logical()`
- `as.numeric()`
- `asin()`
- `asinh()`
- `atan()`
- `atanh()`
- `ceiling()`
- `cos()`
- `cosh()`
- `data.frame()`: `row.names` and `check.rows` arguments not supported; `stringsAsFactors` must be `FALSE`
- `difftime()`: only supports `units = "secs"` (the default); `tz` argument not supported
- `endsWith()`
- `exp()`
- `expm1()`
- `floor()`
- `format()`
- `grepl()`
- `gsub()`
- `ifelse()`
- `is.character()`
- `is.double()`
- `is.factor()`
- `is.finite()`
- `is.infinite()`
- `is.integer()`
- `is.list()`
- `is.logical()`
- `is.na()`
- `is.nan()`
- `is.numeric()`
- `log()`
- `log10()`
- `log1p()`
- `log2()`
- `logb()`
- `max()`
- `mean()`

- `min()`
- `nchar()`: `allowNA = TRUE` and `keepNA = TRUE` not supported
- `paste()`: the `collapse` argument is not yet supported
- `paste0()`: the `collapse` argument is not yet supported
- `pmax()`
- `pmin()`
- `prod()`
- `round()`
- `sign()`
- `sin()`
- `sinh()`
- `sqrt()`
- `startsWith()`
- `strftime()`
- `strptime()`: accepts a `unit` argument not present in the base function. Valid values are "s", "ms" (default), "us", "ns".
- `strrep()`
- `strsplit()`
- `sub()`
- `substr()`: `start` and `stop` must be length 1
- `substring()`
- `sum()`
- `tan()`
- `tanh()`
- `tolower()`
- `toupper()`
- `trunc()`
- |

bit64:

- `as.integer64()`
- `is.integer64()`

dplyr:

- `across()`
- `between()`
- `case_when()`: `.ptype` and `.size` arguments not supported
- `coalesce()`
- `desc()`
- `if_all()`
- `if_any()`
- `if_else()`
- `n()`

- `n_distinct()`
- `recode_values()`: ptype argument and unmatched = "error" not supported
- `replace_values()`
- `replace_when()`
- `when_all()`
- `when_any()`

hms:

- `as_hms()`: subsecond precision not supported for character input
- `hms()`: nanosecond times not supported

lubridate:

- `am()`
- `as_date()`
- `as_datetime()`
- `ceiling_date()`
- `date()`
- `date_decimal()`
- `day()`
- `ddays()`
- `decimal_date()`
- `dhours()`
- `dmicroseconds()`
- `dmilliseconds()`
- `dminutes()`
- `dmonths()`
- `dmy()`: locale argument not supported; see docs for `parse_date_time()` mapping for supported formats
- `dmy_h()`: locale argument not supported; see docs for `parse_date_time()` mapping for supported formats
- `dmy_hm()`: locale argument not supported; see docs for `parse_date_time()` mapping for supported formats
- `dmy_hms()`: locale argument not supported; see docs for `parse_date_time()` mapping for supported formats
- `dnanoseconds()`
- `dpicoseconds()`: not supported
- `dseconds()`
- `dst()`
- `dweeks()`
- `dyears()`
- `dym()`: locale argument not supported; see docs for `parse_date_time()` mapping for supported formats
- `epiweek()`

- `epiyear()`
- `fast_strptime()`: non-default values of `lt` and `cutoff_2000` not supported
- `floor_date()`
- `force_tz()`: Timezone conversion from non-UTC timezone not supported; `roll_dst` values of 'error' and 'boundary' are supported for nonexistent times, `roll_dst` values of 'error', 'pre', and 'post' are supported for ambiguous times.
- `format_ISO8601()`
- `hour()`
- `is.Date()`
- `is.POSIXct()`
- `is.instant()`
- `is.timepoint()`
- `isoweek()`
- `isoyear()`
- `leap_year()`
- `make_date()`
- `make_datetime()`: only supports UTC (default) timezone
- `make_difftime()`: only supports units = "secs" (the default); providing both num and ... is not supported
- `mday()`
- `mdy()`: locale argument not supported; see docs for `parse_date_time()` mapping for supported formats
- `mdy_h()`: locale argument not supported; see docs for `parse_date_time()` mapping for supported formats
- `mdy_hm()`: locale argument not supported; see docs for `parse_date_time()` mapping for supported formats
- `mdy_hms()`: locale argument not supported; see docs for `parse_date_time()` mapping for supported formats
- `minute()`
- `month()`
- `my()`: locale argument not supported; see docs for `parse_date_time()` mapping for supported formats
- `myd()`: locale argument not supported; see docs for `parse_date_time()` mapping for supported formats
- `parse_date_time()`: `quiet = FALSE` is not supported Available formats are H, I, j, M, S, U, w, W, y, Y, R, T. On Linux and macOS additionally a, A, b, B, Om, p, r are available.
- `pm()`
- `qday()`
- `quarter()`
- `round_date()`
- `second()`
- `semester()`
- `tz()`
- `wday()`

- `week()`
- `with_tz()`
- `yday()`
- `ydm()`: locale argument not supported; see docs for `parse_date_time()` mapping for supported formats
- `ydm_h()`: locale argument not supported; see docs for `parse_date_time()` mapping for supported formats
- `ydm_hm()`: locale argument not supported; see docs for `parse_date_time()` mapping for supported formats
- `ydm_hms()`: locale argument not supported; see docs for `parse_date_time()` mapping for supported formats
- `year()`
- `ym()`: locale argument not supported; see docs for `parse_date_time()` mapping for supported formats
- `ymd()`: locale argument not supported; see docs for `parse_date_time()` mapping for supported formats
- `ymd_h()`: locale argument not supported; see docs for `parse_date_time()` mapping for supported formats
- `ymd_hm()`: locale argument not supported; see docs for `parse_date_time()` mapping for supported formats
- `ymd_hms()`: locale argument not supported; see docs for `parse_date_time()` mapping for supported formats
- `yq()`: locale argument not supported; see docs for `parse_date_time()` mapping for supported formats

methods:

- `is()`

rlang:

- `is_character()`
- `is_double()`
- `is_integer()`
- `is_list()`
- `is_logical()`

stats:

- `median()`: approximate median (t-digest) is computed
- `quantile()`: probs must be length 1; approximate quantile (t-digest) is computed
- `sd()`
- `var()`

stringi:

- `stri_reverse()`

stringr:

Pattern modifiers `coll()` and `boundary()` are not supported in any functions.

- `str_c()`: the collapse argument is not yet supported
- `str_count()`: pattern must be a length 1 character vector
- `str_detect()`
- `str_dup()`
- `str_ends()`
- `str_iliike()`
- `str_length()`
- `str_like()`
- `str_pad()`
- `str_remove()`
- `str_remove_all()`
- `str_replace()`
- `str_replace_all()`
- `str_replace_na()`
- `str_split()`: Case-insensitive string splitting and splitting into 0 parts not supported
- `str_starts()`
- `str_sub()`: start and end must be length 1
- `str_to_lower()`
- `str_to_title()`
- `str_to_upper()`
- `str_trim()`

tibble:

- `tibble()`

tidyselect:

- `all_of()`
- `contains()`
- `ends_with()`
- `everything()`
- `last_col()`
- `matches()`
- `num_range()`
- `one_of()`
- `starts_with()`

Array

*Array Classes***Description**

An Array is an immutable data array with some logical type and some length. Most logical types are contained in the base Array class; there are also subclasses for DictionaryArray, ListArray, and StructArray.

Factory

The `Array$create()` factory method instantiates an Array and takes the following arguments:

- `x`: an R vector, list, or `data.frame`
- `type`: an optional [data type](#) for `x`. If omitted, the type will be inferred from the data.

`Array$create()` will return the appropriate subclass of Array, such as `DictionaryArray` when given an R factor.

To compose a `DictionaryArray` directly, call `DictionaryArray$create()`, which takes two arguments:

- `x`: an R vector or Array of integers for the dictionary indices
- `dict`: an R vector or Array of dictionary values (like R factor levels but not limited to strings only)

Usage

```
a <- Array$create(x)
length(a)

print(a)
a == a
```

Methods

- `$IsNull(i)`: Return true if value at index is null. Does not boundscheck
- `$IsValid(i)`: Return true if value at index is valid. Does not boundscheck
- `$length()`: Size in the number of elements this array contains
- `$nbytes()`: Total number of bytes consumed by the elements of the array
- `$offset`: A relative position into another array's data, to enable zero-copy slicing
- `$null_count`: The number of null entries in the array
- `$type`: logical type of data
- `$type_id()`: type id
- `$Equals(other)` : is this array equal to other
- `$ApproxEquals(other)` :
- `$Diff(other)` : return a string expressing the difference between two arrays

- `$data()`: return the underlying [ArrayData](#)
- `$as_vector()`: convert to an R vector
- `$ToString()`: string representation of the array
- `$Slice(offset, length = NULL)`: Construct a zero-copy slice of the array with the indicated offset and length. If length is NULL, the slice goes until the end of the array.
- `$Take(i)`: return an Array with values at positions given by integers (R vector or Array Array) i.
- `$Filter(i, keep_na = TRUE)`: return an Array with values at positions where logical vector (or Arrow boolean Array) i is TRUE.
- `$SortIndices(descending = FALSE)`: return an Array of integer positions that can be used to rearrange the Array in ascending or descending order
- `$RangeEquals(other, start_idx, end_idx, other_start_idx)`:
- `$cast(target_type, safe = TRUE, options = cast_options(safe))`: Alter the data in the array to change its type.
- `$View(type)`: Construct a zero-copy view of this array with the given type.
- `$Validate()` : Perform any validation checks to determine obvious inconsistencies within the array's internal data. This can be an expensive check, potentially $O(\text{length})$

Examples

```
my_array <- Array$create(1:10)
my_array$type
my_array$cast(int8())

# Check if value is null; zero-indexed
na_array <- Array$create(c(1:5, NA))
na_array$IsNull(0)
na_array$IsNull(5)
na_array$IsValid(5)
na_array$null_count

# zero-copy slicing; the offset of the new Array will be the same as the index passed to $Slice
new_array <- na_array$Slice(5)
new_array$offset

# Compare 2 arrays
na_array2 <- na_array
na_array2 == na_array # element-wise comparison
na_array2$Equals(na_array) # overall comparison
```

ArrayData

ArrayData class

Description

The `ArrayData` class allows you to get and inspect the data inside an `arrow::Array`.

Usage

```
data <- Array$create(x)$data()

data$type
data$length
data$null_count
data$offset
data$buffers
```

Methods

```
...
```

arrow_array	<i>Create an Arrow Array</i>
-------------	------------------------------

Description

Create an Arrow Array

Usage

```
arrow_array(x, type = NULL)
```

Arguments

x	An R object representable as an Arrow array, e.g. a vector, list, or <code>data.frame</code> .
type	An optional data type for x. If omitted, the type will be inferred from the data.

Examples

```
my_array <- arrow_array(1:10)

# Compare 2 arrays
na_array <- arrow_array(c(1:5, NA))
na_array2 <- na_array
na_array2 == na_array # element-wise comparison
```

arrow_info	<i>Report information on the package's capabilities</i>
------------	---

Description

This function summarizes a number of build-time configurations and run-time settings for the Arrow package. It may be useful for diagnostics.

Usage

```
arrow_info()

arrow_with_acero()

arrow_with_dataset()

arrow_with_substrait()

arrow_with_parquet()

arrow_with_s3()

arrow_with_gcs()

arrow_with_azure()

arrow_with_json()
```

Value

arrow_info() returns a list including version information, boolean "capabilities", and statistics from Arrow's memory allocator, and also Arrow's run-time information. The _available() functions return a logical value whether or not the C++ library was built with support for them.

See Also

If any capabilities are FALSE, see the [install guide](#) for guidance on reinstalling the package.

arrow_table	<i>Create an Arrow Table</i>
-------------	------------------------------

Description

Create an Arrow Table

Usage

```
arrow_table(..., schema = NULL)
```

Arguments

... A data.frame or a named set of Arrays or vectors. If given a mixture of data.frames and named vectors, the inputs will be autospliced together (see examples). Alternatively, you can provide a single Arrow IPC InputStream, Message, Buffer, or R raw object containing a Buffer.

schema a [Schema](#), or NULL (the default) to infer the schema from the data in ... When providing an Arrow IPC buffer, schema is required.

See Also

[Table](#)

Examples

```
tbl <- arrow_table(name = rownames(mtcars), mtcars)
dim(tbl)
dim(head(tbl))
names(tbl)
tbl$mpg
tbl[["cyl"]]
as.data.frame(tbl[4:8, c("gear", "hp", "wt")])
```

as_arrow_array

Convert an object to an Arrow Array

Description

The `as_arrow_array()` function is identical to `Array$create()` except that it is an S3 generic, which allows methods to be defined in other packages to convert objects to [Array](#). `Array$create()` is slightly faster because it tries to convert in C++ before falling back on `as_arrow_array()`.

Usage

```
as_arrow_array(x, ..., type = NULL)

## S3 method for class 'Array'
as_arrow_array(x, ..., type = NULL)

## S3 method for class 'Scalar'
as_arrow_array(x, ..., type = NULL)

## S3 method for class 'ChunkedArray'
as_arrow_array(x, ..., type = NULL)
```

Arguments

`x` An object to convert to an Arrow Array

`...` Passed to S3 methods

`type` A [type](#) for the final Array. A value of NULL will default to the type guessed by [infer_type\(\)](#).

Value

An [Array](#) with type `type`.

Examples

```
as_arrow_array(1:5)
```

as_arrow_table	<i>Convert an object to an Arrow Table</i>
----------------	--

Description

Whereas [arrow_table\(\)](#) constructs a table from one or more columns, `as_arrow_table()` converts a single object to an Arrow [Table](#).

Usage

```
as_arrow_table(x, ..., schema = NULL)

## Default S3 method:
as_arrow_table(x, ...)

## S3 method for class 'Table'
as_arrow_table(x, ..., schema = NULL)

## S3 method for class 'RecordBatch'
as_arrow_table(x, ..., schema = NULL)

## S3 method for class 'data.frame'
as_arrow_table(x, ..., schema = NULL)

## S3 method for class 'RecordBatchReader'
as_arrow_table(x, ...)

## S3 method for class 'Dataset'
as_arrow_table(x, ...)

## S3 method for class 'arrow_dplyr_query'
```

```
as_arrow_table(x, ...)

## S3 method for class 'Schema'
as_arrow_table(x, ...)
```

Arguments

x	An object to convert to an Arrow Table
...	Passed to S3 methods
schema	a Schema , or NULL (the default) to infer the schema from the data in ... When providing an Arrow IPC buffer, schema is required.

Value

A [Table](#)

Examples

```
# use as_arrow_table() for a single object
as_arrow_table(data.frame(col1 = 1, col2 = "two"))

# use arrow_table() to create from columns
arrow_table(col1 = 1, col2 = "two")
```

as_chunked_array	<i>Convert an object to an Arrow ChunkedArray</i>
------------------	---

Description

Whereas [chunked_array\(\)](#) constructs a [ChunkedArray](#) from zero or more [Arrays](#) or R vectors, [as_chunked_array\(\)](#) converts a single object to a [ChunkedArray](#).

Usage

```
as_chunked_array(x, ..., type = NULL)

## S3 method for class 'ChunkedArray'
as_chunked_array(x, ..., type = NULL)

## S3 method for class 'Array'
as_chunked_array(x, ..., type = NULL)
```

Arguments

x	An object to convert to an Arrow Chunked Array
...	Passed to S3 methods
type	A type for the final Array. A value of NULL will default to the type guessed by infer_type() .

Value

A [ChunkedArray](#).

Examples

```
as_chunked_array(1:5)
```

as_data_type

Convert an object to an Arrow [DataType](#)

Description

Convert an object to an Arrow [DataType](#)

Usage

```
as_data_type(x, ...)
```

```
## S3 method for class 'DataType'  
as_data_type(x, ...)
```

```
## S3 method for class 'Field'  
as_data_type(x, ...)
```

```
## S3 method for class 'Schema'  
as_data_type(x, ...)
```

Arguments

x An object to convert to an Arrow [DataType](#)

... Passed to S3 methods.

Value

A [DataType](#) object.

Examples

```
as_data_type(int32())
```

as_record_batch	<i>Convert an object to an Arrow RecordBatch</i>
-----------------	--

Description

Whereas `record_batch()` constructs a [RecordBatch](#) from one or more columns, `as_record_batch()` converts a single object to an Arrow [RecordBatch](#).

Usage

```
as_record_batch(x, ..., schema = NULL)

## S3 method for class 'RecordBatch'
as_record_batch(x, ..., schema = NULL)

## S3 method for class 'Table'
as_record_batch(x, ..., schema = NULL)

## S3 method for class 'arrow_dplyr_query'
as_record_batch(x, ...)

## S3 method for class 'data.frame'
as_record_batch(x, ..., schema = NULL)
```

Arguments

x	An object to convert to an Arrow RecordBatch
...	Passed to S3 methods
schema	a Schema , or NULL (the default) to infer the schema from the data in When providing an Arrow IPC buffer, schema is required.

Value

A [RecordBatch](#)

Examples

```
# use as_record_batch() for a single object
as_record_batch(data.frame(col1 = 1, col2 = "two"))

# use record_batch() to create from columns
record_batch(col1 = 1, col2 = "two")
```

`as_record_batch_reader`*Convert an object to an Arrow RecordBatchReader*

Description

Convert an object to an Arrow RecordBatchReader

Usage

```
as_record_batch_reader(x, ...)  
  
## S3 method for class 'RecordBatchReader'  
as_record_batch_reader(x, ...)  
  
## S3 method for class 'Table'  
as_record_batch_reader(x, ...)  
  
## S3 method for class 'RecordBatch'  
as_record_batch_reader(x, ...)  
  
## S3 method for class 'data.frame'  
as_record_batch_reader(x, ...)  
  
## S3 method for class 'Dataset'  
as_record_batch_reader(x, ...)  
  
## S3 method for class ``function``  
as_record_batch_reader(x, ..., schema)  
  
## S3 method for class 'arrow_dplyr_query'  
as_record_batch_reader(x, ...)  
  
## S3 method for class 'Scanner'  
as_record_batch_reader(x, ...)
```

Arguments

<code>x</code>	An object to convert to a RecordBatchReader
<code>...</code>	Passed to S3 methods
<code>schema</code>	The schema() that must match the schema returned by each call to <code>x</code> when <code>x</code> is a function.

Value

A [RecordBatchReader](#)

Examples

```
reader <- as_record_batch_reader(data.frame(col1 = 1, col2 = "two"))
reader$read_next_batch()
```

as_schema

Convert an object to an Arrow Schema

Description

Convert an object to an Arrow Schema

Usage

```
as_schema(x, ...)
```

S3 method for class 'Schema'

```
as_schema(x, ...)
```

S3 method for class 'StructType'

```
as_schema(x, ...)
```

Arguments

x	An object to convert to a schema()
...	Passed to S3 methods.

Value

A [Schema](#) object.

Examples

```
as_schema(schema(col1 = int32()))
```

az_container	<i>Connect to an Azure Blob Storage container</i>
--------------	---

Description

az_container is a convenience function to create an AzureFileSystem object that provides a file system interface for blob storage containers in an Azure Storage Account.

Usage

```
az_container(container_path, ...)
```

Arguments

container_path string Container name or path.
... Additional connection options, passed to AzureFileSystem\$create().

Value

A SubTreeFileSystem containing an AzureFileSystem and the container's relative path. Note that this function's success does not guarantee that you are authorized to access the container's contents.

Examples

```
container_fs <- az_container(  
  container_path = "arrow-datasets",  
  account_name = azurite_account_name,  
  account_key = azurite_account_key,  
  blob_storage_authority = azurite_blob_storage_authority,  
  blob_storage_scheme = azurite_blob_storage_scheme  
)
```

buffer	<i>Create a Buffer</i>
--------	------------------------

Description

Create a Buffer

Usage

```
buffer(x)
```

Arguments

`x` R object. Only raw, numeric and integer vectors are currently supported

Value

an instance of `Buffer` that borrows memory from `x`

See Also

[Buffer](#)

Buffer

Buffer class

Description

A `Buffer` is an object containing a pointer to a piece of contiguous memory with a particular size.

Factory

`buffer()` lets you create an `arrow::Buffer` from an R object

Methods

- `$is_mutable` : is this buffer mutable?
- `$ZeroPadding()` : zero bytes in padding, i.e. bytes between size and capacity
- `$size` : size in memory, in bytes
- `$capacity`: possible capacity, in bytes

Examples

```
my_buffer <- buffer(c(1, 2, 3, 4))
my_buffer$is_mutable
my_buffer$ZeroPadding()
my_buffer$size
my_buffer$capacity
```

call_function	<i>Call an Arrow compute function</i>
---------------	---------------------------------------

Description

This function provides a lower-level API for calling Arrow functions by their string function name. You won't use it directly for most applications. Many Arrow compute functions are mapped to R methods, and in a dplyr evaluation context, [all Arrow functions](#) are callable with an arrow_ prefix.

Usage

```
call_function(
  function_name,
  ...,
  args = list(...),
  options = empty_named_list()
)
```

Arguments

function_name	string Arrow compute function name
...	Function arguments, which may include Array, ChunkedArray, Scalar, RecordBatch, or Table.
args	list arguments as an alternative to specifying in ...
options	named list of C++ function options.

Details

When passing indices in ..., args, or options, express them as 0-based integers (consistent with C++).

Value

An Array, ChunkedArray, Scalar, RecordBatch, or Table, whatever the compute function results in.

See Also

[Arrow C++ documentation](#) for the functions and their respective options.

Examples

```
a <- Array$create(c(1L, 2L, 3L, NA, 5L))
s <- Scalar$create(4L)
call_function("coalesce", a, s)

a <- Array$create(rnorm(10000))
call_function("quantile", a, options = list(q = seq(0, 1, 0.25)))
```

chunked_array	Create a Chunked Array
---------------	------------------------

Description

Create a Chunked Array

Usage

```
chunked_array(..., type = NULL)
```

Arguments

- ... R objects to coerce into a ChunkedArray. They must be of the same type.
- type An optional [data type](#). If omitted, the type will be inferred from the data.

See Also

[ChunkedArray](#)

Examples

```
# Pass items into chunked_array as separate objects to create chunks
class_scores <- chunked_array(c(87, 88, 89), c(94, 93, 92), c(71, 72, 73))

# If you pass a list into chunked_array, you get a list of length 1
list_scores <- chunked_array(list(c(9.9, 9.6, 9.5), c(8.2, 8.3, 8.4), c(10.0, 9.9, 9.8)))

# When constructing a ChunkedArray, the first chunk is used to infer type.
infer_type(chunked_array(c(1, 2, 3), c(5L, 6L, 7L)))

# Concatenating chunked arrays returns a new chunked array containing all chunks
a <- chunked_array(c(1, 2), 3)
b <- chunked_array(c(4, 5), 6)
c(a, b)
```

ChunkedArray	ChunkedArray class
--------------	--------------------

Description

A ChunkedArray is a data structure managing a list of primitive Arrow [Arrays](#) logically as one large array. Chunked arrays may be grouped together in a [Table](#).

Factory

The ChunkedArray\$create() factory method instantiates the object from various Arrays or R vectors. chunked_array() is an alias for it.

Methods

- `$length()`: Size in the number of elements this array contains
- `$chunk(i)`: Extract an Array chunk by integer position
- `$nbytes()` : Total number of bytes consumed by the elements of the array
- `$as_vector()`: convert to an R vector
- `$Slice(offset, length = NULL)`: Construct a zero-copy slice of the array with the indicated offset and length. If length is NULL, the slice goes until the end of the array.
- `$Take(i)`: return a ChunkedArray with values at positions given by integers i. If i is an Arrow Array or ChunkedArray, it will be coerced to an R vector before taking.
- `$Filter(i, keep_na = TRUE)`: return a ChunkedArray with values at positions where logical vector or Arrow boolean-type (Chunked)Array i is TRUE.
- `$SortIndices(descending = FALSE)`: return an Array of integer positions that can be used to rearrange the ChunkedArray in ascending or descending order
- `$cast(target_type, safe = TRUE, options = cast_options(safe))`: Alter the data in the array to change its type.
- `$null_count`: The number of null entries in the array
- `$chunks`: return a list of Arrays
- `$num_chunks`: integer number of chunks in the ChunkedArray
- `$type`: logical type of data
- `$View(type)`: Construct a zero-copy view of this ChunkedArray with the given type.
- `$Validate()`: Perform any validation checks to determine obvious inconsistencies within the array's internal data. This can be an expensive check, potentially $O(\text{length})$

See Also

[Array](#)

Examples

```
# Pass items into chunked_array as separate objects to create chunks
class_scores <- chunked_array(c(87, 88, 89), c(94, 93, 92), c(71, 72, 73))
class_scores$num_chunks

# When taking a Slice from a chunked_array, chunks are preserved
class_scores$Slice(2, length = 5)

# You can combine Take and SortIndices to return a ChunkedArray with 1 chunk
# containing all values, ordered.
class_scores$Take(class_scores$SortIndices(descending = TRUE))

# If you pass a list into chunked_array, you get a list of length 1
list_scores <- chunked_array(list(c(9.9, 9.6, 9.5), c(8.2, 8.3, 8.4), c(10.0, 9.9, 9.8)))
list_scores$num_chunks

# When constructing a ChunkedArray, the first chunk is used to infer type.
doubles <- chunked_array(c(1, 2, 3), c(5L, 6L, 7L))
doubles$type
```

```
# Concatenating chunked arrays returns a new chunked array containing all chunks
a <- chunked_array(c(1, 2), 3)
b <- chunked_array(c(4, 5), 6)
c(a, b)
```

Codec

Compression Codec class

Description

Codecs allow you to create [compressed input and output streams](#).

Factory

The `Codec$create()` factory method takes the following arguments:

- `type`: string name of the compression method. Possible values are "uncompressed", "snappy", "gzip", "brotli", "zstd", "lz4", "lzo", or "bz2". `type` may be upper- or lower-cased. Not all methods may be available; support depends on build-time flags for the C++ library. See [codec_is_available\(\)](#). Most builds support at least "snappy" and "gzip". All support "uncompressed".
- `compression_level`: compression level, the default value (NA) uses the default compression level for the selected compression type.

codec_is_available

Check whether a compression codec is available

Description

Support for compression libraries depends on the build-time settings of the Arrow C++ library. This function lets you know which are available for use.

Usage

```
codec_is_available(type)
```

Arguments

`type` A string, one of "uncompressed", "snappy", "gzip", "brotli", "zstd", "lz4", "lzo", or "bz2", case-insensitive.

Value

Logical: is type available?

Examples

```
codec_is_available("gzip")
```

compression	<i>Compressed stream classes</i>
-------------	----------------------------------

Description

CompressedInputStream and CompressedOutputStream allow you to apply a compression [Codec](#) to an input or output stream.

Factory

The CompressedInputStream\$create() and CompressedOutputStream\$create() factory methods instantiate the object and take the following arguments:

- stream An [InputStream](#) or [OutputStream](#), respectively
- codec A Codec, either a [Codec](#) instance or a string
- compression_level compression level for when the codec argument is given as a string

Methods

Methods are inherited from [InputStream](#) and [OutputStream](#), respectively

concat_arrays	<i>Concatenate zero or more Arrays</i>
---------------	--

Description

Concatenates zero or more [Array](#) objects into a single array. This operation will make a copy of its input; if you need the behavior of a single Array but don't need a single object, use [ChunkedArray](#).

Usage

```
concat_arrays(..., type = NULL)

## S3 method for class 'Array'
c(...)
```

Arguments

...	zero or more Array objects to concatenate
type	An optional type describing the desired type for the final Array.

Value

A single [Array](#)

Examples

```
concat_arrays(Array$create(1:3), Array$create(4:5))
```

concat_tables	<i>Concatenate one or more Tables</i>
---------------	---------------------------------------

Description

Concatenate one or more [Table](#) objects into a single table. This operation does not copy array data, but instead creates new chunked arrays for each column that point at existing array data.

Usage

```
concat_tables(..., unify_schemas = TRUE)
```

Arguments

...	One or more Table or RecordBatch objects. RecordBatch objects will be automatically converted to Tables.
unify_schemas	If TRUE, the schemas of the tables will be first unified with fields of the same name being merged, then each table will be promoted to the unified schema before being concatenated. Otherwise, all tables should have the same schema.

Examples

```
tbl <- arrow_table(name = rownames(mtcars), mtcars)
prius <- arrow_table(name = "Prius", mpg = 58, cyl = 4, disp = 1.8)
combined <- concat_tables(tbl, prius)
tail(combined)$to_data_frame()

# Can also pass RecordBatch objects
batch <- record_batch(name = "Volt", mpg = 53, cyl = 4, disp = 1.5)
combined2 <- concat_tables(tbl, batch)
```

copy_files	<i>Copy files between FileSystems</i>
------------	---------------------------------------

Description

Copy files between FileSystems

Usage

```
copy_files(from, to, chunk_size = 1024L * 1024L)
```

Arguments

from	A string path to a local directory or file, a URI, or a SubTreeFileSystem. Files will be copied recursively from this path.
to	A string path to a local directory or file, a URI, or a SubTreeFileSystem. Directories will be created as necessary
chunk_size	The maximum size of block to read before flushing to the destination file. A larger chunk_size will use more memory while copying but may help accommodate high latency FileSystems.

Value

Nothing: called for side effects in the file system

Examples

```
# Copy an S3 bucket's files to a local directory:
copy_files("s3://your-bucket-name", "local-directory")
# Using a FileSystem object
copy_files(s3_bucket("your-bucket-name"), "local-directory")
# Or go the other way, from local to S3
copy_files("local-directory", s3_bucket("your-bucket-name"))
```

cpu_count

Manage the global CPU thread pool in libarrow

Description

Manage the global CPU thread pool in libarrow

Usage

```
cpu_count()

set_cpu_count(num_threads)
```

Arguments

num_threads	integer: New number of threads for thread pool
-------------	--

```
create_package_with_all_dependencies
```

Create a source bundle that includes all thirdparty dependencies

Description

Create a source bundle that includes all thirdparty dependencies

Usage

```
create_package_with_all_dependencies(dest_file = NULL, source_file = NULL)
```

Arguments

<code>dest_file</code>	File path for the new tar.gz package. Defaults to <code>arrow_V.V.V_with_deps.tar.gz</code> in the current directory (V.V.V is the version)
<code>source_file</code>	File path for the input tar.gz package. Defaults to downloading the package from CRAN (or whatever you have set as the first in <code>getOption("repos")</code>)

Value

The full path to `dest_file`, invisibly

This function is used for setting up an offline build. If it's possible to download at build time, don't use this function. Instead, let cmake download the required dependencies for you. These downloaded dependencies are only used in the build if `ARROW_DEPENDENCY_SOURCE` is unset, `BUNDLED`, or `AUTO`. <https://arrow.apache.org/docs/developers/cpp/building.html#offline-builds>

If you're using binary packages you shouldn't need to use this function. You should download the appropriate binary from your package repository, transfer that to the offline computer, and install that. Any OS can create the source bundle, but it cannot be installed on Windows. (Instead, use a standard Windows binary package.)

Note if you're using RStudio Package Manager on Linux: If you still want to make a source bundle with this function, make sure to set the first repo in `options("repos")` to be a mirror that contains source packages (that is: something other than the RSPM binary mirror URLs).

Steps for an offline install with optional dependencies::

Using a computer with internet access, pre-download the dependencies::

- Install the arrow package *or* run `source("https://raw.githubusercontent.com/apache/arrow/main/r/R/install_arrow_pkg.R")`
- Run `create_package_with_all_dependencies("my_arrow_pkg.tar.gz")`
- Copy the newly created `my_arrow_pkg.tar.gz` to the computer without internet access

On the computer without internet access, install the prepared package::

- Install the arrow package from the copied file
 - `install.packages("my_arrow_pkg.tar.gz", dependencies = c("Depends", "Imports", "LinkingTo"))`
 - This installation will build from source, so cmake must be available
- Run `arrow_info()` to check installed capabilities

Examples

```
## Not run:
new_pkg <- create_package_with_all_dependencies()
# Note: this works when run in the same R session, but it's meant to be
# copied to a different computer.
install.packages(new_pkg, dependencies = c("Depends", "Imports", "LinkingTo"))

## End(Not run)
```

csv_convert_options	<i>CSV Convert Options</i>
---------------------	----------------------------

Description

CSV Convert Options

Usage

```
csv_convert_options(
  check_utf8 = TRUE,
  null_values = c("", "NA"),
  true_values = c("T", "true", "TRUE"),
  false_values = c("F", "false", "FALSE"),
  strings_can_be_null = FALSE,
  col_types = NULL,
  auto_dict_encode = FALSE,
  auto_dict_max_cardinality = 50L,
  include_columns = character(),
  include_missing_columns = FALSE,
  timestamp_parsers = NULL,
  decimal_point = "."
)
```

Arguments

check_utf8	Logical: check UTF8 validity of string columns?
null_values	Character vector of recognized spellings for null values. Analogous to the <code>na.strings</code> argument to read.csv() or <code>na</code> in readr::read_csv() .
true_values	Character vector of recognized spellings for TRUE values
false_values	Character vector of recognized spellings for FALSE values
strings_can_be_null	Logical: can string / binary columns have null values? Similar to the <code>quoted_na</code> argument to readr::read_csv()
col_types	A Schema or NULL to infer types

auto_dict_encode	Logical: Whether to try to automatically dictionary-encode string / binary data (think <code>stringsAsFactors</code>). This setting is ignored for non-inferred columns (those in <code>col_types</code>).
auto_dict_max_cardinality	If <code>auto_dict_encode</code> , string/binary columns are dictionary-encoded up to this number of unique values (default 50), after which it switches to regular encoding.
include_columns	If non-empty, indicates the names of columns from the CSV file that should be actually read and converted (in the vector's order).
include_missing_columns	Logical: if <code>include_columns</code> is provided, should columns named in it but not found in the data be included as a column of type <code>null()</code> ? The default (FALSE) means that the reader will instead raise an error.
timestamp_parsers	User-defined timestamp parsers, tried in order when inferring column types and when converting columns of type <code>timestamp()</code> . Possible values are: • NULL: the default, which uses the ISO-8601 parser • a character vector of <code>strptime</code> parse strings • a list of <code>TimestampParser</code> objects and/or parse strings Supplying parsers replaces the default ISO-8601 parser rather than adding to it. If none of the parsers match a value during type inference, the column is read as a string without error; to get an error instead, specify the column as a timestamp in <code>col_types</code> . These parsers are not used for date columns.
decimal_point	Character to use for decimal point in floating point numbers.

Examples

```
tf <- tempfile()
on.exit(unlink(tf))
writeLines("x\n1\nNULL\n2\nNA", tf)
read_csv_arrow(tf, convert_options = csv_convert_options(null_values = c("", "NA", "NULL")))
open_csv_dataset(tf, convert_options = csv_convert_options(null_values = c("", "NA", "NULL")))
```

csv_parse_options	<i>CSV Parsing Options</i>
-------------------	----------------------------

Description

CSV Parsing Options

Usage

```
csv_parse_options(  
  delimiter = ",",  
  quoting = TRUE,  
  quote_char = "\"",  
  double_quote = TRUE,  
  escaping = FALSE,  
  escape_char = "\\",  
  newlines_in_values = FALSE,  
  ignore_empty_lines = TRUE  
)
```

Arguments

delimiter	Field delimiting character
quoting	Logical: are strings quoted?
quote_char	Quoting character, if quoting is TRUE
double_quote	Logical: are quotes inside values double-quoted?
escaping	Logical: whether escaping is used
escape_char	Escaping character, if escaping is TRUE
newlines_in_values	Logical: are values allowed to contain CR (0x0d) and LF (0x0a) characters?
ignore_empty_lines	Logical: should empty lines be ignored (default) or generate a row of missing values (if FALSE)?

Examples

```
tf <- tempfile()  
on.exit(unlink(tf))  
writeLines("x\n1\n\n2", tf)  
read_csv_arrow(tf, parse_options = csv_parse_options(ignore_empty_lines = FALSE))  
open_csv_dataset(tf, parse_options = csv_parse_options(ignore_empty_lines = FALSE))
```

csv_read_options	CSV Reading Options
------------------	---------------------

Description

CSV Reading Options

Usage

```
csv_read_options(
  use_threads = option_use_threads(),
  block_size = 1048576L,
  skip_rows = 0L,
  column_names = character(0),
  autogenerate_column_names = FALSE,
  encoding = "UTF-8",
  skip_rows_after_names = 0L
)
```

Arguments

use_threads	Whether to use the global CPU thread pool
block_size	Block size we request from the IO layer; also determines the size of chunks when use_threads is TRUE.
skip_rows	Number of lines to skip before reading data (default 0).
column_names	Character vector to supply column names. If length-0 (the default), the first non-skipped row will be parsed to generate column names, unless autogenerate_column_names is TRUE.
autogenerate_column_names	Logical: generate column names instead of using the first non-skipped row (the default)? If TRUE, column names will be "f0", "f1", ..., "fN".
encoding	The file encoding. (default "UTF-8")
skip_rows_after_names	Number of lines to skip after the column names (default 0). This number can be larger than the number of rows in one block, and empty rows are counted. The order of application is as follows: • skip_rows is applied (if non-zero); • column names are read (unless column_names is set); • skip_rows_after_names is applied (if non-zero).

Examples

```
tf <- tempfile()
on.exit(unlink(tf))
writeLines("my file has a non-data header\nx\n1\n2", tf)
read_csv_arrow(tf, read_options = csv_read_options(skip_rows = 1))
open_csv_dataset(tf, read_options = csv_read_options(skip_rows = 1))
```

csv_write_options	<i>CSV Writing Options</i>
-------------------	----------------------------

Description

CSV Writing Options

Usage

```
csv_write_options(  
  include_header = TRUE,  
  batch_size = 1024L,  
  null_string = "",  
  delimiter = ",",  
  eol = "\n",  
  quoting_style = c("Needed", "AllValid", "None")  
)
```

Arguments

- include_header Whether to write an initial header line with column names
- batch_size Maximum number of rows processed at a time.
- null_string The string to be written for null values. Must not contain quotation marks.
- delimiter Field delimiter
- eol The end of line character to use for ending rows
- quoting_style How to handle quotes. "Needed" (Only enclose values in quotes which need them, because their CSV rendering can contain quotes itself (e.g. strings or binary values)), "AllValid" (Enclose all valid values in quotes), or "None" (Do not enclose any values in quotes).

Examples

```
tf <- tempfile()  
on.exit(unlink(tf))  
write_csv_arrow(airquality, tf, write_options = csv_write_options(null_string = "-99"))
```

CsvFileFormat	<i>CSV dataset file format</i>
---------------	--------------------------------

Description

A CsvFileFormat is a [FileFormat](#) subclass which holds information about how to read and parse the files included in a CSV Dataset.

Value

A CsvFileFormat object

Factory

CsvFileFormat\$create() can take options in the form of lists passed through as parse_options, read_options, or convert_options parameters. Alternatively, readr-style options can be passed through individually. While it is possible to pass in CSVReadOptions, CSVConvertOptions, and CSVParseOptions objects, this is not recommended as options set in these objects are not validated for compatibility.

See Also

[FileFormat](#)

Examples

```
# Set up directory for examples
tf <- tempfile()
dir.create(tf)
on.exit(unlink(tf))
df <- data.frame(x = c("1", "2", "NULL"))
write.table(df, file.path(tf, "file1.txt"), sep = ",", row.names = FALSE)

# Create CsvFileFormat object with Arrow-style null_values option
format <- CsvFileFormat$create(convert_options = list(null_values = c("", "NA", "NULL")))
open_dataset(tf, format = format)

# Use readr-style options
format <- CsvFileFormat$create(na = c("", "NA", "NULL"))
open_dataset(tf, format = format)
```

CsvReadOptions

File reader options

Description

CsvReadOptions, CsvParseOptions, CsvConvertOptions, JsonReadOptions, JsonParseOptions, and TimestampParser are containers for various file reading options. See their usage in [read_csv_arrow\(\)](#) and [read_json_arrow\(\)](#), respectively.

Factory

The CsvReadOptions\$create() and JsonReadOptions\$create() factory methods take the following arguments:

- use_threads Whether to use the global CPU thread pool

- `block_size` Block size we request from the IO layer; also determines the size of chunks when `use_threads` is TRUE. NB: if FALSE, JSON input must end with an empty line.

`CsvReadOptions$create()` further accepts these additional arguments:

- `skip_rows` Number of lines to skip before reading data (default 0).
- `column_names` Character vector to supply column names. If length-0 (the default), the first non-skipped row will be parsed to generate column names, unless `autogenerate_column_names` is TRUE.
- `autogenerate_column_names` Logical: generate column names instead of using the first non-skipped row (the default)? If TRUE, column names will be "f0", "f1", ..., "fN".
- `encoding` The file encoding. (default "UTF-8")
- `skip_rows_after_names` Number of lines to skip after the column names (default 0). This number can be larger than the number of rows in one block, and empty rows are counted. The order of application is as follows:
 - `skip_rows` is applied (if non-zero);
 - column names are read (unless `column_names` is set);
 - `skip_rows_after_names` is applied (if non-zero).

`CsvParseOptions$create()` takes the following arguments:

- `delimiter` Field delimiting character (default ",")
- `quoting` Logical: are strings quoted? (default TRUE)
- `quote_char` Quoting character, if quoting is TRUE (default "'")
- `double_quote` Logical: are quotes inside values double-quoted? (default TRUE)
- `escaping` Logical: whether escaping is used (default FALSE)
- `escape_char` Escaping character, if escaping is TRUE (default "\\")
- `newlines_in_values` Logical: are values allowed to contain CR (0x0d) and LF (0x0a) characters? (default FALSE)
- `ignore_empty_lines` Logical: should empty lines be ignored (default) or generate a row of missing values (if FALSE)?

`JsonParseOptions$create()` accepts only the `newlines_in_values` argument.

`CsvConvertOptions$create()` takes the following arguments:

- `check_utf8` Logical: check UTF8 validity of string columns? (default TRUE)
- `null_values` character vector of recognized spellings for null values. Analogous to the `na.strings` argument to `read.csv()` or `na` in `readr::read_csv()`.
- `strings_can_be_null` Logical: can string / binary columns have null values? Similar to the `quoted_na` argument to `readr::read_csv()`. (default FALSE)
- `true_values` character vector of recognized spellings for TRUE values
- `false_values` character vector of recognized spellings for FALSE values
- `col_types` A Schema or NULL to infer types

- `auto_dict_encode` Logical: Whether to try to automatically dictionary-encode string / binary data (think `stringsAsFactors`). Default `FALSE`. This setting is ignored for non-inferred columns (those in `col_types`).
- `auto_dict_max_cardinality` If `auto_dict_encode`, string/binary columns are dictionary-encoded up to this number of unique values (default 50), after which it switches to regular encoding.
- `include_columns` If non-empty, indicates the names of columns from the CSV file that should be actually read and converted (in the vector's order).
- `include_missing_columns` Logical: if `include_columns` is provided, should columns named in it but not found in the data be included as a column of type `null()`? The default (`FALSE`) means that the reader will instead raise an error.
- `timestamp_parsers` User-defined timestamp parsers, tried in order when inferring column types and when converting timestamp columns. Possible values are (a) `NULL`, the default, which uses the ISO-8601 parser; (b) a character vector of `strptime` parse strings; or (c) a list of `TimestampParser` objects and/or parse strings. Supplying parsers replaces the default ISO-8601 parser; see `read_delim_arrow()` for details.
- `decimal_point` Character to use for decimal point in floating point numbers. Default: `."`

`TimestampParser$create()` takes an optional format string argument. See `strptime()` for example syntax. The default is to use an ISO-8601 format parser, which is useful as a fallback at the end of a list of `timestamp_parsers`.

The `CsvWriteOptions$create()` factory method takes the following arguments:

- `include_header` Whether to write an initial header line with column names
- `batch_size` Maximum number of rows processed at a time. Default is 1024.
- `null_string` The string to be written for null values. Must not contain quotation marks. Default is an empty string (`""`).
- `eol` The end of line character to use for ending rows.
- `delimiter` Field delimiter
- `quoting_style` Quoting style: "Needed" (Only enclose values in quotes which need them, because their CSV rendering can contain quotes itself (e.g. strings or binary values)), "All-Valid" (Enclose all valid values in quotes), or "None" (Do not enclose any values in quotes).

Active bindings

- `column_names`: from `CsvReadOptions`

CsvTableReader

Arrow CSV and JSON table reader classes

Description

`CsvTableReader` and `JsonTableReader` wrap the Arrow C++ CSV and JSON table readers. See their usage in `read_csv_arrow()` and `read_json_arrow()`, respectively.

Factory

The `CsvTableReader$create()` and `JsonTableReader$create()` factory methods take the following arguments:

- file An Arrow [InputStream](#)
- convert_options (CSV only), parse_options, read_options: see [CsvReadOptions](#)
- ... additional parameters.

Methods

- `$Read()`: returns an Arrow Table.

data-type

Create Arrow data types

Description

These functions create type objects corresponding to Arrow types. Use them when defining a [schema\(\)](#) or as inputs to other types, like `struct`. Most of these functions don't take arguments, but a few do.

Usage

`int8()`

`int16()`

`int32()`

`int64()`

`uint8()`

`uint16()`

`uint32()`

`uint64()`

`float16()`

`halffloat()`

`float32()`

`float()`

```
float64()

boolean()

bool()

utf8()

large_utf8()

binary()

large_binary()

fixed_size_binary(byte_width)

string()

date32()

date64()

time32(unit = c("ms", "s"))

time64(unit = c("ns", "us"))

duration(unit = c("s", "ms", "us", "ns"))

null()

timestamp(unit = c("s", "ms", "us", "ns"), timezone = "")

decimal(precision, scale)

decimal32(precision, scale)

decimal64(precision, scale)

decimal128(precision, scale)

decimal256(precision, scale)

struct(...)

list_of(type)

large_list_of(type)
```

```
fixed_size_list_of(type, list_size)
```

```
map_of(key_type, item_type, .keys_sorted = FALSE)
```

Arguments

byte_width	byte width for FixedSizeBinary type.
unit	For time/timestamp types, the time unit. <code>time32()</code> can take either "s" or "ms", while <code>time64()</code> can be "us" or "ns". <code>timestamp()</code> can take any of those four values.
timezone	For <code>timestamp()</code> , an optional time zone string.
precision	For <code>decimal()</code> , <code>decimal128()</code> , and <code>decimal256()</code> the number of significant digits the arrow decimal type can represent. The maximum precision for <code>decimal128()</code> is 38 significant digits, while for <code>decimal256()</code> it is 76 digits. <code>decimal()</code> will use it to choose which type of decimal to return.
scale	For <code>decimal()</code> , <code>decimal128()</code> , and <code>decimal256()</code> the number of digits after the decimal point. It can be negative.
...	For <code>struct()</code> , a named list of types to define the struct columns
type	For <code>list_of()</code> , a data type to make a list-of-type
list_size	list size for FixedSizeList type.
key_type, item_type	For MapType, the key and item types.
.keys_sorted	Use TRUE to assert that keys of a MapType are sorted.

Details

A few functions have aliases:

- `utf8()` and `string()`
- `float16()` and `halffloat()`
- `float32()` and `float()`
- `bool()` and `boolean()`
- When called inside an arrow function, such as `schema()` or `cast()`, `double()` also is supported as a way of creating a `float64()`

`date32()` creates a datetime type with a "day" unit, like the R Date class. `date64()` has a "ms" unit.

`uint64` (64 bit unsigned integer) is always converted to double ("numeric") in R. Note that doubles cannot exactly represent all `uint64` values; precision may be lost for values above 2^{53} . `uint32` (32 bit unsigned integer) and `int64` (64-bit signed integer) types may contain values that exceed the range of R's integer type (32-bit signed integer). When they do, `uint32` is converted to double ("numeric") and `int64` is converted to `bit64::integer64`. For `int64` types, this conversion can be disabled (so that `int64` always yields a `bit64::integer64` object) by setting `options(arrow.int64_downcast = FALSE)`. For values nested inside a list type, this decision is made once for all values in the column, so every element of the list uses the same R type.

`decimal128()` creates a `Decimal128Type`. Arrow decimals are fixed-point decimal numbers encoded as a scalar integer. The precision is the number of significant digits that the decimal type can represent; the scale is the number of digits after the decimal point. For example, the number 1234.567 has a precision of 7 and a scale of 3. Note that scale can be negative.

As an example, `decimal128(7, 3)` can exactly represent the numbers 1234.567 and -1234.567 (encoded internally as the 128-bit integers 1234567 and -1234567, respectively), but neither 12345.67 nor 123.4567.

`decimal128(5, -3)` can exactly represent the number 12345000 (encoded internally as the 128-bit integer 12345), but neither 123450000 nor 1234500. The scale can be thought of as an argument that controls rounding. When negative, scale causes the number to be expressed using scientific notation and power of 10.

`decimal256()` creates a `Decimal256Type`, which allows for higher maximum precision. For most use cases, the maximum precision offered by `Decimal128Type` is sufficient, and it will result in a more compact and more efficient encoding.

`decimal()` creates either a `Decimal128Type` or a `Decimal256Type` depending on the value for precision. If precision is greater than 38 a `Decimal256Type` is returned, otherwise a `Decimal128Type`.

Use `decimal128()` or `decimal256()` as the names are more informative than `decimal()`.

Value

An Arrow type object inheriting from [DataType](#).

See Also

[dictionary\(\)](#) for creating a dictionary (factor-like) type.

Examples

```
bool()
struct(a = int32(), b = double())
timestamp("ms", timezone = "CEST")
time64("ns")

# Use the cast method to change the type of data contained in Arrow objects.
# Please check the documentation of each data object class for details.
my_scalar <- Scalar$create(0L, type = int64()) # int64
my_scalar$cast(timestamp("ns")) # timestamp[ns]

my_array <- Array$create(0L, type = int64()) # int64
my_array$cast(timestamp("s", timezone = "UTC")) # timestamp[s, tz=UTC]

my_chunked_array <- chunked_array(0L, 1L) # int32
my_chunked_array$cast(date32()) # date32[day]

# You can also use `cast()` in an Arrow dplyr query.
if (requireNamespace("dplyr", quietly = TRUE)) {
  library(dplyr, warn.conflicts = FALSE)
  arrow_table(mtcars) |>
    transmute(
```

```

      col1 = cast(cyl, string()),
      col2 = cast(cyl, int8())
    ) |>
    compute()
  }

```

Dataset

Multi-file datasets

Description

Arrow Datasets allow you to query against data that has been split across multiple files. This sharding of data may indicate partitioning, which can accelerate queries that only touch some partitions (files).

A Dataset contains one or more Fragments, such as files, of potentially differing type and partitioning.

For `Dataset$create()`, see [open_dataset\(\)](#), which is an alias for it.

`DatasetFactory` is used to provide finer control over the creation of Datasets.

Factory

`DatasetFactory` is used to create a Dataset, inspect the [Schema](#) of the fragments contained in it, and declare a partitioning. `FileSystemDatasetFactory` is a subclass of `DatasetFactory` for discovering files in the local file system, the only currently supported file system.

For the `DatasetFactory$create()` factory method, see [dataset_factory\(\)](#), an alias for it. A `DatasetFactory` has:

- `$Inspect(unify_schemas)`: If `unify_schemas` is `TRUE`, all fragments will be scanned and a unified [Schema](#) will be created from them; if `FALSE` (default), only the first fragment will be inspected for its schema. Use this fast path when you know and trust that all fragments have an identical schema.
- `$Finish(schema, unify_schemas)`: Returns a Dataset. If `schema` is provided, it will be used for the Dataset; if omitted, a `Schema` will be created from inspecting the fragments (files) in the dataset, following `unify_schemas` as described above.

`FileSystemDatasetFactory$create()` is a lower-level factory method and takes the following arguments:

- `filesystem`: A [FileSystem](#)
- `selector`: Either a [FileSelector](#) or `NULL`
- `paths`: Either a character vector of file paths or `NULL`
- `format`: A [FileFormat](#)
- `partitioning`: Either `Partitioning`, `PartitioningFactory`, or `NULL`

Methods

A Dataset has the following methods:

- `$NewScan()`: Returns a [ScannerBuilder](#) for building a query
- `$WithSchema()`: Returns a new Dataset with the specified schema. This method currently supports only adding, removing, or reordering fields in the schema: you cannot alter or cast the field types.
- `$schema`: Active binding that returns the [Schema](#) of the Dataset; you may also replace the dataset's schema by using `ds$schema <- new_schema`.

FileSystemDataset has the following methods:

- `$files`: Active binding, returns the files of the FileSystemDataset
- `$format`: Active binding, returns the [FileFormat](#) of the FileSystemDataset

UnionDataset has the following methods:

- `$children`: Active binding, returns all child Datasets.

See Also

[open_dataset\(\)](#) for a simple interface to creating a Dataset

dataset_factory	Create a DatasetFactory
-----------------	-------------------------

Description

A [Dataset](#) can be constructed using one or more [DatasetFactory](#)s. This function helps you construct a DatasetFactory that you can pass to [open_dataset\(\)](#).

Usage

```
dataset_factory(
  x,
  filesystem = NULL,
  format = c("parquet", "arrow", "ipc", "feather", "csv", "tsv", "text", "json"),
  partitioning = NULL,
  hive_style = NA,
  factory_options = list(),
  ...
)
```

Arguments

x	A string path to a directory containing data files, a vector of one or more string paths to data files, or a list of DatasetFactory objects whose datasets should be combined. If this argument is specified it will be used to construct a UnionDatasetFactory and other arguments will be ignored.
filesystem	A FileSystem object; if omitted, the FileSystem will be detected from x
format	A FileFormat object, or a string identifier of the format of the files in x. Currently supported values: • "parquet" • "ipc"/"arrow" for the Arrow IPC format (also supported as "feather" but this is deprecated) • "csv"/"text", aliases for the same thing (because comma is the default delimiter for text files) • "tsv", equivalent to passing format = "text", delimiter = "\t" Default is "parquet", unless a delimiter is also specified, in which case it is assumed to be "text".
partitioning	One of • A Schema, in which case the file paths relative to sources will be parsed, and path segments will be matched with the schema fields. For example, schema(year = int16(), month = int8()) would create partitions for file paths like "2019/01/file.parquet", "2019/02/file.parquet", etc. • A character vector that defines the field names corresponding to those path segments (that is, you're providing the names that would correspond to a Schema but the types will be autodetected) • A HivePartitioning or HivePartitioningFactory, as returned by hive_partition() which parses explicit or autodetected fields from Hive-style path segments • NULL for no partitioning
hive_style	Logical: if partitioning is a character vector or a Schema, should it be interpreted as specifying Hive-style partitioning? Default is NA, which means to inspect the file paths for Hive-style partitioning and behave accordingly.
factory_options	list of optional FileSystemFactoryOptions: • partition_base_dir: string path segment prefix to ignore when discovering partition information with DirectoryPartitioning. Not meaningful (ignored with a warning) for HivePartitioning, nor is it valid when providing a vector of file paths. • exclude_invalid_files: logical: should files that are not valid data files be excluded? Default is FALSE because checking all files up front incurs I/O and thus will be slower, especially on remote filesystems. If false and there are invalid files, there will be an error at scan time. This is the only FileSystemFactoryOption that is valid for both when providing a directory path in which to discover files and when providing a vector of file paths. • selector_ignore_prefixes: character vector of file prefixes to ignore when discovering files in a directory. If invalid files can be excluded by a

common filename prefix this way, you can avoid the I/O cost of `exclude_invalid_files`.
 Not valid when providing a vector of file paths (but if you're providing the
 file list, you can filter invalid files yourself).

... Additional format-specific options, passed to `FileFormat$create()`. For CSV
 options, note that you can specify them either with the Arrow C++ library nam-
 ing ("delimiter", "quoting", etc.) or the readr-style naming used in `read_csv_arrow()`
 ("delim", "quote", etc.). Not all readr options are currently supported; please
 file an issue if you encounter one that arrow should support.

Details

If you would only have a single `DatasetFactory` (for example, you have a single directory con-
 taining Parquet files), you can call `open_dataset()` directly. Use `dataset_factory()` when you
 want to combine different directories, file systems, or file formats.

Value

A `DatasetFactory` object. Pass this to `open_dataset()`, in a list potentially with other `DatasetFactory`
 objects, to create a `Dataset`.

DataType

DataType class

Description

DataType class

R6 Methods

- `$ToString()`: String representation of the `DataType`
- `$Equals(other)`: Is the `DataType` equal to other
- `$fields()`: The children fields associated with this type
- `$code(namespace)`: Produces an R call of the data type. Use `namespace=TRUE` to call with
`arrow::`.

There are also some active bindings:

- `$id`: integer Arrow type id.
- `$name`: string Arrow type name.
- `$num_fields`: number of child fields.

See Also

[infer_type\(\)](#)

[data-type](#)

dictionary	<i>Create a dictionary type</i>
------------	---------------------------------

Description

Create a dictionary type

Usage

```
dictionary(index_type = int32(), value_type = utf8(), ordered = FALSE)
```

Arguments

index_type	A <code>DataType</code> for the indices (default <code>int32()</code>)
value_type	A <code>DataType</code> for the values (default <code>utf8()</code>)
ordered	Is this an ordered dictionary (default <code>FALSE</code>)?

Value

A `DictionaryType`

See Also

[Other Arrow data types](#)

DictionaryType	<i>DictionaryType class</i>
----------------	-----------------------------

Description

`DictionaryType` is a `FixedWidthType` that represents dictionary-encoded data. Dictionary encoding stores unique values in a dictionary and uses integer-type indices to reference them, which can be more memory-efficient for data with many repeated values.

R6 Methods

- `$ToString()`: Return a string representation of the dictionary type
- `$code(namespace = FALSE)`: Return R code to create this dictionary type

Active bindings

- `$index_type`: The `DataType` for the dictionary indices (must be an integer type, signed or unsigned)
- `$value_type`: The `DataType` for the dictionary values
- `$name`: The name of the type.
- `$ordered`: Whether the dictionary is ordered.

Factory

DictionaryType\$create() takes the following arguments:

- index_type: A [DataType](#) for the indices (default [int32\(\)](#))
- value_type: A [DataType](#) for the values (default [utf8\(\)](#))
- ordered: Is this an ordered dictionary (default FALSE)?

Expression	<i>Arrow expressions</i>
------------	--------------------------

Description

Expressions are used to define filter logic for passing to a [Dataset Scanner](#).

Expression\$scalar(x) constructs an Expression which always evaluates to the provided scalar (length-1) R value.

Expression\$field_ref(name) is used to construct an Expression which evaluates to the named column in the Dataset against which it is evaluated.

Expression\$create(function_name, ..., options) builds a function-call Expression containing one or more Expressions. Anything in ... that is not already an expression will be wrapped in Expression\$scalar().

Expression\$op(FUN, ...) is for logical and arithmetic operators. Scalar inputs in ... will be attempted to be cast to the common type of the Expressions in the call so that the types of the columns in the Dataset are preserved and not unnecessarily upcast, which may be expensive.

ExtensionArray	<i>ExtensionArray class</i>
----------------	-----------------------------

Description

ExtensionArray class

Methods

The ExtensionArray class inherits from Array, but also provides access to the underlying storage of the extension.

- \$storage(): Returns the underlying [Array](#) used to store values.

The ExtensionArray is not intended to be subclassed for extension types.

ExtensionType

ExtensionType class

Description

ExtensionType class

Methods

The ExtensionType class inherits from DataType, but also defines extra methods specific to extension types:

- `$storage_type()`: Returns the underlying [DataType](#) used to store values.
- `$storage_id()`: Returns the [Type](#) identifier corresponding to the `$storage_type()`.
- `$extension_name()`: Returns the extension name.
- `$extension_metadata()`: Returns the serialized version of the extension metadata as a [raw\(\)](#) vector.
- `$extension_metadata_utf8()`: Returns the serialized version of the extension metadata as a UTF-8 encoded string.
- `$WrapArray(array)`: Wraps a storage [Array](#) into an [ExtensionArray](#) with this extension type.

In addition, subclasses may override the following methods to customize the behaviour of extension classes.

- `$deserialize_instance()`: This method is called when a new [ExtensionType](#) is initialized and is responsible for parsing and validating the serialized `extension_metadata` (a [raw\(\)](#) vector) such that its contents can be inspected by fields and/or methods of the R6 ExtensionType subclass. Implementations must also check the `storage_type` to make sure it is compatible with the extension type.
- `$as_vector(extension_array)`: Convert an [Array](#) or [ChunkedArray](#) to an R vector. This method is called by `as.vector()` on [ExtensionArray](#) objects, when a [RecordBatch](#) containing an [ExtensionArray](#) is converted to a `data.frame()`, or when a [ChunkedArray](#) (e.g., a column in a [Table](#)) is converted to an R vector. The default method returns the converted storage array.
- `$ToString()` Return a string representation that will be printed to the console when this type or an Array of this type is printed.

FeatherReader

FeatherReader class

Description

This class enables you to interact with Feather files. Create one to connect to a file or other Input-Stream, and call `Read()` on it to make an `arrow::Table`. See its usage in [read_feather\(\)](#).

Factory

The FeatherReader\$create() factory method instantiates the object and takes the following argument:

- file an Arrow file connection object inheriting from RandomAccessFile.

Methods

- \$Read(columns): Returns a Table of the selected columns, a vector of integer indices
- \$column_names: Active binding, returns the column names in the Feather file
- \$schema: Active binding, returns the schema of the Feather file
- \$version: Active binding, returns 1 or 2, according to the Feather file version

field	Create a Field
-------	----------------

Description

Create a Field

Usage

```
field(name, type, metadata = NULL, nullable = TRUE)
```

Arguments

name	field name
type	logical type, instance of DataType
metadata	a named character vector or list to attach as field metadata. All values will be coerced to character.
nullable	TRUE if field is nullable

See Also

[Field](#)

Examples

```
field("x", int32())
field("x", int32(), metadata = list(key = "value"))
```

Field

*Field class***Description**

`field()` lets you create an `arrow::Field` that maps a [DataType](#) to a column name. Fields are contained in [Schemas](#).

Methods

- `f$ToString()`: convert to a string
- `f$Equals(other, check_metadata = FALSE)`: test for equality. More naturally called as `f == other`
- `f$WithMetadata(metadata)`: returns a new `Field` with the key-value metadata set. Note that all list elements in `metadata` will be coerced to character.
- `f$RemoveMetadata()`: returns a new `Field` without metadata.

Active bindings

- `$HasMetadata`: logical: does this `Field` have extra metadata?
- `$metadata`: returns the key-value metadata as a named list, or `NULL` if no metadata is set.

FileFormat

*Dataset file formats***Description**

A `FileFormat` holds information about how to read and parse the files included in a `Dataset`. There are subclasses corresponding to the supported file formats (`ParquetFileFormat` and `IpcFileFormat`).

Factory

`FileFormat$create()` takes the following arguments:

- `format`: A string identifier of the file format. Currently supported values:
 - `"parquet"`
 - `"ipc"/"arrow"` for the Arrow IPC format (also supported as `"feather"` but this is deprecated)
 - `"csv"/"text"`, aliases for the same thing (because comma is the default delimiter for text files)
 - `"tsv"`, equivalent to passing `format = "text"`, `delimiter = "\t"`
- `...`: Additional format-specific options
 - `format = "parquet"`:
 - `dict_columns`: Names of columns which should be read as dictionaries.

- Any Parquet options from [FragmentScanOptions](#).
- format = "text": see [CsvParseOptions](#). Note that you can specify them either with the Arrow C++ library naming ("delimiter", "quoting", etc.) or the readr-style naming used in [read_csv_arrow\(\)](#) ("delim", "quote", etc.). Not all readr options are currently supported; please file an issue if you encounter one that arrow should support. Also, the following options are supported. From [CsvReadOptions](#):
- skip_rows
 - column_names. Note that if a [Schema](#) is specified, column_names must match those specified in the schema.
 - autogenerate_column_names From [CsvFragmentScanOptions](#) (these values can be overridden at scan time):
 - convert_options: a [CsvConvertOptions](#)
 - block_size

It returns the appropriate subclass of FileFormat (e.g. ParquetFileFormat)

Examples

```
## Semi-colon delimited files
# Set up directory for examples
tf <- tempfile()
dir.create(tf)
on.exit(unlink(tf))
write.table(mtcars, file.path(tf, "file1.txt"), sep = ";", row.names = FALSE)

# Create FileFormat object
format <- FileFormat$create(format = "text", delimiter = ";")

open_dataset(tf, format = format)
```

FileInfo

FileSystem entry info

Description

FileSystem entry info

Methods

- `base_name()` : The file base name (component after the last directory separator).
- `extension()` : The file extension

Active bindings

- `$type`: The file type
- `$path`: The full file path in the filesystem
- `$size`: The size in bytes, if available. Only regular files are guaranteed to have a size.
- `$mtime`: The time of last modification, if available.

FileSelector	<i>file selector</i>
--------------	----------------------

Description

file selector

Factory

The `$create()` factory method instantiates a `FileSelector` given the 3 fields described below.

Fields

- `base_dir`: The directory in which to select files. If the path exists but doesn't point to a directory, this should be an error.
- `allow_not_found`: The behavior if `base_dir` doesn't exist in the filesystem. If `FALSE`, an error is returned. If `TRUE`, an empty selection is returned
- `recursive`: Whether to recurse into subdirectories.

FileSystem	<i>FileSystem classes</i>
------------	---------------------------

Description

`FileSystem` is an abstract file system API, `LocalFileSystem` is an implementation accessing files on the local machine. `SubTreeFileSystem` is an implementation that delegates to another implementation after prepending a fixed base path

Factory

`LocalFileSystem$create()` returns the object and takes no arguments.

`SubTreeFileSystem$create()` takes the following arguments:

- `base_path`, a string path
- `base_fs`, a `FileSystem` object

`S3FileSystem$create()` optionally takes arguments:

- `anonymous`: logical, default `FALSE`. If true, will not attempt to look up credentials using standard AWS configuration methods.
- `access_key`, `secret_key`: authentication credentials. If one is provided, the other must be as well. If both are provided, they will override any AWS configuration set at the environment level.
- `session_token`: optional string for authentication along with `access_key` and `secret_key`

- `role_arn`: string AWS ARN of an AccessRole. If provided instead of `access_key` and `secret_key`, temporary credentials will be fetched by assuming this role.
- `session_name`: optional string identifier for the assumed role session.
- `external_id`: optional unique string identifier that might be required when you assume a role in another account.
- `load_frequency`: integer, frequency (in seconds) with which temporary credentials from an assumed role session will be refreshed. Default is 900 (i.e. 15 minutes)
- `region`: AWS region to connect to. If omitted, the AWS library will provide a sensible default based on client configuration, falling back to "us-east-1" if no other alternatives are found.
- `endpoint_override`: If non-empty, override region with a connect string such as "localhost:9000". This is useful for connecting to file systems that emulate S3.
- `scheme`: S3 connection transport (default "https")
- `proxy_options`: optional string, URI of a proxy to use when connecting to S3
- `background_writes`: logical, whether OutputStream writes will be issued in the background, without blocking (default TRUE)
- `allow_bucket_creation`: logical, if TRUE, the filesystem will create buckets if `$CreateDir()` is called on the bucket level (default FALSE).
- `allow_bucket_deletion`: logical, if TRUE, the filesystem will delete buckets if `$DeleteDir()` is called on the bucket level (default FALSE).
- `check_directory_existence_before_creation`: logical, check if directory already exists or not before creation. Helpful for cloud storage operations where object mutation operations are rate limited or existing directories are read-only. (default FALSE).
- `request_timeout`: Socket read time on Windows and macOS in seconds. If negative, the AWS SDK default (typically 3 seconds).
- `connect_timeout`: Socket connection timeout in seconds. If negative, AWS SDK default is used (typically 1 second).

`GcsFileSystem$create()` optionally takes arguments:

- `anonymous`: logical, default FALSE. If true, will not attempt to look up credentials using standard GCS configuration methods.
- `access_token`: optional string for authentication. Should be provided along with `expiration`
- `expiration`: POSIXct. optional datetime representing point at which `access_token` will expire.
- `json_credentials`: optional string for authentication. Either a string containing JSON credentials or a path to their location on the filesystem. If a path to credentials is given, the file should be UTF-8 encoded.
- `endpoint_override`: if non-empty, will connect to provided host name / port, such as "localhost:9001", instead of default GCS ones. This is primarily useful for testing purposes.
- `scheme`: connection transport (default "https")
- `default_bucket_location`: the default location (or "region") to create new buckets in.
- `retry_limit_seconds`: the maximum amount of time to spend retrying if the filesystem encounters errors. Default is 15 seconds.

- `default_metadata`: default metadata to write in new objects.
- `project_id`: the project to use for creating buckets.

`AzureFileSystem$create()` takes following required argument:

- `account_name`: Azure Blob Storage account name.

`AzureFileSystem$create()` takes following optional arguments:

- `account_key`: Account key of the storage account. Cannot be used with `sas_token`.
- `blob_storage_authority`: Hostname of the blob service, defaulting to `"blob.core.windows.net"`.
- `blob_storage_scheme`: Either `"http"` or `"https"` (the default).
- `client_id`: The client/application ID for Azure Active Directory authentication. If used with `client_secret` and `tenant_id` then it is the application ID for a registered Azure AD application. Otherwise, it is the client ID of a user-assigned managed identity.
- `client_secret`: Client secret for Azure Active Directory authentication. Must be provided with both `client_id` and `tenant_id`.
- `dfs_storage_authority`: Hostname of the data lake (gen 2) service, defaulting to `"dfs.core.windows.net"`.
- `dfs_storage_scheme`: Either `"http"` or `"https"` (the default).
- `sas_token`: Shared access signature (SAS) token for the storage account. Cannot be used with `account_key`.
- `tenant_id`: Tenant ID for Azure Active Directory authentication. Must be provided with both `client_id` and `client_secret`.

Methods

- `path(x)`: Create a `SubTreeFileSystem` from the current `FileSystem` rooted at the specified path `x`.
- `cd(x)`: Create a `SubTreeFileSystem` from the current `FileSystem` rooted at the specified path `x`.
- `ls(path, ...)`: List files or objects at the given path or from the root of the `FileSystem` if path is not provided. Additional arguments passed to `FileSelector$create`, see [FileSelector](#).
- `$GetFileInfo(x)`: `x` may be a [FileSelector](#) or a character vector of paths. Returns a list of [FileInfo](#).
- `$CreateDir(path, recursive = TRUE)`: Create a directory and subdirectories.
- `$DeleteDir(path)`: Delete a directory and its contents, recursively.
- `$DeleteDirContents(path)`: Delete a directory's contents, recursively. Like `$DeleteDir()`, but doesn't delete the directory itself. Passing an empty path `""` will wipe the entire filesystem tree.
- `$DeleteFile(path)`: Delete a file.
- `$DeleteFiles(paths)`: Delete many files. The default implementation issues individual delete operations in sequence.
- `$Move(src, dest)`: Move / rename a file or directory. If the destination exists: if it is a non-empty directory, an error is returned otherwise, if it has the same type as the source, it is replaced otherwise, behavior is unspecified (implementation-dependent).

- `$CopyFile(src, dest)`: Copy a file. If the destination exists and is a directory, an error is returned. Otherwise, it is replaced.
- `$OpenInputStream(path)`: Open an [input stream](#) for sequential reading.
- `$OpenInputFile(path)`: Open an [input file](#) for random access reading.
- `$OpenOutputStream(path)`: Open an [output stream](#) for sequential writing.
- `$OpenAppendStream(path)`: Open an [output stream](#) for appending.

Active bindings

- `$type_name`: string filesystem type name, such as "local", "s3", etc.
- `$region`: string AWS region, for `S3FileSystem` and `SubTreeFileSystem` containing a `S3FileSystem`
- `$base_fs`: for `SubTreeFileSystem`, the `FileSystem` it contains
- `$base_path`: for `SubTreeFileSystem`, the path in `$base_fs` which is considered root in this `SubTreeFileSystem`.
- `$options`: for `GcsFileSystem`, the options used to create the `GcsFileSystem` instance as a list

Notes

On `S3FileSystem`, `$CreateDir()` on a top-level directory creates a new bucket. When `S3FileSystem` creates new buckets (assuming `allow_bucket_creation` is `TRUE`), it does not pass any non-default settings. In AWS S3, the bucket and all objects will be not publicly visible, and will have no bucket policies and no resource tags. To have more control over how buckets are created, use a different API to create them.

On `S3FileSystem`, output is only produced for fatal errors or when printing return values. For troubleshooting, the log level can be set using the environment variable `ARROW_S3_LOG_LEVEL` (e.g., `Sys.setenv("ARROW_S3_LOG_LEVEL"="DEBUG")`). The log level must be set prior to running any code that interacts with S3. Possible values include 'FATAL' (the default), 'ERROR', 'WARN', 'INFO', 'DEBUG' (recommended), 'TRACE', and 'OFF'.

On `AzureFileSystem`, passing no arguments for authentication uses the `AzureDefaultCredential` for authentication, so that several authentication types are tried until one succeeds.

`AzureFileSystem` is not presently supported on Windows due to upstream compatibility issues between the Azure C++ SDK and the MinGW toolchain.

FileWriteOptions	<i>Format-specific write options</i>
------------------	--------------------------------------

Description

A `FileWriteOptions` holds write options specific to a `FileFormat`.

flight_connect	<i>Connect to a Flight server</i>
----------------	-----------------------------------

Description

Connect to a Flight server

Usage

```
flight_connect(host = "localhost", port, scheme = "grpc+tcp")
```

Arguments

host	string hostname to connect to
port	integer port to connect on
scheme	URL scheme, default is "grpc+tcp"

Value

A `pyarrow.flight.FlightClient`.

flight_disconnect	<i>Explicitly close a Flight client</i>
-------------------	---

Description

Explicitly close a Flight client

Usage

```
flight_disconnect(client)
```

Arguments

client	The client to disconnect
--------	--------------------------

flight_get	Get data from a Flight server
------------	-------------------------------

Description

Get data from a Flight server

Usage

flight_get(client, path)

Arguments

client	pyarrow.flight.FlightClient, as returned by flight_connect()
path	string identifier under which data is stored

Value

A [Table](#)

flight_put	Send data to a Flight server
------------	------------------------------

Description

Send data to a Flight server

Usage

flight_put(client, data, path, overwrite = TRUE, max_chunksize = NULL)

Arguments

client	pyarrow.flight.FlightClient, as returned by flight_connect()
data	data.frame, RecordBatch , or Table to upload
path	string identifier to store the data under
overwrite	logical: if path exists on client already, should we replace it with the contents of data? Default is TRUE; if FALSE and path exists, the function will error.
max_chunksize	integer: Maximum number of rows for RecordBatch chunks when a data.frame is sent. Individual chunks may be smaller depending on the chunk layout of individual columns.

Value

client, invisibly.

FragmentScanOptions	<i>Format-specific scan options</i>
---------------------	-------------------------------------

Description

A FragmentScanOptions holds options specific to a FileFormat and a scan operation.

Factory

FragmentScanOptions#create() takes the following arguments:

- format: A string identifier of the file format. Currently supported values:
 - "parquet"
 - "csv"/"text", aliases for the same format.
- ...: Additional format-specific options
 - format = "parquet":
 - use_buffered_stream: Read files through buffered input streams rather than loading entire row groups at once. This may be enabled to reduce memory overhead. Disabled by default.
 - buffer_size: Size of buffered stream, if enabled. Default is 8KB.
 - pre_buffer: Pre-buffer the raw Parquet data. This can improve performance on high-latency filesystems. Disabled by default.
 - thrift_string_size_limit: Maximum string size allocated for decoding thrift strings. May need to be increased in order to read files with especially large headers. Default value 100000000.
 - thrift_container_size_limit: Maximum size of thrift containers. May need to be increased in order to read files with especially large headers. Default value 1000000.
 - format = "text": see [CsvConvertOptions](#). Note that options can only be specified with the Arrow C++ library naming. Also, "block_size" from [CsvReadOptions](#) may be given.

It returns the appropriate subclass of FragmentScanOptions (e.g. CsvFragmentScanOptions).

gs_bucket	<i>Connect to a Google Cloud Storage (GCS) bucket</i>
-----------	---

Description

gs_bucket() is a convenience function to create an GcsFileSystem object that holds onto its relative path

Usage

```
gs_bucket(bucket, ...)
```

Arguments

bucket string GCS bucket name or path
 ... Additional connection options, passed to `GcsFileSystem$create()`

Value

A `SubTreeFileSystem` containing an `GcsFileSystem` and the bucket's relative path. Note that this function's success does not guarantee that you are authorized to access the bucket's contents.

Examples

```
bucket <- gs_bucket("arrow-datasets")
```

hive_partition	<i>Construct Hive partitioning</i>
----------------	------------------------------------

Description

Hive partitioning embeds field names and values in path segments, such as `"/year=2019/month=2/data.parquet"`.

Usage

```
hive_partition(..., null_fallback = NULL, segment_encoding = "uri")
```

Arguments

... named list of [data types](#), passed to [schema\(\)](#)
 null_fallback character to be used in place of missing values (NA or NULL) in partition columns. Default is `"__HIVE_DEFAULT_PARTITION__"`, which is what Hive uses.
 segment_encoding Decode partition segments after splitting paths. Default is `"uri"` (URI-decode segments). May also be `"none"` (leave as-is).

Details

Because fields are named in the path segments, order of fields passed to `hive_partition()` does not matter.

Value

A [HivePartitioning](#), or a `HivePartitioningFactory` if calling `hive_partition()` with no arguments.

Examples

```
hive_partition(year = int16(), month = int8())
```

infer_schema	<i>Extract a schema from an object</i>
--------------	--

Description

Extract a schema from an object

Usage

```
infer_schema(x)
```

Arguments

x	An object which has a schema, e.g. a Dataset
---	--

infer_type	<i>Infer the arrow Array type from an R object</i>
------------	--

Description

Infer the arrow Array type from an R object

Usage

```
infer_type(x, ...)
```

Arguments

x	an R object (usually a vector) to be converted to an Array or ChunkedArray .
...	Passed to S3 methods

Value

An arrow [data type](#)

Examples

```
infer_type(1:10)
infer_type(1L:10L)
infer_type(c(1, 1.5, 2))
infer_type(c("A", "B", "C"))
infer_type(mtcars)
infer_type(Sys.Date())
infer_type(as.POSIXlt(Sys.Date()))
infer_type(vctrs::new_vctr(1:5, class = "my_custom_vctr_class"))
```

 InputStream

InputStream classes

Description

RandomAccessFile inherits from InputStream and is a base class for: ReadableFile for reading from a file; MemoryMappedFile for the same but with memory mapping; and BufferedReader for reading from a buffer. Use these with the various table readers.

Factory

The \$create() factory methods instantiate the InputStream object and take the following arguments, depending on the subclass:

- path For ReadableFile, a character file name
- x For BufferedReader, a [Buffer](#) or an object that can be made into a buffer via buffer().

To instantiate a MemoryMappedFile, call [mmap_open\(\)](#).

Methods

- \$GetSize():
- \$supports_zero_copy(): Logical
- \$seek(position): go to that position in the stream
- \$tell(): return the position in the stream
- \$close(): close the stream
- \$Read(nbytes): read data from the stream, either a specified nbytes or all, if nbytes is not provided
- \$ReadAt(position, nbytes): similar to \$seek(position)\$Read(nbytes)
- \$Resize(size): for a MemoryMappedFile that is writeable

 install_arrow

Install or upgrade the Arrow library

Description

Use this function to install the latest release of arrow, to switch to or from a nightly development version, or on Linux to try reinstalling with all necessary C++ dependencies.

Usage

```
install_arrow(
  nightly = FALSE,
  binary = Sys.getenv("LIBARROW_BINARY", TRUE),
  use_system = Sys.getenv("ARROW_USE_PKG_CONFIG", FALSE),
  minimal = Sys.getenv("LIBARROW_MINIMAL", FALSE),
  verbose = Sys.getenv("ARROW_R_DEV", FALSE),
  repos = getOption("repos"),
  ...
)
```

Arguments

nightly	logical: Should we install a development version of the package, or should we install from CRAN (the default).
binary	On Linux, value to set for the environment variable LIBARROW_BINARY, which governs how C++ binaries are used, if at all. The default value, TRUE, tells the installation script to detect the Linux distribution and version and find an appropriate C++ library. FALSE would tell the script not to retrieve a binary and instead build Arrow C++ from source. Other valid values are strings corresponding to a Linux distribution-version, to override the value that would be detected. See the install guide for further details.
use_system	logical: Should we use pkg-config to look for Arrow system packages? Default is FALSE. If TRUE, source installation may be faster, but there is a risk of version mismatch. This sets the ARROW_USE_PKG_CONFIG environment variable.
minimal	logical: If building from source, should we build without optional dependencies (compression libraries, for example)? Default is FALSE. This sets the LIBARROW_MINIMAL environment variable.
verbose	logical: Print more debugging output when installing? Default is FALSE. This sets the ARROW_R_DEV environment variable.
repos	character vector of base URLs of the repositories to install from (passed to <code>install.packages()</code>)
...	Additional arguments passed to <code>install.packages()</code>

Details

Note that, unlike packages like `tensorflow`, `blogdown`, and others that require external dependencies, you do not need to run `install_arrow()` after a successful arrow installation.

See Also

[arrow_info\(\)](#) to see if the package was configured with necessary C++ dependencies. [install guide](#) for more ways to tune installation on Linux.

install_pyarrow	<i>Install pyarrow for use with reticulate</i>
-----------------	--

Description

pyarrow is the Python package for Apache Arrow. This function helps with installing it for use with reticulate.

Usage

```
install_pyarrow(envname = NULL, nightly = FALSE, ...)
```

Arguments

envname	The name or full path of the Python environment to install into. This can be a virtualenv or conda environment created by reticulate. See <code>reticulate::py_install()</code> .
nightly	logical: Should we install a development version of the package? Default is to use the official release version.
...	additional arguments passed to <code>reticulate::py_install()</code> .

io_thread_count	<i>Manage the global I/O thread pool in libarrow</i>
-----------------	--

Description

Manage the global I/O thread pool in libarrow

Usage

```
io_thread_count()

set_io_thread_count(num_threads)
```

Arguments

num_threads	integer: New number of threads for thread pool. At least two threads are recommended to support all operations in the arrow package.
-------------	--

JsonFileFormat	<i>JSON dataset file format</i>
----------------	---------------------------------

Description

A `JsonFileFormat` is a [FileFormat](#) subclass which holds information about how to read and parse the files included in a JSON Dataset.

Value

A `JsonFileFormat` object

Factory

`JsonFileFormat$create()` can take options in the form of lists passed through as `parse_options`, or `read_options` parameters.

Available `read_options` parameters:

- `use_threads`: Whether to use the global CPU thread pool. Default `TRUE`. If `FALSE`, JSON input must end with an empty line.
- `block_size`: Block size we request from the IO layer; also determines size of chunks when `use_threads` is `TRUE`.

Available `parse_options` parameters:

- `newlines_in_values`: Logical: are values allowed to contain CR (`0x0d` or `\r`) and LF (`0x0a` or `\n`) characters? (default `FALSE`)

See Also

[FileFormat](#)

<code>list_compute_functions</code>

List available Arrow C++ compute functions

Description

This function lists the names of all available Arrow C++ library compute functions. These can be called by passing to [call_function\(\)](#), or they can be called by name with an `arrow_` prefix inside a `dplyr` verb.

Usage

```
list_compute_functions(pattern = NULL, ...)
```

Arguments

pattern	Optional regular expression to filter the function list
...	Additional parameters passed to <code>grep()</code>

Details

The resulting list describes the capabilities of your arrow build. Some functions, such as string and regular expression functions, require optional build-time C++ dependencies. If your arrow package was not compiled with those features enabled, those functions will not appear in this list.

Some functions take options that need to be passed when calling them (in a list called `options`). These options require custom handling in C++; many functions already have that handling set up but not all do. If you encounter one that needs special handling for options, please report an issue.

Note that this list does *not* enumerate all of the R bindings for these functions. The package includes Arrow methods for many base R functions that can be called directly on Arrow objects, as well as some tidyverse-flavored versions available inside `dplyr` verbs.

Value

A character vector of available Arrow C++ function names

See Also

[acero](#) for R bindings for Arrow functions

Examples

```
available_funcs <- list_compute_functions()
utf8_funcs <- list_compute_functions(pattern = "^UTF8", ignore.case = TRUE)
```

list_flights	<i>See available resources on a Flight server</i>
--------------	---

Description

See available resources on a Flight server

Usage

```
list_flights(client)

flight_path_exists(client, path)
```

Arguments

client	pyarrow.flight.FlightClient, as returned by flight_connect()
path	string identifier under which data is stored

Value

list_flights() returns a character vector of paths. flight_path_exists() returns a logical value, the equivalent of path %in% list_flights()

load_flight_server	<i>Load a Python Flight server</i>
--------------------	------------------------------------

Description

Load a Python Flight server

Usage

```
load_flight_server(name, path = system.file(package = "arrow"))
```

Arguments

name	string Python module name
path	file system path where the Python module is found. Default is to look in the inst/ directory for included modules.

Examples

```
load_flight_server("demo_flight_server")
```

map_batches	<i>Apply a function to a stream of RecordBatches</i>
-------------	--

Description

As an alternative to calling collect() on a Dataset query, you can use this function to access the stream of RecordBatches in the Dataset. This lets you do more complex operations in R that operate on chunks of data without having to hold the entire Dataset in memory at once. You can include map_batches() in a dplyr pipeline and do additional dplyr methods on the stream of data in Arrow after it.

Usage

```
map_batches(X, FUN, ..., .schema = NULL, .lazy = TRUE)
```

Arguments

x	A Dataset or arrow_dplyr_query object, as returned by the dplyr methods on Dataset.
FUN	A function or purrr-style lambda expression to apply to each batch. It must return a RecordBatch or something coercible to one via 'as_record_batch()'.
...	Additional arguments passed to FUN
.schema	An optional <code>schema()</code> . If NULL, the schema will be inferred from the first batch.
.lazy	Use TRUE to evaluate FUN lazily as batches are read from the result; use FALSE to evaluate FUN on all batches before returning the reader.

Details

This is experimental and not recommended for production use. It is also single-threaded and runs in R not C++, so it won't be as fast as core Arrow methods.

Value

An arrow_dplyr_query.

match_arrow	<i>Value matching for Arrow objects</i>
-------------	---

Description

`base::match()` and `base::%in%` are not generics, so we can't just define Arrow methods for them. These functions expose the analogous functions in the Arrow C++ library.

Usage

```
match_arrow(x, table, ...)

is_in(x, table, ...)
```

Arguments

x	Scalar, Array or ChunkedArray
table	Scalar, Array, 'ChunkedArray', or R vector lookup table.
...	additional arguments, ignored

Value

`match_arrow()` returns an int32-type Arrow object of the same length and type as x with the (0-based) indexes into table. `is_in()` returns a boolean-type Arrow object of the same length and type as x with values indicating per element of x it is present in table.

Examples

```
# note that the returned value is 0-indexed
cars_tbl <- arrow_table(name = rownames(mtcars), mtcars)
match_arrow(Scalar$create("Mazda RX4 Wag"), cars_tbl$name)

is_in(Array$create("Mazda RX4 Wag"), cars_tbl$name)

# Although there are multiple matches, you are returned the index of the first
# match, as with the base R equivalent
match(4, mtcars$cyl) # 1-indexed
match_arrow(Scalar$create(4), cars_tbl$cyl) # 0-indexed

# If `x` contains multiple values, you are returned the indices of the first
# match for each value.
match(c(4, 6, 8), mtcars$cyl)
match_arrow(Array$create(c(4, 6, 8)), cars_tbl$cyl)

# Return type matches type of `x`
is_in(c(4, 6, 8), mtcars$cyl) # returns vector
is_in(Scalar$create(4), mtcars$cyl) # returns Scalar
is_in(Array$create(c(4, 6, 8)), cars_tbl$cyl) # returns Array
is_in(ChunkedArray$create(c(4, 6, 8), cars_tbl$cyl) # returns ChunkedArray
```

Message

Message class

Description

Message holds an Arrow IPC message, which includes metadata and an optional message body.

R6 Methods

- `$Equals(other)`: Check if this Message is equal to another Message
- `$body_length()`: Return the length of the message body in bytes
- `$Verify()`: Check if the Message metadata is valid Flatbuffer format

Active bindings

- `$type`: The message type
- `$metadata`: The message metadata
- `$body`: The message body as a [Buffer](#)

MessageReader	<i>MessageReader class</i>
---------------	----------------------------

Description

MessageReader reads Message objects from an input stream.

R6 Methods

- `$ReadNextMessage()`: Read the next Message from the stream. Returns NULL if there are no more messages.

Factory

MessageReader\$create() takes the following argument:

- stream: An [InputStream](#) or object coercible to one (e.g., a raw vector)

mmap_create	<i>Create a new read/write memory mapped file of a given size</i>
-------------	---

Description

Create a new read/write memory mapped file of a given size

Usage

```
mmap_create(path, size)
```

Arguments

path	file path
size	size in bytes

Value

a [arrow::io::MemoryMappedFile](#)

mmap_open	<i>Open a memory mapped file</i>
-----------	----------------------------------

Description

Open a memory mapped file

Usage

```
mmap_open(path, mode = c("read", "write", "readwrite"))
```

Arguments

path	file path
mode	file mode (read/write/readwrite)

new_extension_type	<i>Extension types</i>
--------------------	------------------------

Description

Extension arrays are wrappers around regular Arrow [Array](#) objects that provide some customized behaviour and/or storage. A common use-case for extension types is to define a customized conversion between an Arrow [Array](#) and an R object when the default conversion is slow or loses metadata important to the interpretation of values in the array. For most types, the built-in [vctrs extension type](#) is probably sufficient.

Usage

```
new_extension_type(
  storage_type,
  extension_name,
  extension_metadata = raw(),
  type_class = ExtensionType
)

new_extension_array(storage_array, extension_type)

register_extension_type(extension_type)

reregister_extension_type(extension_type)

unregister_extension_type(extension_name)
```

Arguments

storage_type	The data type of the underlying storage array.
extension_name	The extension name. This should be namespaced using "dot" syntax (i.e., "some_package.some_type"). The namespace "arrow" is reserved for extension types defined by the Apache Arrow libraries.
extension_metadata	A raw() or character() vector containing the serialized version of the type. Character vectors must be length 1 and are converted to UTF-8 before converting to raw() .
type_class	An R6::R6Class whose <code>\$new()</code> class method will be used to construct a new instance of the type.
storage_array	An Array object of the underlying storage.
extension_type	An ExtensionType instance.

Details

These functions create, register, and unregister [ExtensionType](#) and [ExtensionArray](#) objects. To use an extension type you will have to:

- Define an [R6::R6Class](#) that inherits from [ExtensionType](#) and reimplement one or more methods (e.g., `deserialize_instance()`).
- Make a type constructor function (e.g., `my_extension_type()`) that calls [new_extension_type\(\)](#) to create an R6 instance that can be used as a [data type](#) elsewhere in the package.
- Make an array constructor function (e.g., `my_extension_array()`) that calls [new_extension_array\(\)](#) to create an [Array](#) instance of your extension type.
- Register a dummy instance of your extension type created using your constructor function using [register_extension_type\(\)](#).

If defining an extension type in an R package, you will probably want to use [reregister_extension_type\(\)](#) in that package's `.onLoad()` hook since your package will probably get reloaded in the same R session during its development and [register_extension_type\(\)](#) will error if called twice for the same `extension_name`. For an example of an extension type that uses most of these features, see [vctrs_extension_type\(\)](#).

Value

- `new_extension_type()` returns an [ExtensionType](#) instance according to the `type_class` specified.
- `new_extension_array()` returns an [ExtensionArray](#) whose `$type` corresponds to `extension_type`.
- `register_extension_type()`, `unregister_extension_type()` and `reregister_extension_type()` return `NULL`, invisibly.

Examples

```

# Create the R6 type whose methods control how Array objects are
# converted to R objects, how equality between types is computed,
# and how types are printed.
QuantizedType <- R6::R6Class(
  "QuantizedType",
  inherit = ExtensionType,
  public = list(
    # methods to access the custom metadata fields
    center = function() private$.center,
    scale = function() private$.scale,

    # called when an Array of this type is converted to an R vector
    as_vector = function(extension_array) {
      if (inherits(extension_array, "ExtensionArray")) {
        unquantized_arrow <-
          (extension_array$storage()$cast(float64()) / private$.scale) +
            private$.center

        as_vector(unquantized_arrow)
      } else {
        super$as_vector(extension_array)
      }
    },

    # populate the custom metadata fields from the serialized metadata
    deserialize_instance = function() {
      vals <- as.numeric(strsplit(self$extension_metadata_utf8(), ";")[[1]])
      private$.center <- vals[1]
      private$.scale <- vals[2]
    }
  ),
  private = list(
    .center = NULL,
    .scale = NULL
  )
)

# Create a helper type constructor that calls new_extension_type()
quantized <- function(center = 0, scale = 1, storage_type = int32()) {
  new_extension_type(
    storage_type = storage_type,
    extension_name = "arrow.example.quantized",
    extension_metadata = paste(center, scale, sep = ";"),
    type_class = QuantizedType
  )
}

# Create a helper array constructor that calls new_extension_array()
quantized_array <- function(x, center = 0, scale = 1,
  storage_type = int32()) {
  type <- quantized(center, scale, storage_type)

```

```

    new_extension_array(
      Array$create((x - center) * scale, type = storage_type),
      type
    )
  }

  # Register the extension type so that Arrow knows what to do when
  # it encounters this extension type
  reregister_extension_type(quantized())

  # Create Array objects and use them!
  (vals <- runif(5, min = 19, max = 21))

  (array <- quantized_array(
    vals,
    center = 20,
    scale = 2^15 - 1,
    storage_type = int16()
  )
  )

  array$type$center()
  array$type$scale()

  as.vector(array)

```

open_dataset

Open a multi-file dataset

Description

Arrow Datasets allow you to query against data that has been split across multiple files. This sharding of data may indicate partitioning, which can accelerate queries that only touch some partitions (files). Call `open_dataset()` to point to a directory of data files and return a Dataset, then use dplyr methods to query it.

Usage

```

open_dataset(
  sources,
  schema = NULL,
  partitioning = hive_partition(),
  hive_style = NA,
  unify_schemas = NULL,
  format = c("parquet", "arrow", "ipc", "feather", "csv", "tsv", "text", "json"),
  factory_options = list(),
  ...
)

```

Arguments

sources	One of: • a string path or URI to a directory containing data files • a FileSystem that references a directory containing data files (such as what is returned by s3_bucket()) • a string path or URI to a single file • a character vector of paths or URIs to individual data files • a list of Dataset objects as created by this function • a list of DatasetFactory objects as created by dataset_factory(). When sources is a vector of file URIs, they must all use the same protocol and point to files located in the same file system and having the same format.
schema	Schema for the Dataset. If NULL (the default), the schema will be inferred from the data sources.
partitioning	When sources is a directory path/URI, one of: • a Schema, in which case the file paths relative to sources will be parsed, and path segments will be matched with the schema fields. • a character vector that defines the field names corresponding to those path segments (that is, you're providing the names that would correspond to a Schema but the types will be autodetected) • a Partitioning or PartitioningFactory, such as returned by hive_partition() • NULL for no partitioning The default is to autodetect Hive-style partitions unless <code>hive_style = FALSE</code>. See the "Partitioning" section for details. When sources is not a directory path/URI, partitioning is ignored.
hive_style	Logical: should partitioning be interpreted as Hive-style? Default is NA, which means to inspect the file paths for Hive-style partitioning and behave accordingly.
unify_schemas	logical: should all data fragments (files, Datasets) be scanned in order to create a unified schema from them? If FALSE, only the first fragment will be inspected for its schema. Use this fast path when you know and trust that all fragments have an identical schema. The default is FALSE when creating a dataset from a directory path/URI or vector of file paths/URIs (because there may be many files and scanning may be slow) but TRUE when sources is a list of Datasets (because there should be few Datasets in the list and their Schemas are already in memory).
format	A FileFormat object, or a string identifier of the format of the files in x. This argument is ignored when sources is a list of Dataset objects. Currently supported values: • "parquet" • "ipc"/"arrow" for the Arrow IPC format (also supported as "feather" but this is deprecated) • "csv"/"text", aliases for the same thing (because comma is the default delimiter for text files)

- "tsv", equivalent to passing `format = "text"`, `delimiter = "\t"`
- "json", for JSON format datasets Note: only newline-delimited JSON (aka ND-JSON) datasets are currently supported Default is "parquet", unless a delimiter is also specified, in which case it is assumed to be "text".

factory_options

list of optional `FileSystemFactoryOptions`:

- `partition_base_dir`: string path segment prefix to ignore when discovering partition information with `DirectoryPartitioning`. Not meaningful (ignored with a warning) for `HivePartitioning`, nor is it valid when providing a vector of file paths.
- `exclude_invalid_files`: logical: should files that are not valid data files be excluded? Default is FALSE because checking all files up front incurs I/O and thus will be slower, especially on remote filesystems. If false and there are invalid files, there will be an error at scan time. This is the only `FileSystemFactoryOption` that is valid for both when providing a directory path in which to discover files and when providing a vector of file paths.
- `selector_ignore_prefixes`: character vector of file prefixes to ignore when discovering files in a directory. If invalid files can be excluded by a common filename prefix this way, you can avoid the I/O cost of `exclude_invalid_files`. Not valid when providing a vector of file paths (but if you're providing the file list, you can filter invalid files yourself).

...

additional arguments passed to `dataset_factory()` when sources is a directory path/URI or vector of file paths/URIs, otherwise ignored. These may include format to indicate the file format, or other format-specific options (see [read_csv_arrow\(\)](#), [read_parquet\(\)](#) and [read_ipc_file\(\)](#) on how to specify these).

Value

A [Dataset](#) R6 object. Use dplyr methods on it to query the data, or call [\\$NewScan\(\)](#) to construct a query directly.

Partitioning

Data is often split into multiple files and nested in subdirectories based on the value of one or more columns in the data. It may be a column that is commonly referenced in queries, or it may be time-based, for some examples. Data that is divided this way is "partitioned," and the values for those partitioning columns are encoded into the file path segments. These path segments are effectively virtual columns in the dataset, and because their values are known prior to reading the files themselves, we can greatly speed up filtered queries by skipping some files entirely.

Arrow supports reading partition information from file paths in two forms:

- "Hive-style", deriving from the Apache Hive project and common to some database systems. Partitions are encoded as "key=value" in path segments, such as `"year=2019/month=1/file.parquet"`. While they may be awkward as file names, they have the advantage of being self-describing.
- "Directory" partitioning, which is Hive without the key names, like `"2019/01/file.parquet"`. In order to use these, we need know at least what names to give the virtual columns that come from the path segments.

The default behavior in `open_dataset()` is to inspect the file paths contained in the provided directory, and if they look like Hive-style, parse them as Hive. If your dataset has Hive-style partitioning in the file paths, you do not need to provide anything in the `partitioning` argument to `open_dataset()` to use them. If you do provide a character vector of partition column names, they will be ignored if they match what is detected, and if they don't match, you'll get an error. (If you want to rename partition columns, do that using `select()` or `rename()` after opening the dataset.). If you provide a Schema and the names match what is detected, it will use the types defined by the Schema. In the example file path above, you could provide a Schema to specify that "month" should be `int8()` instead of the `int32()` it will be parsed as by default. This is also useful for keeping leading zeros, so that a value such as 001 isn't parsed as the integer 1.

If your file paths do not appear to be Hive-style, or if you pass `hive_style = FALSE`, the `partitioning` argument will be used to create Directory partitioning. A character vector of names is required to create partitions; you may instead provide a Schema to map those names to desired column types, as described above. If neither are provided, no partitioning information will be taken from the file paths.

Adding the source filename as a column

Partitioning only recovers information encoded in directory names. If you need to know which file each row came from, call `add_filename()` inside a dplyr query on the dataset:

```
open_dataset("nyc-taxi") |>
  mutate(file = add_filename()) |>
  collect()
```

This is useful, for example, when you have opened a subdirectory of a partitioned dataset directly (so the partition columns above that directory are not inferred) and want to recover the partition values from the path. See `add_filename()` for details and limitations.

See Also

[datasets article](#), `add_filename()`

Examples

```
# Set up directory for examples
tf <- tempfile()
dir.create(tf)

write_dataset(mtcars, tf, partitioning = "cyl")

# You can specify a directory containing the files for your dataset and
# open_dataset will scan all files in your directory.
open_dataset(tf)

# You can also supply a vector of paths
open_dataset(c(file.path(tf, "cyl=4/part-0.parquet"), file.path(tf, "cyl=8/part-0.parquet")))

## You must specify the file format if using a format other than parquet.
tf2 <- tempfile()
```

```

dir.create(tf2)
write_dataset(mtcars, tf2, format = "ipc")
# This line will results in errors when you try to work with the data
## Not run:
open_dataset(tf2)

## End(Not run)
# This line will work
open_dataset(tf2, format = "ipc")

## You can specify file partitioning to include it as a field in your dataset
# Create a temporary directory and write example dataset
tf3 <- tempfile()
dir.create(tf3)
write_dataset(airquality, tf3, partitioning = c("Month", "Day"), hive_style = FALSE)

# View files - you can see the partitioning means that files have been written
# to folders based on Month/Day values
tf3_files <- list.files(tf3, recursive = TRUE)

# With no partitioning specified, dataset contains all files but doesn't include
# directory names as field names
open_dataset(tf3)

# Now that partitioning has been specified, your dataset contains columns for Month and Day
open_dataset(tf3, partitioning = c("Month", "Day"))

# If you want to specify the data types for your fields, you can pass in a Schema
open_dataset(tf3, partitioning = schema(Month = int8(), Day = int8()))

# Specifying the type also keeps leading zeros, so "001" stays a string
# instead of becoming the integer 1
products <- data.frame(x = 1:3, product_id = c("001", "002", "010"))
tf4 <- tempfile()
write_dataset(products, tf4, partitioning = "product_id")
open_dataset(tf4, partitioning = schema(product_id = string()))

```

open_delim_dataset	<i>Open a multi-file dataset of CSV or other delimiter-separated format</i>
--------------------	---

Description

A wrapper around [open_dataset](#) which explicitly includes parameters mirroring [read_csv_arrow\(\)](#), [read_delim_arrow\(\)](#), and [read_tsv_arrow\(\)](#) to allow for easy switching between functions for opening single files and functions for opening datasets.

Usage

```
open_delim_dataset(
```

```
sources,  
schema = NULL,  
partitioning = hive_partition(),  
hive_style = NA,  
unify_schemas = NULL,  
factory_options = list(),  
delim = ",",  
quote = "\"",  
escape_double = TRUE,  
escape_backslash = FALSE,  
col_names = TRUE,  
col_types = NULL,  
na = c("", "NA"),  
skip_empty_rows = TRUE,  
skip = 0L,  
convert_options = NULL,  
read_options = NULL,  
timestamp_parsers = NULL,  
quoted_na = TRUE,  
parse_options = NULL  
)
```

```
open_csv_dataset(  
sources,  
schema = NULL,  
partitioning = hive_partition(),  
hive_style = NA,  
unify_schemas = NULL,  
factory_options = list(),  
quote = "\"",  
escape_double = TRUE,  
escape_backslash = FALSE,  
col_names = TRUE,  
col_types = NULL,  
na = c("", "NA"),  
skip_empty_rows = TRUE,  
skip = 0L,  
convert_options = NULL,  
read_options = NULL,  
timestamp_parsers = NULL,  
quoted_na = TRUE,  
parse_options = NULL  
)
```

```
open_tsv_dataset(  
sources,  
schema = NULL,  
partitioning = hive_partition(),
```

```

hive_style = NA,
unify_schemas = NULL,
factory_options = list(),
quote = "\"",
escape_double = TRUE,
escape_backslash = FALSE,
col_names = TRUE,
col_types = NULL,
na = c("", "NA"),
skip_empty_rows = TRUE,
skip = 0L,
convert_options = NULL,
read_options = NULL,
timestamp_parsers = NULL,
quoted_na = TRUE,
parse_options = NULL
)

```

Arguments

sources

One of:

- a string path or URI to a directory containing data files
- a [FileSystem](#) that references a directory containing data files (such as what is returned by [s3_bucket\(\)](#))
- a string path or URI to a single file
- a character vector of paths or URIs to individual data files
- a list of Dataset objects as created by this function
- a list of DatasetFactory objects as created by [dataset_factory\(\)](#).

When sources is a vector of file URIs, they must all use the same protocol and point to files located in the same file system and having the same format.

schema

[Schema](#) for the Dataset. If NULL (the default), the schema will be inferred from the data sources.

partitioning

When sources is a directory path/URI, one of:

- a Schema, in which case the file paths relative to sources will be parsed, and path segments will be matched with the schema fields.
- a character vector that defines the field names corresponding to those path segments (that is, you're providing the names that would correspond to a Schema but the types will be autodetected)
- a Partitioning or PartitioningFactory, such as returned by [hive_partition\(\)](#)
- NULL for no partitioning

The default is to autodetect Hive-style partitions unless `hive_style = FALSE`. See the "Partitioning" section for details. When sources is not a directory path/URI, partitioning is ignored.

hive_style

Logical: should partitioning be interpreted as Hive-style? Default is NA, which means to inspect the file paths for Hive-style partitioning and behave accordingly.

unify_schemas	logical: should all data fragments (files, Datasets) be scanned in order to create a unified schema from them? If FALSE, only the first fragment will be inspected for its schema. Use this fast path when you know and trust that all fragments have an identical schema. The default is FALSE when creating a dataset from a directory path/URI or vector of file paths/URIs (because there may be many files and scanning may be slow) but TRUE when sources is a list of Datasets (because there should be few Datasets in the list and their Schemas are already in memory).
factory_options	list of optional FileSystemFactoryOptions: • partition_base_dir: string path segment prefix to ignore when discovering partition information with DirectoryPartitioning. Not meaningful (ignored with a warning) for HivePartitioning, nor is it valid when providing a vector of file paths. • exclude_invalid_files: logical: should files that are not valid data files be excluded? Default is FALSE because checking all files up front incurs I/O and thus will be slower, especially on remote filesystems. If false and there are invalid files, there will be an error at scan time. This is the only FileSystemFactoryOption that is valid for both when providing a directory path in which to discover files and when providing a vector of file paths. • selector_ignore_prefixes: character vector of file prefixes to ignore when discovering files in a directory. If invalid files can be excluded by a common filename prefix this way, you can avoid the I/O cost of exclude_invalid_files. Not valid when providing a vector of file paths (but if you're providing the file list, you can filter invalid files yourself).
delim	Single character used to separate fields within a record.
quote	Single character used to quote strings.
escape_double	Does the file escape quotes by doubling them? i.e. If this option is TRUE, the value """" represents a single quote, \".
escape_backslash	Does the file use backslashes to escape special characters? This is more general than escape_double as backslashes can be used to escape the delimiter character, the quote character, or to add special characters like \\n.
col_names	If TRUE, the first row of the input will be used as the column names and will not be included in the data frame. If FALSE, column names will be generated by Arrow, starting with "f0", "f1", ..., "fN". Alternatively, you can specify a character vector of column names.
col_types	A compact string representation of the column types, an Arrow Schema , or NULL (the default) to infer types from the data.
na	A character vector of strings to interpret as missing values.
skip_empty_rows	Should blank rows be ignored altogether? If TRUE, blank rows will not be represented at all. If FALSE, they will be filled with missings.
skip	Number of lines to skip before reading data.

convert_options	see CSV conversion options . If given, this overrides any conversion options provided in other arguments (e.g. na, col_types, timestamp_parsers, etc.).
read_options	see CSV reading options
timestamp_parsers	User-defined timestamp parsers, tried in order when inferring column types and when converting columns of type <code>timestamp()</code>. Possible values are: • NULL: the default, which uses the ISO-8601 parser • a character vector of <code>strptime</code> parse strings • a list of <code>TimestampParser</code> objects and/or parse strings Supplying parsers replaces the default ISO-8601 parser rather than adding to it. If none of the parsers match a value during type inference, the column is read as a string without error; to get an error instead, specify the column as a timestamp in <code>col_types</code>. These parsers are not used for date columns.
quoted_na	Should missing values inside quotes be treated as missing values (the default) or strings. (Note that this is different from the the Arrow C++ default for the corresponding convert option, <code>strings_can_be_null</code> .)
parse_options	see CSV parsing options . If given, this overrides any parsing options provided in other arguments (e.g. delim, quote, etc.).

Options currently supported by `read_delim_arrow()` which are not supported here

- file (instead, please specify files in sources)
- col_select (instead, subset columns after dataset creation)
- as_data_frame (instead, convert to data frame after dataset creation)
- parse_options

See Also

[open_dataset\(\)](#)

Examples

```
# Set up directory for examples
tf <- tempfile()
dir.create(tf)

df <- data.frame(x = c("1", "2", "NULL"))
file_path <- file.path(tf, "file1.txt")
write.table(df, file_path, sep = ",", row.names = FALSE)

# Use readr-style params identically in both `read_csv_dataset()` and `open_csv_dataset()`
read_csv_arrow(file_path, na = c("", "NA", "NULL"), col_names = "y", skip = 1)
open_csv_dataset(file_path, na = c("", "NA", "NULL"), col_names = "y", skip = 1)

# Use `col_types` to specify a schema, partial schema, or compact representation
tf2 <- tempfile()
write_csv_dataset(cars, tf2)
```

```
open_csv_dataset(tf2, col_types = schema(speed = int32(), dist = int32()))
open_csv_dataset(tf2, col_types = schema(speed = int32()))
open_csv_dataset(tf2, col_types = "ii", col_names = c("speed", "dist"), skip = 1)
```

OutputStream

OutputStream classes

Description

FileOutputStream is for writing to a file; BufferedOutputStream writes to a buffer; You can create one and pass it to any of the table writers, for example.

Factory

The `$create()` factory methods instantiate the OutputStream object and take the following arguments, depending on the subclass:

- `path` For FileOutputStream, a character file name
- `initial_capacity` For BufferedOutputStream, the size in bytes of the buffer.

Methods

- `$tell()`: return the position in the stream
- `$close()`: close the stream
- `$write(x)`: send `x` to the stream
- `$capacity()`: for BufferedOutputStream
- `$finish()`: for BufferedOutputStream
- `$GetExtentBytesWritten()`: for MockOutputStream, report how many bytes were sent.

ParquetArrowReaderProperties

ParquetArrowReaderProperties class

Description

This class holds settings to control how a Parquet file is read by [ParquetFileReader](#).

Factory

The `ParquetArrowReaderProperties$create()` factory method instantiates the object and takes the following arguments:

- `use_threads` Logical: whether to use multithreading (default TRUE)

Methods

- `$read_dictionary(column_index)`
- `$set_read_dictionary(column_index, read_dict)`
- `$use_threads(use_threads)`

ParquetFileReader	<i>ParquetFileReader class</i>
-------------------	--------------------------------

Description

This class enables you to interact with Parquet files.

Factory

The `ParquetFileReader$create()` factory method instantiates the object and takes the following arguments:

- `file` A character file name, raw vector, or Arrow file connection object (e.g. `RandomAccessFile`).
- `props` Optional [ParquetArrowReaderProperties](#)
- `mmap` Logical: whether to memory-map the file (default TRUE)
- `reader_props` Optional [ParquetReaderProperties](#)
- ... Additional arguments, currently ignored

Methods

- `$ReadTable(column_indices)`: get an `arrow::Table` from the file. The optional `column_indices=` argument is a 0-based integer vector indicating which columns to retain.
- `$ReadRowGroup(i, column_indices)`: get an `arrow::Table` by reading the *i*th row group (0-based). The optional `column_indices=` argument is a 0-based integer vector indicating which columns to retain.
- `$ReadRowGroups(row_groups, column_indices)`: get an `arrow::Table` by reading several row groups (0-based integers). The optional `column_indices=` argument is a 0-based integer vector indicating which columns to retain.
- `$GetSchema()`: get the `arrow::Schema` of the data in the file
- `$ReadColumn(i)`: read the *i*th column (0-based) as a [ChunkedArray](#).

Active bindings

- `$num_rows`: number of rows.
- `$num_columns`: number of columns.
- `$num_row_groups`: number of row groups.

Examples

```
f <- system.file("v0.7.1.parquet", package = "arrow")
pq <- ParquetFileReader$create(f)
pq$GetSchema()
if (codec_is_available("snappy")) {
  # This file has compressed data columns
  tab <- pq$ReadTable()
  tab$schema
}
```

ParquetFileWriter	<i>ParquetFileWriter class</i>
-------------------	--------------------------------

Description

This class enables you to interact with Parquet files.

Factory

The `ParquetFileWriter$create()` factory method instantiates the object and takes the following arguments:

- schema A [Schema](#)
- sink An [arrow::io::OutputStream](#)
- properties An instance of [ParquetWriterProperties](#)
- arrow_properties An instance of [ParquetArrowWriterProperties](#)

Methods

- `WriteTable` Write a [Table](#) to sink
- `WriteBatch` Write a [RecordBatch](#) to sink
- `Close` Close the writer. Note: does not close the sink. [arrow::io::OutputStream](#) has its own `close()` method.

ParquetReaderProperties*ParquetReaderProperties class*

Description

This class holds settings to control how a Parquet file is read by [ParquetFileReader](#).

Factory

The `ParquetReaderProperties#create()` factory method instantiates the object and takes no arguments.

Methods

- `$thrift_string_size_limit()`
- `$set_thrift_string_size_limit()`
- `$thrift_container_size_limit()`
- `$set_thrift_container_size_limit()`

ParquetWriterProperties*ParquetWriterProperties class*

Description

This class holds settings to control how a Parquet file is read by [ParquetFileWriter](#).

Details

The parameters `compression`, `compression_level`, `use_dictionary` and `write_statistics` support various patterns:

- The default NULL leaves the parameter unspecified, and the C++ library uses an appropriate default for each column (defaults listed above)
- A single, unnamed, value (e.g. a single string for compression) applies to all columns
- An unnamed vector, of the same size as the number of columns, to specify a value for each column, in positional order
- A named vector, to specify the value for the named columns, the default value for the setting is used when not supplied

Unlike the high-level [write_parquet](#), `ParquetWriterProperties` arguments use the C++ defaults. Currently this means "uncompressed" rather than "snappy" for the `compression` argument.

Factory

The `ParquetWriterProperties#create()` factory method instantiates the object and takes the following arguments:

- `table`: table to write (required)
- `version`: Parquet version, "1.0" or "2.0". Default "1.0"
- `compression`: Compression type, algorithm "uncompressed"
- `compression_level`: Compression level; meaning depends on compression algorithm
- `use_dictionary`: Specify if we should use dictionary encoding. Default TRUE
- `write_statistics`: Specify if we should write statistics. Default TRUE
- `data_page_size`: Set a target threshold for the approximate encoded size of data pages within a column chunk (in bytes). Default 1 MiB.

See Also

[write_parquet](#)

[Schema](#) for information about schemas and metadata handling.

 Partitioning

Define Partitioning for a Dataset

Description

Pass a `Partitioning` object to a [FileSystemDatasetFactory](#)'s `$create()` method to indicate how the file's paths should be interpreted to define partitioning.

`DirectoryPartitioning` describes how to interpret raw path segments, in order. For example, `schema(year = int16(), month = int8())` would define partitions for file paths like "2019/01/file.parquet", "2019/02/file.parquet", etc. In this scheme NULL values will be skipped. In the previous example: when writing a dataset if the month was NA (or NULL), the files would be placed in "2019/file.parquet". When reading, the rows in "2019/file.parquet" would return an NA for the month column. An error will be raised if an outer directory is NULL and an inner directory is not.

`HivePartitioning` is for Hive-style partitioning, which embeds field names and values in path segments, such as "/year=2019/month=2/data.parquet". Because fields are named in the path segments, order does not matter. This partitioning scheme allows NULL values. They will be replaced by a configurable `null_fallback` which defaults to the string "__HIVE_DEFAULT_PARTITION__" when writing. When reading, the `null_fallback` string will be replaced with NAs as appropriate.

`PartitioningFactory` subclasses instruct the `DatasetFactory` to detect partition features from the file paths.

Factory

Both `DirectoryPartitioning$create()` and `HivePartitioning$create()` methods take a [Schema](#) as a single input argument. The helper function `hive_partition(...)` is shorthand for `HivePartitioning$create(schema)`.

With `DirectoryPartitioningFactory$create()`, you can provide just the names of the path segments (in our example, `c("year", "month")`), and the `DatasetFactory` will infer the data types for those partition variables. `HivePartitioningFactory$create()` takes no arguments: both variable names and their types can be inferred from the file paths. `hive_partition()` with no arguments returns a `HivePartitioningFactory`.

<code>read_delim_arrow</code>	<i>Read a CSV or other delimited file with Arrow</i>
-------------------------------	--

Description

These functions uses the Arrow C++ CSV reader to read into a tibble. Arrow C++ options have been mapped to argument names that follow those of `readr::read_delim()`, and `col_select` was inspired by `vroom::vroom()`.

Usage

```
read_delim_arrow(
  file,
  delim = ",",
  quote = "\"",
  escape_double = TRUE,
  escape_backslash = FALSE,
  schema = NULL,
  col_names = TRUE,
  col_types = NULL,
  col_select = NULL,
  na = c("", "NA"),
  quoted_na = TRUE,
  skip_empty_rows = TRUE,
  skip = 0L,
  parse_options = NULL,
  convert_options = NULL,
  read_options = NULL,
  as_data_frame = TRUE,
  timestamp_parsers = NULL,
  decimal_point = "."
)

read_csv_arrow(
  file,
  quote = "\"",
  escape_double = TRUE,
```

```
    escape_backslash = FALSE,  
    schema = NULL,  
    col_names = TRUE,  
    col_types = NULL,  
    col_select = NULL,  
    na = c("", "NA"),  
    quoted_na = TRUE,  
    skip_empty_rows = TRUE,  
    skip = 0L,  
    parse_options = NULL,  
    convert_options = NULL,  
    read_options = NULL,  
    as_data_frame = TRUE,  
    timestamp_parsers = NULL  
  )
```

```
read_csv2_arrow(  
  file,  
  quote = "\"",  
  escape_double = TRUE,  
  escape_backslash = FALSE,  
  schema = NULL,  
  col_names = TRUE,  
  col_types = NULL,  
  col_select = NULL,  
  na = c("", "NA"),  
  quoted_na = TRUE,  
  skip_empty_rows = TRUE,  
  skip = 0L,  
  parse_options = NULL,  
  convert_options = NULL,  
  read_options = NULL,  
  as_data_frame = TRUE,  
  timestamp_parsers = NULL  
)
```

```
read_tsv_arrow(  
  file,  
  quote = "\"",  
  escape_double = TRUE,  
  escape_backslash = FALSE,  
  schema = NULL,  
  col_names = TRUE,  
  col_types = NULL,  
  col_select = NULL,  
  na = c("", "NA"),  
  quoted_na = TRUE,  
  skip_empty_rows = TRUE,
```

```

    skip = 0L,
    parse_options = NULL,
    convert_options = NULL,
    read_options = NULL,
    as_data_frame = TRUE,
    timestamp_parsers = NULL
  )

```

Arguments

file	A character file name or URI, connection, literal data (either a single string or a raw vector), an Arrow input stream, or a <code>FileSystem</code> with path (<code>SubTreeFileSystem</code>). If a file name, a memory-mapped Arrow InputStream will be opened and closed when finished; compression will be detected from the file extension and handled automatically. If an input stream is provided, it will be left open. To be recognised as literal data, the input must be wrapped with <code>I()</code> .
delim	Single character used to separate fields within a record.
quote	Single character used to quote strings.
escape_double	Does the file escape quotes by doubling them? i.e. If this option is <code>TRUE</code> , the value <code>"\""</code> represents a single quote, <code>\</code> .
escape_backslash	Does the file use backslashes to escape special characters? This is more general than <code>escape_double</code> as backslashes can be used to escape the delimiter character, the quote character, or to add special characters like <code>\\n</code> .
schema	Schema that describes the table. If provided, it will be used to satisfy both <code>col_names</code> and <code>col_types</code> .
col_names	If <code>TRUE</code> , the first row of the input will be used as the column names and will not be included in the data frame. If <code>FALSE</code> , column names will be generated by Arrow, starting with <code>"f0"</code> , <code>"f1"</code> , ..., <code>"fN"</code> . Alternatively, you can specify a character vector of column names.
col_types	A compact string representation of the column types, an Arrow Schema , or <code>NULL</code> (the default) to infer types from the data.
col_select	A character vector of column names to keep, as in the <code>"select"</code> argument to <code>data.table::fread()</code> , or a tidy selection specification of columns, as used in <code>dplyr::select()</code> .
na	A character vector of strings to interpret as missing values.
quoted_na	Should missing values inside quotes be treated as missing values (the default) or strings. (Note that this is different from the the Arrow C++ default for the corresponding convert option, <code>strings_can_be_null</code> .)
skip_empty_rows	Should blank rows be ignored altogether? If <code>TRUE</code> , blank rows will not be represented at all. If <code>FALSE</code> , they will be filled with missings.
skip	Number of lines to skip before reading data.
parse_options	see CSV parsing options . If given, this overrides any parsing options provided in other arguments (e.g. <code>delim</code> , <code>quote</code> , etc.).

convert_options	see CSV conversion options . If given, this overrides any conversion options provided in other arguments (e.g. na, col_types, timestamp_parsers, etc.).
read_options	see CSV reading options
as_data_frame	Should the function return a tibble (default) or an Arrow Table ?
timestamp_parsers	User-defined timestamp parsers, tried in order when inferring column types and when converting columns of type timestamp(). Possible values are: • NULL: the default, which uses the ISO-8601 parser • a character vector of strptime parse strings • a list of TimestampParser objects and/or parse strings Supplying parsers replaces the default ISO-8601 parser rather than adding to it. If none of the parsers match a value during type inference, the column is read as a string without error; to get an error instead, specify the column as a timestamp in col_types. These parsers are not used for date columns.
decimal_point	Character to use for decimal point in floating point numbers.

Details

read_csv_arrow() and read_tsv_arrow() are wrappers around read_delim_arrow() that specify a delimiter. read_csv2_arrow() uses ; for the delimiter and , for the decimal point.

Note that not all readr options are currently implemented here. Please file an issue if you encounter one that arrow should support.

If you need to control Arrow-specific reader parameters that don't have an equivalent in readr::read_csv(), you can either provide them in the parse_options, convert_options, or read_options arguments, or you can use [CsvTableReader](#) directly for lower-level access.

Value

A tibble, or a Table if as_data_frame = FALSE.

Specifying column types and names

By default, the CSV reader will infer the column names and data types from the file, but there are a few ways you can specify them directly.

One way is to provide an Arrow [Schema](#) in the schema argument, which is an ordered map of column name to type. When provided, it satisfies both the col_names and col_types arguments. This is good if you know all of this information up front.

You can also pass a Schema to the col_types argument. If you do this, column names will still be inferred from the file unless you also specify col_names. In either case, the column names in the Schema must match the data's column names, whether they are explicitly provided or inferred. That said, this Schema does not have to reference all columns: those omitted will have their types inferred.

Alternatively, you can declare column types by providing the compact string representation that readr uses to the col_types argument. This means you provide a single string, one character per column, where the characters map to Arrow types analogously to the readr type mapping:

- "c": `utf8()`
- "i": `int32()`
- "n": `float64()`
- "d": `float64()`
- "l": `bool()`
- "f": `dictionary()`
- "D": `date32()`
- "T": `timestamp(unit = "ns")`
- "t": `time32()` (The unit arg is set to the default value "ms")
- "_": `null()`
- "-": `null()`
- "?": infer the type from the data

If you use the compact string representation for `col_types`, you must also specify `col_names`.

Regardless of how types are specified, all columns with a `null()` type will be dropped.

Note that if you are specifying column names, whether by `schema` or `col_names`, and the CSV file has a header row that would otherwise be used to identify column names, you'll need to add `skip = 1` to skip that row.

Examples

```
tf <- tempfile()
on.exit(unlink(tf))
write.csv(mtcars, file = tf)
df <- read_csv_arrow(tf)
dim(df)
# Can select columns
df <- read_csv_arrow(tf, col_select = starts_with("d"))

# Specifying column types and names
write.csv(data.frame(x = c(1, 3), y = c(2, 4)), file = tf, row.names = FALSE)
read_csv_arrow(tf, schema = schema(x = int32(), y = utf8()), skip = 1)
read_csv_arrow(tf, col_types = schema(y = utf8()))
read_csv_arrow(tf, col_types = "ic", col_names = c("x", "y"), skip = 1)

# Note that if a timestamp column contains time zones,
# the string "T" `col_types` specification won't work.
# To parse timestamps with time zones, provide a [Schema] to `col_types`
# and specify the time zone in the type object:
tf <- tempfile()
write.csv(data.frame(x = "1970-01-01T12:00:00+12:00"), file = tf, row.names = FALSE)
read_csv_arrow(
  tf,
  col_types = schema(x = timestamp(unit = "us", timezone = "UTC"))
)

# Parse non-ISO timestamps with `timestamp_parsers`. Supplying parsers
```

```
# replaces the default ISO-8601 parser, so include `TimestampParser$create()`
# to keep it as a fallback:
write.csv(
  data.frame(x = c("16/01/2023 19:47", "2023-01-17 08:00:00")),
  file = tf,
  row.names = FALSE
)
read_csv_arrow(tf, timestamp_parsers = list("%d/%m/%Y %H:%M", TimestampParser$create()))

# Read directly from strings with `I()`
read_csv_arrow(I("x,y\n1,2\n3,4"))
read_delim_arrow(I(c("x y", "1 2", "3 4")), delim = " ")
```

read_feather	<i>Read a Feather file (deprecated)</i>
--------------	---

Description

`read_feather()` is deprecated and will be removed in a future release. Use [read_ipc_file\(\)](#) instead.

`read_feather()` can read both the Feather V1 format (a legacy format which is also being deprecated) and the Feather V2 format (which is the Arrow IPC format). `read_ipc_file()` can also read both formats.

Usage

```
read_feather(file, col_select = NULL, as_data_frame = TRUE, mmap = TRUE)
```

Arguments

file	A character file name or URI, connection, raw vector, an Arrow input stream, or a <code>FileSystem</code> with path (<code>SubTreeFileSystem</code>). If a file name or URI, an Arrow InputStream will be opened and closed when finished. If an input stream is provided, it will be left open.
col_select	A character vector of column names to keep, as in the "select" argument to <code>data.table::fread()</code> , or a tidy selection specification of columns, as used in <code>dplyr::select()</code> .
as_data_frame	Should the function return a tibble (default) or an Arrow Table ?
mmap	Logical: whether to memory-map the file (default TRUE)

Value

A tibble if `as_data_frame` is TRUE (the default), or an Arrow [Table](#) otherwise

See Also

[read_ipc_file\(\)](#)

read_ipc_file	<i>Read an Arrow IPC file</i>
---------------	-------------------------------

Description

The Arrow IPC file format provides binary columnar serialization for data frames. It is designed to make reading and writing data frames efficient, and to make sharing data across data analysis languages easy.

Usage

```
read_ipc_file(file, col_select = NULL, as_data_frame = TRUE, mmap = TRUE)
```

Arguments

file	A character file name or URI, connection, raw vector, an Arrow input stream, or a FileSystem with path (SubTreeFileSystem). If a file name or URI, an Arrow InputStream will be opened and closed when finished. If an input stream is provided, it will be left open.
col_select	A character vector of column names to keep, as in the "select" argument to <code>data.table::fread()</code> , or a tidy selection specification of columns, as used in <code>dplyr::select()</code> .
as_data_frame	Should the function return a tibble (default) or an Arrow Table ?
mmap	Logical: whether to memory-map the file (default TRUE)

Details

This function can also read the legacy Feather V1 format.

Value

A tibble if `as_data_frame` is TRUE (the default), or an Arrow [Table](#) otherwise

See Also

[FeatherReader](#) and [RecordBatchReader](#) for lower-level access to reading Arrow IPC data.

Examples

```
tf <- tempfile(fileext = ".arrow")
on.exit(unlink(tf))
write_ipc_file(mtcars, tf)
df <- read_ipc_file(tf)
dim(df)
# Can select columns
df <- read_ipc_file(tf, col_select = starts_with("d"))
```

read_ipc_stream	<i>Read Arrow IPC stream format</i>
-----------------	-------------------------------------

Description

Apache Arrow defines two formats for **serializing data for interprocess communication (IPC)**: a "stream" format and a "file" format. `read_ipc_stream()` and `read_ipc_file()` read those formats, respectively.

Usage

```
read_ipc_stream(file, as_data_frame = TRUE, ...)
```

Arguments

<code>file</code>	A character file name or URI, connection, raw vector, an Arrow input stream, or a <code>FileSystem</code> with path (<code>SubTreeFileSystem</code>). If a file name or URI, an Arrow InputStream will be opened and closed when finished. If an input stream is provided, it will be left open.
<code>as_data_frame</code>	Should the function return a tibble (default) or an Arrow Table ?
<code>...</code>	deprecated and ignored.

Value

A tibble if `as_data_frame` is `TRUE` (the default), or an Arrow [Table](#) otherwise

Untrusted data

If reading from an untrusted source, you can validate the data by reading with `as_data_frame = FALSE` and calling `$ValidateFull()` on the Table before processing.

See Also

[write_ipc_file\(\)](#) for writing IPC files. [RecordBatchReader](#) for a lower-level interface.

read_json_arrow	<i>Read a JSON file</i>
-----------------	-------------------------

Description

Wrapper around [JsonTableReader](#) to read a newline-delimited JSON (ndjson) file into a data frame or Arrow Table.

Usage

```
read_json_arrow(
  file,
  col_select = NULL,
  as_data_frame = TRUE,
  schema = NULL,
  ...
)
```

Arguments

file	A character file name or URI, connection, literal data (either a single string or a raw vector), an Arrow input stream, or a <code>FileSystem</code> with path (<code>SubTreeFileSystem</code>). If a file name, a memory-mapped Arrow InputStream will be opened and closed when finished; compression will be detected from the file extension and handled automatically. If an input stream is provided, it will be left open. To be recognised as literal data, the input must be wrapped with <code>I()</code> .
col_select	A character vector of column names to keep, as in the "select" argument to <code>data.table::fread()</code> , or a tidy selection specification of columns, as used in <code>dplyr::select()</code> .
as_data_frame	Should the function return a tibble (default) or an Arrow Table ?
schema	Schema that describes the table.
...	Additional options passed to <code>JsonTableReader\$create()</code>

Details

If passed a path, will detect and handle compression from the file extension (e.g. `.json.gz`).

If schema is not provided, Arrow data types are inferred from the data:

- JSON null values convert to the [null\(\)](#) type, but can fall back to any other type.
- JSON booleans convert to [boolean\(\)](#).
- JSON numbers convert to [int64\(\)](#), falling back to [float64\(\)](#) if a non-integer is encountered.
- JSON strings of the kind "YYYY-MM-DD" and "YYYY-MM-DD hh:mm:ss" convert to [timestamp\(unit = "s"\)](#), falling back to [utf8\(\)](#) if a conversion error occurs.
- JSON arrays convert to a [list_of\(\)](#) type, and inference proceeds recursively on the JSON arrays' values.
- Nested JSON objects convert to a [struct\(\)](#) type, and inference proceeds recursively on the JSON objects' values.

When `as_data_frame = TRUE`, Arrow types are further converted to R types.

Value

A tibble, or a Table if `as_data_frame = FALSE`.

Examples

```
tf <- tempfile()
on.exit(unlink(tf))
writeLines('
  { "hello": 3.5, "world": false, "yo": "thing" }
  { "hello": 3.25, "world": null }
  { "hello": 0.0, "world": true, "yo": null }
', tf, useBytes = TRUE)

read_json_arrow(tf)

# Read directly from strings with `I()`
read_json_arrow(I(c('{ "x": 1, "y": 2}', '{ "x": 3, "y": 4}')))
```

read_message	<i>Read a Message from a stream</i>
--------------	-------------------------------------

Description

Read a Message from a stream

Usage

```
read_message(stream)
```

Arguments

stream an InputStream

read_parquet	<i>Read a Parquet file</i>
--------------	----------------------------

Description

'**Parquet**' is a columnar storage file format. This function enables you to read Parquet files into R.

Usage

```
read_parquet(
  file,
  col_select = NULL,
  as_data_frame = TRUE,
  props = ParquetArrowReaderProperties$create(),
  mmap = TRUE,
  ...
)
```

Arguments

file	A character file name or URI, connection, raw vector, an Arrow input stream, or a FileSystem with path (SubTreeFileSystem). If a file name or URI, an Arrow InputStream will be opened and closed when finished. If an input stream is provided, it will be left open.
col_select	A character vector of column names to keep, as in the "select" argument to <code>data.table::fread()</code> , or a tidy selection specification of columns, as used in <code>dplyr::select()</code> .
as_data_frame	Should the function return a tibble (default) or an Arrow Table ?
props	ParquetArrowReaderProperties
mmap	Use TRUE to use memory mapping where possible
...	Additional arguments passed to <code>ParquetFileReader\$create()</code>

Value

A tibble if `as_data_frame` is TRUE (the default), or an Arrow [Table](#) otherwise.

Examples

```
tf <- tempfile()
on.exit(unlink(tf))
write_parquet(mtcars, tf)
df <- read_parquet(tf, col_select = starts_with("d"))
head(df)
```

read_schema	<i>Read a Schema from a stream</i>
-------------	------------------------------------

Description

Read a Schema from a stream

Usage

```
read_schema(stream, ...)
```

Arguments

stream	a Message, InputStream, or Buffer
...	currently ignored

Value

A [Schema](#)

record_batch	<i>Create a RecordBatch</i>
--------------	-----------------------------

Description

Create a RecordBatch

Usage

```
record_batch(..., schema = NULL)
```

Arguments

...	A data.frame or a named set of Arrays or vectors. If given a mixture of data.frames and vectors, the inputs will be autospliced together (see examples). Alternatively, you can provide a single Arrow IPC InputStream, Message, Buffer, or R raw object containing a Buffer.
schema	a Schema , or NULL (the default) to infer the schema from the data in ... When providing an Arrow IPC buffer, schema is required.

Examples

```
batch <- record_batch(name = rownames(mtcars), mtcars)
dim(batch)
dim(head(batch))
names(batch)
batch$mpg
batch[["cyl"]]
as.data.frame(batch[4:8, c("gear", "hp", "wt")])
```

RecordBatch	<i>RecordBatch class</i>
-------------	--------------------------

Description

A record batch is a collection of equal-length arrays matching a particular [Schema](#). It is a table-like data structure that is semantically a sequence of [fields](#), each a contiguous Arrow [Array](#).

S3 Methods and Usage

Record batches are data-frame-like, and many methods you expect to work on a data.frame are implemented for RecordBatch. This includes [, [[, \$, names, dim, nrow, ncol, head, and tail. You can also pull the data from an Arrow record batch into R with as.data.frame(). See the examples.

A caveat about the \$ method: because RecordBatch is an R6 object, \$ is also used to access the object's methods (see below). Methods take precedence over the table's columns. So, batch\$Slice would return the "Slice" method function even if there were a column in the table called "Slice".

R6 Methods

In addition to the more R-friendly S3 methods, a RecordBatch object has the following R6 methods that map onto the underlying C++ methods:

- `$Equals(other)`: Returns TRUE if the other record batch is equal
- `$column(i)`: Extract an Array by integer position from the batch
- `$column_name(i)`: Get a column's name by integer position
- `$names()`: Get all column names (called by `names(batch)`)
- `$nbytes()`: Total number of bytes consumed by the elements of the record batch
- `$RenameColumns(value)`: Set all column names (called by `names(batch) <- value`)
- `$GetColumnByName(name)`: Extract an Array by string name
- `$RemoveColumn(i)`: Drops a column from the batch by integer position
- `$SelectColumns(indices)`: Return a new record batch with a selection of columns, expressed as 0-based integers.
- `$Slice(offset, length = NULL)`: Create a zero-copy view starting at the indicated integer offset and going for the given length, or to the end of the table if NULL, the default.
- `$Take(i)`: return an RecordBatch with rows at positions given by integers (R vector or Array Array) `i`.
- `$Filter(i, keep_na = TRUE)`: return an RecordBatch with rows at positions where logical vector (or Arrow boolean Array) `i` is TRUE.
- `$SortIndices(names, descending = FALSE)`: return an Array of integer row positions that can be used to rearrange the RecordBatch in ascending or descending order by the first named column, breaking ties with further named columns. `descending` can be a logical vector of length one or of the same length as `names`.
- `$serialize()`: Returns a raw vector suitable for interprocess communication
- `$cast(target_schema, safe = TRUE, options = cast_options(safe))`: Alter the schema of the record batch.

There are also some active bindings

- `$num_columns`
- `$num_rows`
- `$schema`
- `$metadata`: Returns the key-value metadata of the Schema as a named list. Modify or replace by assigning in (`batch$metadata <- new_metadata`). All list elements are coerced to string. See `schema()` for more information.
- `$columns`: Returns a list of Arrays

RecordBatchReader

RecordBatchReader classes

Description

Apache Arrow defines two formats for **serializing data for interprocess communication (IPC)**: a "stream" format and a "file" format. RecordBatchStreamReader and RecordBatchFileReader are interfaces for accessing record batches from input sources in those formats, respectively.

For guidance on how to use these classes, see the examples section.

Factory

The `RecordBatchFileReader$create()` and `RecordBatchStreamReader$create()` factory methods instantiate the object and take a single argument, named according to the class:

- file A character file name, raw vector, or Arrow file connection object (e.g. [RandomAccessFile](#)).
- stream A raw vector, [Buffer](#), or [InputStream](#).

Methods

- `$read_next_batch()`: Returns a `RecordBatch`, iterating through the Reader. If there are no further batches in the Reader, it returns `NULL`.
- `$schema`: Returns a [Schema](#) (active binding)
- `$batches()`: Returns a list of `RecordBatches`
- `$read_table()`: Collects the reader's `RecordBatches` into a [Table](#)
- `$get_batch(i)`: For `RecordBatchFileReader`, return a particular batch by an integer index.
- `$num_record_batches()`: For `RecordBatchFileReader`, see how many batches are in the file.

See Also

[read_ipc_stream\(\)](#) and [read_ipc_file\(\)](#) provide a much simpler interface for reading data from these formats and are sufficient for many use cases.

Examples

```
tf <- tempfile()
on.exit(unlink(tf))

batch <- record_batch(chickwts)

# This opens a connection to the file in Arrow
file_obj <- FileOutputStream$create(tf)
# Pass that to a RecordBatchWriter to write data conforming to a schema
writer <- RecordBatchFileWriter$create(file_obj, batch$schema)
writer$write(batch)
# You may write additional batches to the stream, provided that they have
# the same schema.
# Call "close" on the writer to indicate end-of-file/stream
writer$close()
# Then, close the connection--closing the IPC message does not close the file
file_obj$close()

# Now, we have a file we can read from. Same pattern: open file connection,
# then pass it to a RecordBatchReader
read_file_obj <- ReadableFile$create(tf)
reader <- RecordBatchFileReader$create(read_file_obj)
# RecordBatchFileReader knows how many batches it has (StreamReader does not)
reader$num_record_batches
# We could consume the Reader by calling $read_next_batch() until all are,
```

```
# consumed, or we can call $read_table() to pull them all into a Table
tab <- reader$read_table()
# Call as.data.frame to turn that Table into an R data.frame
df <- as.data.frame(tab)
# This should be the same data we sent
all.equal(df, chickwts, check.attributes = FALSE)
# Unlike the Writers, we don't have to close RecordBatchReaders,
# but we do still need to close the file connection
read_file_obj$close()
```

RecordBatchWriter	<i>RecordBatchWriter classes</i>
-------------------	----------------------------------

Description

Apache Arrow defines two formats for **serializing data for interprocess communication (IPC)**: a "stream" format and a "file" format. RecordBatchStreamWriter and RecordBatchFileWriter are interfaces for writing record batches to those formats, respectively.

For guidance on how to use these classes, see the examples section.

Factory

The RecordBatchFileWriter\$create() and RecordBatchStreamWriter\$create() factory methods instantiate the object and take the following arguments:

- sink An OutputStream
- schema A [Schema](#) for the data to be written
- use_legacy_format logical: write data formatted so that Arrow libraries versions 0.14 and lower can read it. Default is FALSE. You can also enable this by setting the environment variable ARROW_PRE_0_15_IPC_FORMAT=1.
- metadata_version: A string like "V5" or the equivalent integer indicating the Arrow IPC MetadataVersion. Default (NULL) will use the latest version, unless the environment variable ARROW_PRE_1_0_METADATA_VERSION=1, in which case it will be V4.

Methods

- \$write(x): Write a [RecordBatch](#), [Table](#), or data.frame, dispatching to the methods below appropriately
- \$write_batch(batch): Write a RecordBatch to stream
- \$write_table(table): Write a Table to stream
- \$close(): close stream. Note that this indicates end-of-file or end-of-stream—it does not close the connection to the sink. That needs to be closed separately.

See Also

[write_ipc_stream\(\)](#) and [write_ipc_file\(\)](#) provide a much simpler interface for writing data to these formats and are sufficient for many use cases. [write_to_raw\(\)](#) is a version that serializes data to a buffer.

Examples

```
tf <- tempfile()
on.exit(unlink(tf))

batch <- record_batch(chickwts)

# This opens a connection to the file in Arrow
file_obj <- FileOutputStream$create(tf)
# Pass that to a RecordBatchWriter to write data conforming to a schema
writer <- RecordBatchFileWriter$create(file_obj, batch$schema)
writer$write(batch)
# You may write additional batches to the stream, provided that they have
# the same schema.
# Call "close" on the writer to indicate end-of-file/stream
writer$close()
# Then, close the connection--closing the IPC message does not close the file
file_obj$close()

# Now, we have a file we can read from. Same pattern: open file connection,
# then pass it to a RecordBatchReader
read_file_obj <- ReadableFile$create(tf)
reader <- RecordBatchFileReader$create(read_file_obj)
# RecordBatchFileReader knows how many batches it has (StreamReader does not)
reader$num_record_batches
# We could consume the Reader by calling $read_next_batch() until all are,
# consumed, or we can call $read_table() to pull them all into a Table
tab <- reader$read_table()
# Call as.data.frame to turn that Table into an R data.frame
df <- as.data.frame(tab)
# This should be the same data we sent
all.equal(df, chickwts, check.attributes = FALSE)
# Unlike the Writers, we don't have to close RecordBatchReaders,
# but we do still need to close the file connection
read_file_obj$close()
```

register_scalar_function

Register user-defined functions

Description

These functions support calling R code from query engine execution (i.e., a `dplyr::mutate()` or `dplyr::filter()` on a `Table` or `Dataset`). Use `register_scalar_function()` attach Arrow input and output types to an R function and make it available for use in the dplyr interface and/or `call_function()`. Scalar functions are currently the only type of user-defined function supported. In Arrow, scalar functions must be stateless and return output with the same shape (i.e., the same number of rows) as the input.

Usage

```
register_scalar_function(name, fun, in_type, out_type, auto_convert = FALSE)
```

Arguments

name	The function name to be used in the dplyr bindings
fun	An R function or rlang-style lambda expression. The function will be called with a first argument context which is a <code>list()</code> with elements <code>batch_size</code> (the expected length of the output) and <code>output_type</code> (the required DataType of the output) that may be used to ensure that the output has the correct type and length. Subsequent arguments are passed by position as specified by <code>in_types</code> . If <code>auto_convert</code> is <code>TRUE</code> , subsequent arguments are converted to R vectors before being passed to <code>fun</code> and the output is automatically constructed with the expected output type via <code>as_arrow_array()</code> .
in_type	A DataType of the input type or a <code>schema()</code> for functions with more than one argument. This signature will be used to determine if this function is appropriate for a given set of arguments. If this function is appropriate for more than one signature, pass a <code>list()</code> of the above. Arguments are passed to <code>fun</code> by position, so if the schema (or field) is named, the names must match the argument names of <code>fun</code> (after context). Fields that would be passed to <code>...</code> in <code>fun</code> can be named anything.
out_type	A DataType of the output type or a function accepting a single argument (types), which is a <code>list()</code> of DataTypes . If a function it must return a DataType .
auto_convert	Use <code>TRUE</code> to convert inputs before passing to <code>fun</code> and construct an <code>Array</code> of the correct type from the output. Use this option to write functions of R objects as opposed to functions of Arrow R6 objects.

Value

NULL, invisibly

Examples

```
library(dplyr, warn.conflicts = FALSE)

some_model <- lm(mpg ~ disp + cyl, data = mtcars)
register_scalar_function(
  "mtcars_predict_mpg",
  function(context, disp, cyl) {
    predict(some_model, newdata = data.frame(disp, cyl))
  },
  in_type = schema(disp = float64(), cyl = float64()),
  out_type = float64(),
  auto_convert = TRUE
)

as_arrow_table(mtcars) |>
  transmute(mpg, mpg_predicted = mtcars_predict_mpg(disp, cyl)) |>
  collect() |>
```

```
head()
```

s3_bucket

Connect to an AWS S3 bucket

Description

s3_bucket() is a convenience function to create an S3FileSystem object that automatically detects the bucket's AWS region and holding onto the its relative path.

Usage

```
s3_bucket(bucket, ...)
```

Arguments

bucket	string S3 bucket name or path
...	Additional connection options, passed to S3FileSystem\$create()

Details

By default, [s3_bucket](#) and other [S3FileSystem](#) functions only produce output for fatal errors or when printing their return values. When troubleshooting problems, it may be useful to increase the log level. See the Notes section in [S3FileSystem](#) for more information or see Examples below.

Value

A SubTreeFileSystem containing an S3FileSystem and the bucket's relative path. Note that this function's success does not guarantee that you are authorized to access the bucket's contents.

Examples

```
bucket <- s3_bucket("arrow-datasets")

# Turn on debug logging. The following line of code should be run in a fresh
# R session prior to any calls to `s3_bucket()` (or other S3 functions)
Sys.setenv("ARROW_S3_LOG_LEVEL" = "DEBUG")
bucket <- s3_bucket("arrow-datasets")
```

scalar	<i>Create an Arrow Scalar</i>
--------	-------------------------------

Description

Create an Arrow Scalar

Usage

```
scalar(x, type = NULL)
```

Arguments

x	An R vector, list, or data.frame
type	An optional data type for x. If omitted, the type will be inferred from the data.

Examples

```
scalar(pi)
scalar(404)
# If you pass a vector into scalar(), you get a list containing your items
scalar(c(1, 2, 3))

scalar(9) == scalar(10)
```

Scalar	<i>Arrow scalars</i>
--------	----------------------

Description

A Scalar holds a single value of an Arrow type.

Factory

The `Scalar$create()` factory method instantiates a `Scalar` and takes the following arguments:

- x: an R vector, list, or data.frame
- type: an optional [data type](#) for x. If omitted, the type will be inferred from the data.

Usage

```
a <- Scalar$create(x)
length(a)

print(a)
a == a
```

Methods

- `$ToString()`: convert to a string
- `$as_vector()`: convert to an R vector
- `$as_array()`: convert to an Arrow Array
- `$Equals(other)`: is this Scalar equal to other
- `$ApproxEquals(other)`: is this Scalar approximately equal to other
- `$is_valid`: is this Scalar valid
- `$null_count`: number of invalid values - 1 or 0
- `$type`: Scalar type
- `$cast(target_type, safe = TRUE, options = cast_options(safe))`: cast value to a different type

Examples

```
Scalar$create(pi)
Scalar$create(404)
# If you pass a vector into Scalar$create, you get a list containing your items
Scalar$create(c(1, 2, 3))

# Comparisons
my_scalar <- Scalar$create(99)
my_scalar$ApproxEquals(Scalar$create(99.00001)) # FALSE
my_scalar$ApproxEquals(Scalar$create(99.000009)) # TRUE
my_scalar$Equals(Scalar$create(99.000009)) # FALSE
my_scalar$Equals(Scalar$create(99L)) # FALSE (types don't match)

my_scalar$ToString()
```

Scanner

Scan the contents of a dataset

Description

A Scanner iterates over a [Dataset](#)'s fragments and returns data according to given row filtering and column projection. A `ScannerBuilder` can help create one.

Factory

`Scanner$create()` wraps the `ScannerBuilder` interface to make a `Scanner`. It takes the following arguments:

- `dataset`: A `Dataset` or `arrow_dplyr_query` object, as returned by the `dplyr` methods on `Dataset`.
- `projection`: A character vector of column names to select columns or a named list of expressions
- `filter`: A Expression to filter the scanned rows by, or `TRUE` (default) to keep all rows.
- `use_threads`: logical: should scanning use multithreading? Default `TRUE`
- `...`: Additional arguments, currently ignored

Methods

ScannerBuilder has the following methods:

- `$Project(cols)`: Indicate that the scan should only return columns given by cols, a character vector of column names or a named list of [Expression](#).
- `$Filter(expr)`: Filter rows by an [Expression](#).
- `$UseThreads(threads)`: logical: should the scan use multithreading? The method's default input is TRUE, but you must call the method to enable multithreading because the scanner default is FALSE.
- `$BatchSize(batch_size)`: integer: Maximum row count of scanned record batches, default is 32K. If scanned record batches are overflowing memory then this method can be called to reduce their size.
- `$schema`: Active binding, returns the [Schema](#) of the Dataset
- `$Finish()`: Returns a Scanner

Scanner currently has a single method, `$ToTable()`, which evaluates the query and returns an Arrow [Table](#).

Examples

```
# Set up directory for examples
tf <- tempfile()
dir.create(tf)
on.exit(unlink(tf))

write_dataset(mtcars, tf, partitioning="cyl")

ds <- open_dataset(tf)

scan_builder <- ds$NewScan()
scan_builder$Filter(Expression$field_ref("hp") > 100)
scan_builder$Project(list(hp_times_ten = 10 * Expression$field_ref("hp")))

# Once configured, call $Finish()
scanner <- scan_builder$Finish()

# Can get results as a table
as.data.frame(scanner$ToTable())

# Or as a RecordBatchReader
scanner$ToRecordBatchReader()
```

schema

Create a schema or extract one from an object.

Description

Create a schema or extract one from an object.

Usage

```
schema(...)
```

Arguments

```
...           fields, field name/data type pairs (or a list of), or object from which to extract a
              schema
```

See Also

[Schema](#) for detailed documentation of the Schema R6 object

Examples

```
# Create schema using pairs of field names and data types
schema(a = int32(), b = float64())

# Create a schema using a list of pairs of field names and data types
schema(list(a = int8(), b = string()))

# Create schema using fields
schema(
  field("b", double()),
  field("c", bool(), nullable = FALSE),
  field("d", string())
)

# Extract schemas from objects
df <- data.frame(col1 = 2:4, col2 = c(0.1, 0.3, 0.5))
tab1 <- arrow_table(df)
schema(tab1)
tab2 <- arrow_table(df, schema = schema(col1 = int8(), col2 = float32()))
schema(tab2)
```

Schema

Schema class

Description

A Schema is an Arrow object containing [Fields](#), which map names to Arrow [data types](#). Create a Schema when you want to convert an R `data.frame` to Arrow but don't want to rely on the default mapping of R types to Arrow types, such as when you want to choose a specific numeric precision, or when creating a [Dataset](#) and you want to ensure a specific schema rather than inferring it from the various files.

Many Arrow objects, including [Table](#) and [Dataset](#), have a `$schema` method (active binding) that lets you access their schema.

Methods

- `$ToString()`: convert to a string
- `$field(i)`: returns the field at index `i` (0-based)
- `$GetFieldByName(x)`: returns the field with name `x`
- `$WithMetadata(metadata)`: returns a new Schema with the key-value metadata set. Note that all list elements in metadata will be coerced to character.
- `$code(namespace)`: returns the R code needed to generate this schema. Use `namespace=TRUE` to call with `arrow::`.

Active bindings

- `$names`: returns the field names (called in `names(Schema)`)
- `$num_fields`: returns the number of fields (called in `length(Schema)`)
- `$fields`: returns the list of Fields in the Schema, suitable for iterating over
- `$HasMetadata`: logical: does this Schema have extra metadata?
- `$metadata`: returns the key-value metadata as a named list. Modify or replace by assigning in (`sch$metadata <- new_metadata`). All list elements are coerced to string.

R Metadata

When converting a `data.frame` to an Arrow Table or RecordBatch, attributes from the `data.frame` are saved alongside tables so that the object can be reconstructed faithfully in R (e.g. with `as.data.frame()`). This metadata can be both at the top-level of the `data.frame` (e.g. `attributes(df)`) or at the column (e.g. `attributes(df$col_a)`) or for list columns only: element level (e.g. `attributes(df[1, "col_a"])`). For example, this allows for storing haven columns in a table and being able to faithfully re-create them when pulled back into R. This metadata is separate from the schema (column names and types) which is compatible with other Arrow clients. The R metadata is only read by R and is ignored by other clients (e.g. Pandas has its own custom metadata). This metadata is stored in `$metadata[["r"]]`.

Since Schema metadata keys and values must be strings, this metadata is saved by serializing R's attribute list structure to a string. If the serialized metadata exceeds 100Kb in size, by default it is compressed starting in version 3.0.0. To disable this compression (e.g. for tables that are compatible with Arrow versions before 3.0.0 and include large amounts of metadata), set the option `arrow.compress_metadata` to `FALSE`. Files with compressed metadata are readable by older versions of arrow, but the metadata is dropped.

show_exec_plan

Show the details of an Arrow Execution Plan

Description

This is a function which gives more details about the logical query plan that will be executed when evaluating an `arrow_dplyr_query` object. It calls the C++ `ExecPlan` object's `print` method. Functionally, it is similar to `dplyr::explain()`. This function is used as the `dplyr::explain()` and `dplyr::show_query()` methods.

Usage

```
show_exec_plan(x)
```

Arguments

`x` an `arrow_dplyr_query` to print the `ExecPlan` for.

Value

`x`, invisibly.

Examples

```
library(dplyr)
mtcars |>
  arrow_table() |>
  filter(mpg > 20) |>
  mutate(x = gear / carb) |>
  show_exec_plan()
```

Table	<i>Table class</i>
-------	--------------------

Description

A `Table` is a sequence of [chunked arrays](#). They have a similar interface to [record batches](#), but they can be composed from multiple record batches or chunked arrays.

S3 Methods and Usage

Tables are data-frame-like, and many methods you expect to work on a `data.frame` are implemented for `Table`. This includes `[], [[, $, names, dim, nrow, ncol, head, and tail`. You can also pull the data from an Arrow table into R with `as.data.frame()`. See the examples.

A caveat about the `$` method: because `Table` is an R6 object, `$` is also used to access the object's methods (see below). Methods take precedence over the table's columns. So, `tab$Slice` would return the "Slice" method function even if there were a column in the table called "Slice".

R6 Methods

In addition to the more R-friendly S3 methods, a `Table` object has the following R6 methods that map onto the underlying C++ methods:

- `$column(i)`: Extract a `ChunkedArray` by integer position from the table
- `$ColumnNames()`: Get all column names (called by `names(tab)`)
- `$nbytes()`: Total number of bytes consumed by the elements of the table
- `$RenameColumns(value)`: Set all column names (called by `names(tab) <- value`)
- `$GetColumnByName(name)`: Extract a `ChunkedArray` by string name

- `$field(i)`: Extract a Field from the table schema by integer position
- `$SelectColumns(indices)`: Return new Table with specified columns, expressed as 0-based integers.
- `$Slice(offset, length = NULL)`: Create a zero-copy view starting at the indicated integer offset and going for the given length, or to the end of the table if NULL, the default.
- `$Take(i)`: return an Table with rows at positions given by integers `i`. If `i` is an Arrow Array or ChunkedArray, it will be coerced to an R vector before taking.
- `$Filter(i, keep_na = TRUE)`: return an Table with rows at positions where logical vector or Arrow boolean-type (Chunked)Array `i` is TRUE.
- `$SortIndices(names, descending = FALSE)`: return an Array of integer row positions that can be used to rearrange the Table in ascending or descending order by the first named column, breaking ties with further named columns. `descending` can be a logical vector of length one or of the same length as `names`.
- `$serialize(output_stream, ...)`: Write the table to the given [OutputStream](#)
- `$cast(target_schema, safe = TRUE, options = cast_options(safe))`: Alter the schema of the record batch.

There are also some active bindings:

- `$num_columns`
- `$num_rows`
- `$schema`
- `$metadata`: Returns the key-value metadata of the Schema as a named list. Modify or replace by assigning in `(tab$metadata <- new_metadata)`. All list elements are coerced to string. See `schema()` for more information.
- `$columns`: Returns a list of ChunkedArrays

to_arrow

Create an Arrow object from a DuckDB connection

Description

This can be used in pipelines that pass data back and forth between Arrow and DuckDB.

Usage

```
to_arrow(.data)
```

Arguments

`.data` the object to be converted

Details

Note that you can only call `collect()` or `compute()` on the result of this function once. To work around this limitation, you should either only call `collect()` as the final step in a pipeline or call `as_arrow_table()` on the result to materialize the entire Table in-memory.

Value

A RecordBatchReader.

Examples

```
library(dplyr)

ds <- InMemoryDataset$create(mtcars)

ds |>
  filter(mpg < 30) |>
  to_duckdb() |>
  group_by(cyl) |>
  summarize(mean_mpg = mean(mpg, na.rm = TRUE)) |>
  to_arrow() |>
  collect()
```

to_duckdb

Create a (virtual) DuckDB table from an Arrow object

Description

This will do the necessary configuration to create a (virtual) table in DuckDB that is backed by the Arrow object given. No data is copied or modified until `collect()` or `compute()` are called or a query is run against the table.

Usage

```
to_duckdb(
  .data,
  con = arrow_duck_connection(),
  table_name = unique_arrow_tablename(),
  auto_disconnect = TRUE
)
```

Arguments

<code>.data</code>	the Arrow object (e.g. Dataset, Table) to use for the DuckDB table
<code>con</code>	a DuckDB connection to use (default will create one and store it in <code>options("arrow_duck_con")</code>)
<code>table_name</code>	a name to use in DuckDB for this object. The default is a unique string "arrow_" followed by numbers.
<code>auto_disconnect</code>	should the table be automatically cleaned up when the resulting object is removed (and garbage collected)? Default: TRUE

Details

The result is a dbplyr-compatible object that can be used in d(b)plyr pipelines.

If `auto_disconnect = TRUE`, the DuckDB table that is created will be configured to be unregistered when the `tbl` object is garbage collected. This is helpful if you don't want to have extra table objects in DuckDB after you've finished using them.

Value

A `tbl` of the new table in DuckDB

Examples

```
library(dplyr)

ds <- InMemoryDataset$create(mtcars)

ds |>
  filter(mpg < 30) |>
  group_by(cyl) |>
  to_duckdb() |>
  slice_min(displ)
```

unify_schemas

Combine and harmonize schemas

Description

Combine and harmonize schemas

Usage

```
unify_schemas(..., schemas = list(...))
```

Arguments

`...` [Schemas](#) to unify
`schemas` Alternatively, a list of schemas

Value

A Schema with the union of fields contained in the inputs, or NULL if any of schemas is NULL

Examples

```
a <- schema(b = double(), c = bool())
z <- schema(b = double(), k = utf8())
unify_schemas(a, z)
```

value_counts	table for Arrow objects
--------------	-------------------------

Description

This function tabulates the values in the array and returns a table of counts.

Usage

```
value_counts(x)
```

Arguments

x Array or ChunkedArray

Value

A StructArray containing "values" (same type as x) and "counts" Int64.

Examples

```
cyl_vals <- Array$create(mtcars$cyl)
counts <- value_counts(cyl_vals)
```

vctrs_extension_array *Extension type for generic typed vectors*

Description

Most common R vector types are converted automatically to a suitable Arrow [data type](#) without the need for an extension type. For vector types whose conversion is not suitably handled by default, you can create a [vctrs_extension_array\(\)](#), which passes [vctrs::vec_data\(\)](#) to [Array\\$create\(\)](#) and calls [vctrs::vec_restore\(\)](#) when the [Array](#) is converted back into an R vector.

Usage

```
vctrs_extension_array(x, ptype = vctrs::vec_ptype(x), storage_type = NULL)
```

```
vctrs_extension_type(x, storage_type = infer_type(vctrs::vec_data(x)))
```

Arguments

x	A vctr (i.e., vctrs::vec_is() returns TRUE).
ptype	A vctrs::vec_ptype() , which is usually a zero-length version of the object with the appropriate attributes set. This value will be serialized using serialize() , so it should not refer to any R object that can't be saved/reloaded.
storage_type	The data type of the underlying storage array.

Value

- `vctrs_extension_array()` returns an [ExtensionArray](#) instance with a `vctrs_extension_type()`.
- `vctrs_extension_type()` returns an [ExtensionType](#) instance for the extension name "arrow.r.vctrs".

Examples

```
(array <- vctrs_extension_array(as.POSIXlt("2022-01-02 03:45", tz = "UTC")))  
array$type  
as.vector(array)  
  
temp_ipc <- tempfile()  
write_ipc_file(arrow_table(col = array), temp_ipc)  
read_ipc_file(temp_ipc)  
unlink(temp_ipc)
```

write_csv_arrow	Write CSV file to disk
-----------------	------------------------

Description

Write CSV file to disk

Usage

```
write_csv_arrow(  
  x,  
  sink,  
  file = NULL,  
  include_header = TRUE,  
  col_names = NULL,  
  batch_size = 1024L,  
  na = "",  
  write_options = NULL,  
  ...  
)
```

Arguments

x	data.frame, RecordBatch , or Table
sink	A string file path, connection, URI, or OutputStream , or path in a file system (SubTreeFileSystem)
file	file name. Specify this or sink, not both.
include_header	Whether to write an initial header line with column names
col_names	identical to include_header. Specify this or include_headers, not both.
batch_size	Maximum number of rows processed at a time. Default is 1024.

na value to write for NA values. Must not contain quote marks. Default is "".

write_options see [CSV write options](#)

... additional parameters

Value

The input x, invisibly. Note that if sink is an [OutputStream](#), the stream will be left open.

Examples

```
tf <- tempfile()
on.exit(unlink(tf))
write_csv_arrow(mtcars, tf)
```

write_dataset	<i>Write a dataset</i>
---------------	------------------------

Description

This function allows you to write a dataset. By writing to more efficient binary storage formats, and by specifying relevant partitioning, you can make it much faster to read and query.

Usage

```
write_dataset(
  dataset,
  path,
  format = c("parquet", "feather", "arrow", "ipc", "csv", "tsv", "txt", "text"),
  partitioning = dplyr::group_vars(dataset),
  basename_template = paste0("part-{i}.", as.character(format)),
  hive_style = TRUE,
  existing_data_behavior = c("overwrite", "error", "delete_matching"),
  max_partitions = 1024L,
  max_open_files = 900L,
  max_rows_per_file = 0L,
  min_rows_per_group = 0L,
  max_rows_per_group = bitwShiftL(1, 20),
  create_directory = TRUE,
  preserve_order = FALSE,
  ...
)
```

Arguments

dataset [Dataset](#), [RecordBatch](#), [Table](#), [arrow_dplyr_query](#), or [data.frame](#). If an [arrow_dplyr_query](#), the query will be evaluated and the result will be written. This means that you can [select\(\)](#), [filter\(\)](#), [mutate\(\)](#), etc. to transform the data before it is written if you need to.

path	string path, URI, or SubTreeFileSystem referencing a directory to write to (directory will be created if it does not exist)
format	a string identifier of the file format. Default is to use "parquet" (see FileFormat)
partitioning	Partitioning or a character vector of columns to use as partition keys (to be written as path segments). Default is to use the current group_by() columns.
basename_template	string template for the names of files to be written. Must contain "{i}", which will be replaced with an autoincremented integer to generate basenames of datafiles. For example, "part-{i}.arrow" will yield "part-0.arrow", If not specified, it defaults to "part-{i}.<default extension>".
hive_style	logical: write partition segments as Hive-style (key1=value1/key2=value2/file.ext) or as just bare values. Default is TRUE.
existing_data_behavior	The behavior to use when there is already data in the destination directory. Must be one of "overwrite", "error", or "delete_matching". • "overwrite" (the default) then any new files created will overwrite existing files • "error" then the operation will fail if the destination directory is not empty • "delete_matching" then the writer will delete any existing partitions if data is going to be written to those partitions and will leave alone partitions which data is not written to.
max_partitions	maximum number of partitions any batch may be written into. Default is 1024L.
max_open_files	maximum number of files that can be left opened during a write operation. If greater than 0 then this will limit the maximum number of files that can be left open. If an attempt is made to open too many files then the least recently used file will be closed. If this setting is set too low you may end up fragmenting your data into many small files. The default is 900 which also allows some # of files to be open by the scanner before hitting the default Linux limit of 1024.
max_rows_per_file	maximum number of rows per file. If greater than 0 then this will limit how many rows are placed in any single file. Default is 0L.
min_rows_per_group	write the row groups to the disk when this number of rows have accumulated. Default is 0L.
max_rows_per_group	maximum rows allowed in a single group and when this number of rows is exceeded, it is split and the next set of rows is written to the next group. This value must be set such that it is greater than min_rows_per_group. Default is 1024 * 1024.
create_directory	whether to create the directories written into. Requires appropriate permissions on the storage backend. If set to FALSE, directories are assumed to be already present if writing on a classic hierarchical filesystem. Default is TRUE.
preserve_order	Preserve the order of the rows.
...	additional format-specific arguments. For available Parquet options, see write_parquet() . The available IPC options are:

- `use_legacy_format` logical: write data formatted so that Arrow libraries versions 0.14 and lower can read it. Default is FALSE. You can also enable this by setting the environment variable `ARROW_PRE_0_15_IPC_FORMAT=1`.
- `metadata_version`: A string like "V5" or the equivalent integer indicating the Arrow IPC MetadataVersion. Default (NULL) will use the latest version, unless the environment variable `ARROW_PRE_1_0_METADATA_VERSION=1`, in which case it will be V4.
- `codec`: A [Codec](#) which will be used to compress body buffers of written files. Default (NULL) will not compress body buffers.
- `null_fallback`: character to be used in place of missing values (NA or NULL) when using Hive-style partitioning. See [hive_partition\(\)](#).

Value

The input dataset, invisibly

Examples

```
# You can write datasets partitioned by the values in a column (here: "cyl").
# This creates a structure of the form cyl=X/part-Z.parquet.
one_level_tree <- tempfile()
write_dataset(mtcars, one_level_tree, partitioning = "cyl")
list.files(one_level_tree, recursive = TRUE)

# You can also partition by the values in multiple columns
# (here: "cyl" and "gear").
# This creates a structure of the form cyl=X/gear=Y/part-Z.parquet.
two_levels_tree <- tempfile()
write_dataset(mtcars, two_levels_tree, partitioning = c("cyl", "gear"))
list.files(two_levels_tree, recursive = TRUE)

# In the two previous examples we would have:
# X = {4,6,8}, the number of cylinders.
# Y = {3,4,5}, the number of forward gears.
# Z = {0,1,2}, the number of saved parts, starting from 0.

# You can obtain the same result as as the previous examples using arrow with
# a dplyr pipeline. This will be the same as two_levels_tree above, but the
# output directory will be different.
library(dplyr)
two_levels_tree_2 <- tempfile()
mtcars |>
  group_by(cyl, gear) |>
  write_dataset(two_levels_tree_2)
list.files(two_levels_tree_2, recursive = TRUE)

# And you can also turn off the Hive-style directory naming where the column
# name is included with the values by using `hive_style = FALSE`.

# Write a structure X/Y/part-Z.parquet.
two_levels_tree_no_hive <- tempfile()
```

```
mtcars |>
  group_by(cyl, gear) |>
  write_dataset(two_levels_tree_no_hive, hive_style = FALSE)
list.files(two_levels_tree_no_hive, recursive = TRUE)
```

write_delim_dataset	<i>Write a dataset into partitioned flat files.</i>
---------------------	---

Description

The `write_*_dataset()` are a family of wrappers around [write_dataset](#) to allow for easy switching between functions for writing datasets.

Usage

```
write_delim_dataset(
  dataset,
  path,
  partitioning = dplyr::group_vars(dataset),
  basename_template = "part-{i}.txt",
  hive_style = TRUE,
  existing_data_behavior = c("overwrite", "error", "delete_matching"),
  max_partitions = 1024L,
  max_open_files = 900L,
  max_rows_per_file = 0L,
  min_rows_per_group = 0L,
  max_rows_per_group = bitwShiftL(1, 20),
  col_names = TRUE,
  batch_size = 1024L,
  delim = ",",
  na = "",
  eol = "\n",
  quote = c("needed", "all", "none"),
  preserve_order = FALSE
)

write_csv_dataset(
  dataset,
  path,
  partitioning = dplyr::group_vars(dataset),
  basename_template = "part-{i}.csv",
  hive_style = TRUE,
  existing_data_behavior = c("overwrite", "error", "delete_matching"),
  max_partitions = 1024L,
  max_open_files = 900L,
  max_rows_per_file = 0L,
```

```

    min_rows_per_group = 0L,
    max_rows_per_group = bitwShiftL(1, 20),
    col_names = TRUE,
    batch_size = 1024L,
    delim = ",",
    na = "",
    eol = "\n",
    quote = c("needed", "all", "none"),
    preserve_order = FALSE
  )

write_tsv_dataset(
  dataset,
  path,
  partitioning = dplyr::group_vars(dataset),
  basename_template = "part-{i}.tsv",
  hive_style = TRUE,
  existing_data_behavior = c("overwrite", "error", "delete_matching"),
  max_partitions = 1024L,
  max_open_files = 900L,
  max_rows_per_file = 0L,
  min_rows_per_group = 0L,
  max_rows_per_group = bitwShiftL(1, 20),
  col_names = TRUE,
  batch_size = 1024L,
  na = "",
  eol = "\n",
  quote = c("needed", "all", "none"),
  preserve_order = FALSE
)

```

Arguments

dataset	Dataset , RecordBatch , Table , <code>arrow_dplyr_query</code> , or <code>data.frame</code> . If an <code>arrow_dplyr_query</code> , the query will be evaluated and the result will be written. This means that you can <code>select()</code> , <code>filter()</code> , <code>mutate()</code> , etc. to transform the data before it is written if you need to.
path	string path, URI, or <code>SubTreeFileSystem</code> referencing a directory to write to (directory will be created if it does not exist)
partitioning	Partitioning or a character vector of columns to use as partition keys (to be written as path segments). Default is to use the current <code>group_by()</code> columns.
basename_template	string template for the names of files to be written. Must contain "{i}", which will be replaced with an autoincremented integer to generate basenames of datafiles. For example, "part-{i}.arrow" will yield "part-0.arrow", If not specified, it defaults to "part-{i}.<default extension>".
hive_style	logical: write partition segments as Hive-style (key1=value1/key2=value2/file.ext) or as just bare values. Default is TRUE.

existing_data_behavior	The behavior to use when there is already data in the destination directory. Must be one of "overwrite", "error", or "delete_matching". • "overwrite" (the default) then any new files created will overwrite existing files • "error" then the operation will fail if the destination directory is not empty • "delete_matching" then the writer will delete any existing partitions if data is going to be written to those partitions and will leave alone partitions which data is not written to.
max_partitions	maximum number of partitions any batch may be written into. Default is 1024L.
max_open_files	maximum number of files that can be left opened during a write operation. If greater than 0 then this will limit the maximum number of files that can be left open. If an attempt is made to open too many files then the least recently used file will be closed. If this setting is set too low you may end up fragmenting your data into many small files. The default is 900 which also allows some # of files to be open by the scanner before hitting the default Linux limit of 1024.
max_rows_per_file	maximum number of rows per file. If greater than 0 then this will limit how many rows are placed in any single file. Default is 0L.
min_rows_per_group	write the row groups to the disk when this number of rows have accumulated. Default is 0L.
max_rows_per_group	maximum rows allowed in a single group and when this number of rows is exceeded, it is split and the next set of rows is written to the next group. This value must be set such that it is greater than min_rows_per_group. Default is 1024 * 1024.
col_names	Whether to write an initial header line with column names.
batch_size	Maximum number of rows processed at a time. Default is 1024L.
delim	Delimiter used to separate values. Defaults to "," for write_delim_dataset() and write_csv_dataset(), and "\t" for write_tsv_dataset(). Cannot be changed for write_tsv_dataset().
na	a character vector of strings to interpret as missing values. Quotes are not allowed in this string. The default is an empty string "".
eol	the end of line character to use for ending rows. The default is "\n".
quote	How to handle fields which contain characters that need to be quoted. • needed - Enclose all strings and binary values in quotes which need them, because their CSV rendering can contain quotes itself (the default) • all - Enclose all valid values in quotes. Nulls are not quoted. May cause readers to interpret all values as strings if schema is inferred. • none - Do not enclose any values in quotes. Prevents values from containing quotes ("), cell delimiters (,) or line endings (\r, \n), (following RFC4180). If values contain these characters, an error is caused when attempting to write.
preserve_order	Preserve the order of the rows.

Value

The input dataset, invisibly.

See Also

[write_dataset\(\)](#)

write_feather	<i>Write a Feather file (deprecated)</i>
---------------	--

Description

`write_feather()` is deprecated and will be removed in a future release. Use [write_ipc_file\(\)](#) instead.

Column-oriented file format designed for fast reading and writing of data frames. Feather V2 is the Arrow IPC file format. Feather V1 is a legacy format available starting in 2016 that lacks many features, such as the ability to store all Arrow data types, and compression support. Feather V1 is deprecated; use [write_ipc_file\(\)](#) for new files.

Usage

```
write_feather(
  x,
  sink,
  version = 2,
  chunk_size = 65536L,
  compression = c("default", "lz4", "lz4_frame", "uncompressed", "zstd"),
  compression_level = NULL
)
```

Arguments

<code>x</code>	data.frame, RecordBatch , or Table
<code>sink</code>	A string file path, connection, URI, or OutputStream , or path in a file system (SubTreeFileSystem)
<code>version</code>	integer Feather file version, Version 1 or Version 2. Version 2 is the default.
<code>chunk_size</code>	The number of rows that each chunk of data should have in the file. Use a smaller <code>chunk_size</code> when you need faster random row access. Default is 64K.
<code>compression</code>	Name of compression codec to use, if any. Default is "lz4" if LZ4 is available in your build of the Arrow C++ library, otherwise "uncompressed". "zstd" is the other available codec and generally has better compression ratios in exchange for slower read and write performance. "lz4" is shorthand for the "lz4_frame" codec. See codec_is_available() for details. TRUE and FALSE can also be used in place of "default" and "uncompressed".
<code>compression_level</code>	If compression is "zstd", you may specify an integer compression level. If omitted, the compression codec's default compression level is used.

Value

The input `x`, invisibly. Note that if `sink` is an [OutputStream](#), the stream will be left open.

See Also

[write_ipc_file\(\)](#)

write_ipc_file

Write an Arrow IPC file

Description

The Arrow IPC file format provides binary columnar serialization for data frames. It is designed to make reading and writing data frames efficient, and to make sharing data across data analysis languages easy.

Usage

```
write_ipc_file(
  x,
  sink,
  chunk_size = 65536L,
  compression = c("default", "lz4", "lz4_frame", "uncompressed", "zstd"),
  compression_level = NULL
)
```

Arguments

<code>x</code>	data.frame, RecordBatch , or Table
<code>sink</code>	A string file path, connection, URI, or OutputStream , or path in a file system (SubTreeFileSystem)
<code>chunk_size</code>	The number of rows that each chunk of data should have in the file. Use a smaller <code>chunk_size</code> when you need faster random row access. Default is 64K.
<code>compression</code>	Name of compression codec to use, if any. Default is "lz4" if LZ4 is available in your build of the Arrow C++ library, otherwise "uncompressed". "zstd" is the other available codec and generally has better compression ratios in exchange for slower read and write performance. "lz4" is shorthand for the "lz4_frame" codec. See codec_is_available() for details. TRUE and FALSE can also be used in place of "default" and "uncompressed".
<code>compression_level</code>	If compression is "zstd", you may specify an integer compression level. If omitted, the compression codec's default compression level is used.

Value

The input `x`, invisibly. Note that if `sink` is an [OutputStream](#), the stream will be left open.

See Also

[RecordBatchWriter](#) for lower-level access to writing Arrow IPC data.

[Schema](#) for information about schemas and metadata handling.

Examples

```
tf <- tempfile(fileext = ".arrow")
on.exit(unlink(tf))
write_ipc_file(mtcars, tf)
```

write_ipc_stream	<i>Write Arrow IPC stream format</i>
------------------	--------------------------------------

Description

Apache Arrow defines two formats for **serializing data for interprocess communication (IPC)**: a "stream" format and a "file" format. `write_ipc_stream()` and `write_ipc_file()` write those formats, respectively.

Usage

```
write_ipc_stream(x, sink, ...)
```

Arguments

x	data.frame, RecordBatch , or Table
sink	A string file path, connection, URI, or OutputStream , or path in a file system (SubTreeFileSystem)
...	deprecated and ignored.

Value

x, invisibly.

See Also

[write_ipc_file\(\)](#) for writing IPC files. [write_to_raw\(\)](#) to serialize data to a buffer. [RecordBatchWriter](#) for a lower-level interface.

Examples

```
tf <- tempfile()
on.exit(unlink(tf))
write_ipc_stream(mtcars, tf)
```

write_parquet	<i>Write Parquet file to disk</i>
---------------	-----------------------------------

Description

Parquet is a columnar storage file format. This function enables you to write Parquet files from R.

Usage

```
write_parquet(
  x,
  sink,
  chunk_size = NULL,
  version = "2.4",
  compression = default_parquet_compression(),
  compression_level = NULL,
  use_dictionary = NULL,
  write_statistics = NULL,
  data_page_size = NULL,
  use_deprecated_int96_timestamps = FALSE,
  coerce_timestamps = NULL,
  allow_truncated_timestamps = FALSE
)
```

Arguments

x	data.frame, RecordBatch , or Table
sink	A string file path, connection, URI, or OutputStream , or path in a file system (SubTreeFileSystem)
chunk_size	how many rows of data to write to disk at once. This directly corresponds to how many rows will be in each row group in parquet. If NULL, a best guess will be made for optimal size (based on the number of columns and number of rows), though if the data has fewer than 250 million cells (rows x cols), then the total number of rows is used.
version	parquet version: "1.0", "2.4" (default), "2.6", or "latest" (currently equivalent to 2.6). Numeric values are coerced to character.
compression	compression algorithm. Default "snappy". See details.
compression_level	compression level. Meaning depends on compression algorithm
use_dictionary	logical: use dictionary encoding? Default TRUE
write_statistics	logical: include statistics? Default TRUE
data_page_size	Set a target threshold for the approximate encoded size of data pages within a column chunk (in bytes). Default 1 MiB.

```

use_deprecated_int96_timestamps
    logical: write timestamps to INT96 Parquet format, which has been deprecated?
    Default FALSE.
coerce_timestamps
    Cast timestamps a particular resolution. Can be NULL, "ms" or "us". Default
    NULL (no casting)
allow_truncated_timestamps
    logical: Allow loss of data when coercing timestamps to a particular resolution.
    E.g. if microsecond or nanosecond data is lost when coercing to "ms", do not
    raise an exception. Default FALSE.

```

Details

Due to features of the format, Parquet files cannot be appended to. If you want to use the Parquet format but also want the ability to extend your dataset, you can write to additional Parquet files and then treat the whole directory of files as a [Dataset](#) you can query. See the [dataset article](#) for examples of this.

The parameters `compression`, `compression_level`, `use_dictionary` and `write_statistics` support various patterns:

- The default NULL leaves the parameter unspecified, and the C++ library uses an appropriate default for each column (defaults listed above)
- A single, unnamed, value (e.g. a single string for compression) applies to all columns
- An unnamed vector, of the same size as the number of columns, to specify a value for each column, in positional order
- A named vector, to specify the value for the named columns, the default value for the setting is used when not supplied

The `compression` argument can be any of the following (case-insensitive): "uncompressed", "snappy", "gzip", "brotli", "zstd", "lz4", "lzo" or "bz2". Only "uncompressed" is guaranteed to be available, but "snappy" and "gzip" are almost always included. See `codec_is_available()`. The default "snappy" is used if available, otherwise "uncompressed". To disable compression, set `compression = "uncompressed"`. Note that "uncompressed" columns may still have dictionary encoding.

Value

the input `x` invisibly.

See Also

[ParquetFileWriter](#) for a lower-level interface to Parquet writing.

Examples

```

tf1 <- tempfile(fileext = ".parquet")
write_parquet(data.frame(x = 1:5), tf1)

# using compression
if (codec_is_available("gzip")) {

```

```
tf2 <- tempfile(fileext = ".gz.parquet")
write_parquet(data.frame(x = 1:5), tf2, compression = "gzip", compression_level = 5)
}
```

write_to_raw	<i>Write Arrow data to a raw vector</i>
--------------	---

Description

[write_ipc_stream\(\)](#) and [write_ipc_file\(\)](#) write data to a sink and return the data (`data.frame`, `RecordBatch`, or `Table`) they were given. This function wraps those so that you can serialize data to a buffer and access that buffer as a raw vector in R.

Usage

```
write_to_raw(x, format = c("stream", "file"))
```

Arguments

x	<code>data.frame</code> , RecordBatch , or Table
format	one of <code>c("stream", "file")</code> , indicating the IPC format to use

Value

A raw vector containing the bytes of the IPC serialized data.

Examples

```
# The default format is "stream"
mtcars_raw <- write_to_raw(mtcars)
```

Index

`*`, 7
`+`, 7
`.onLoad()`, 76
`/`, 7
`<`, 7
`<=`, 7
`==`, 7
`>`, 7
`>=`, 7
`$NewScan()`, 80
`&`, 7
`%/%`, 7
`%%`, 7
`%in%`, 7
`^`, 7

`abs()`, 7
`acero`, 5, 70
`acos()`, 7
`acosh()`, 7
`across()`, 9
`add_filename()`, 7, 81
`all` Arrow functions, 27
`all()`, 7
`all_of()`, 13
`am()`, 10
`anti_join()`, 5
`any()`, 7
`arrange()`, 5
`Array`, 14, 18–20, 29, 31, 52, 53, 65, 75, 76, 103, 119
`ArrayData`, 15, 15
`Arrays`, 28
`arrow-dplyr` (`acero`), 5
`arrow-functions` (`acero`), 5
`arrow-verbs` (`acero`), 5
`arrow::io::MemoryMappedFile`, 74
`arrow::io::OutputStream`, 89
`arrow_array`, 16
`arrow_info`, 17

`arrow_info()`, 34, 67
`arrow_table`, 17
`arrow_table()`, 19
`arrow_with_acero` (`arrow_info`), 17
`arrow_with_azure` (`arrow_info`), 17
`arrow_with_dataset` (`arrow_info`), 17
`arrow_with_gcs` (`arrow_info`), 17
`arrow_with_json` (`arrow_info`), 17
`arrow_with_parquet` (`arrow_info`), 17
`arrow_with_s3` (`arrow_info`), 17
`arrow_with_substrait` (`arrow_info`), 17
`as.character()`, 8
`as.Date()`, 7
`as.diffftime()`, 8
`as.double()`, 8
`as.integer()`, 8
`as.integer64()`, 9
`as.logical()`, 8
`as.numeric()`, 8
`as.vector()`, 53
`as_arrow_array`, 18
`as_arrow_array()`, 108
`as_arrow_table`, 19
`as_chunked_array`, 20
`as_data_type`, 21
`as_date()`, 10
`as_datetime()`, 10
`as_hms()`, 10
`as_record_batch`, 22
`as_record_batch_reader`, 23
`as_schema`, 24
`asin()`, 8
`asinh()`, 8
`atan()`, 8
`atanh()`, 8
`az_container`, 25
`AzureFileSystem` (`FileSystem`), 57

`between()`, 9
`binary` (data-type), 43

- bool (data-type), 43
- bool(), 96
- boolean (data-type), 43
- boolean(), 100
- Buffer, 26, 26, 66, 73, 105
- buffer, 25
- BufferOutputStream (OutputStream), 87
- BufferedReader (InputStream), 66
- c.Array (concat_arrays), 31
- call_function, 27
- call_function(), 69, 107
- case_when(), 9
- cast(), 7
- ceiling(), 8
- ceiling_date(), 10
- character(), 76
- chunked arrays, 115
- chunked_array, 28
- chunked_array(), 20
- ChunkedArray, 6, 20, 21, 28, 28, 31, 53, 65, 88
- coalesce(), 9
- Codec, 30, 31, 123
- codec_is_available, 30
- codec_is_available(), 30, 127, 128, 131
- collapse(), 5
- collect(), 5
- compressed input and output streams, 30
- CompressedInputStream (compression), 31
- CompressedOutputStream (compression), 31
- compression, 31
- compute(), 5
- concat_arrays, 31
- concat_tables, 32
- contains(), 13
- copy_files, 32
- cos(), 8
- cosh(), 8
- count(), 5
- cpu_count, 33
- create_package_with_all_dependencies, 34
- CSV conversion options, 86, 95
- CSV parsing options, 86, 94
- CSV reading options, 86, 95
- CSV write options, 121
- csv_convert_options, 35
- csv_parse_options, 36
- csv_read_options, 37
- csv_write_options, 39
- CsvConvertOptions, 56, 63
- CsvConvertOptions (CsvReadOptions), 40
- CsvFileFormat, 39
- CsvFragmentScanOptions, 56
- CsvFragmentScanOptions (FragmentScanOptions), 63
- CsvParseOptions, 56
- CsvParseOptions (CsvReadOptions), 40
- CsvReadOptions, 40, 43, 56, 63
- CsvTableReader, 42, 95
- CsvWriteOptions (CsvReadOptions), 40
- data type, 14, 16, 28, 65, 76, 110, 113, 119
- data types, 64, 113
- data-type, 43
- data.frame(), 8, 53
- Dataset, 47, 48, 52, 80, 107, 111, 113, 121, 125, 131
- dataset_factory, 48
- dataset_factory(), 47, 79, 84
- DatasetFactory, 48
- DatasetFactory (Dataset), 47
- DataType, 21, 46, 50, 51–55, 108
- date(), 10
- date32 (data-type), 43
- date32(), 96
- date64 (data-type), 43
- date_decimal(), 10
- day(), 10
- ddays(), 10
- decimal (data-type), 43
- decimal128 (data-type), 43
- decimal256 (data-type), 43
- decimal32 (data-type), 43
- decimal64 (data-type), 43
- decimal_date(), 10
- desc(), 9
- dhours(), 10
- dictionary, 51
- dictionary(), 46, 96
- DictionaryArray (Array), 14
- DictionaryType, 51, 51
- difftime(), 8
- DirectoryPartitioning (Partitioning), 91
- DirectoryPartitioningFactory (Partitioning), 91
- distinct(), 5
- dmicroseconds(), 10

- dmilliseconds(), 10
- dminutes(), 10
- dmonths(), 10
- dmy(), 10
- dmy_h(), 10
- dmy_hm(), 10
- dmy_hms(), 10
- dnanoseconds(), 10
- dpicoseconds(), 10
- dplyr::filter(), 107
- dplyr::mutate(), 107
- dseconds(), 10
- dst(), 10
- duration (data-type), 43
- dweeks(), 10
- dyears(), 10
- dym(), 10
- ends_with(), 13
- endsWith(), 8
- epiweek(), 10
- epiyear(), 11
- everything(), 13
- exp(), 8
- explain(), 5
- expm1(), 8
- Expression, 52, 112
- ExtensionArray, 52, 53, 76, 120
- ExtensionType, 53, 53, 76, 120
- fast_strptime(), 11
- FeatherReader, 53, 98
- Field, 54, 55, 113
- field, 54
- fields, 103, 113
- FileFormat, 39, 40, 47–49, 55, 69, 79, 122
- FileFormat\$create(), 50
- FileInfo, 56, 59
- FileOutputStream (OutputStream), 87
- FileSelector, 47, 57, 59
- FileSystem, 47, 49, 57, 79, 84
- FileSystemDataset (Dataset), 47
- FileSystemDatasetFactory, 91
- FileSystemDatasetFactory (Dataset), 47
- FileWriteOptions, 60
- filter(), 5
- filter_out(), 5
- fixed_size_binary (data-type), 43
- fixed_size_list_of (data-type), 43
- FixedSizeListArray (Array), 14
- FixedSizeListType (data-type), 43
- FixedWidthType, 51
- flight_connect, 61
- flight_connect(), 62, 70
- flight_disconnect, 61
- flight_get, 62
- flight_path_exists (list_flights), 70
- flight_put, 62
- float (data-type), 43
- float16 (data-type), 43
- float32 (data-type), 43
- float64 (data-type), 43
- float64(), 96, 100
- floor(), 8
- floor_date(), 11
- force_tz(), 11
- format(), 8
- format_ISO8601(), 11
- FragmentScanOptions, 56, 63
- full_join(), 6
- GcsFileSystem (FileSystem), 57
- glimpse(), 6
- grepl(), 8
- group_by(), 6
- group_by_drop_default(), 6
- group_vars(), 6
- groups(), 6
- gs_bucket, 63
- gsub(), 8
- halffloat (data-type), 43
- hive_partition, 64
- hive_partition(), 49, 79, 84, 123
- hive_partition(...), 92
- HivePartitioning, 64
- HivePartitioning (Partitioning), 91
- HivePartitioningFactory (Partitioning), 91
- hms(), 10
- hour(), 11
- if_all(), 9
- if_any(), 9
- if_else(), 9
- ifelse(), 8
- infer_schema, 65
- infer_type, 65

- `infer_type()`, [19](#), [20](#), [50](#)
- `InMemoryDataset (Dataset)`, [47](#)
- `inner_join()`, [6](#)
- `input file`, [60](#)
- `input stream`, [60](#)
- `InputStream`, [31](#), [43](#), [66](#), [74](#), [94](#), [97–100](#), [102](#), [105](#)
- `install_arrow`, [66](#)
- `install_pyarrow`, [68](#)
- `int16 (data-type)`, [43](#)
- `int32 (data-type)`, [43](#)
- `int32()`, [51](#), [52](#), [96](#)
- `int64 (data-type)`, [43](#)
- `int64()`, [100](#)
- `int8 (data-type)`, [43](#)
- `io_thread_count`, [68](#)
- `IpcFileFormat (FileFormat)`, [55](#)
- `is()`, [12](#)
- `is.character()`, [8](#)
- `is.Date()`, [11](#)
- `is.double()`, [8](#)
- `is.factor()`, [8](#)
- `is.finite()`, [8](#)
- `is.infinite()`, [8](#)
- `is.instant()`, [11](#)
- `is.integer()`, [8](#)
- `is.integer64()`, [9](#)
- `is.list()`, [8](#)
- `is.logical()`, [8](#)
- `is.na()`, [8](#)
- `is.nan()`, [8](#)
- `is.numeric()`, [8](#)
- `is.POSIXct()`, [11](#)
- `is.timepoint()`, [11](#)
- `is_character()`, [12](#)
- `is_double()`, [12](#)
- `is_in (match_arrow)`, [72](#)
- `is_integer()`, [12](#)
- `is_list()`, [12](#)
- `is_logical()`, [12](#)
- `ISOdate()`, [7](#)
- `ISOdatetime()`, [7](#)
- `isoweek()`, [11](#)
- `isoyear()`, [11](#)
- `JsonFileFormat`, [69](#)
- `JsonFragmentScanOptions (FragmentScanOptions)`, [63](#)
- `JsonParseOptions (CsvReadOptions)`, [40](#)
- `JsonReadOptions (CsvReadOptions)`, [40](#)
- `JsonTableReader`, [99](#)
- `JsonTableReader (CsvTableReader)`, [42](#)
- `large_binary (data-type)`, [43](#)
- `large_list_of (data-type)`, [43](#)
- `large_utf8 (data-type)`, [43](#)
- `LargeListArray (Array)`, [14](#)
- `last_col()`, [13](#)
- `leap_year()`, [11](#)
- `left_join()`, [6](#)
- `list_compute_functions`, [69](#)
- `list_flights`, [70](#)
- `list_of (data-type)`, [43](#)
- `list_of()`, [100](#)
- `ListArray (Array)`, [14](#)
- `load_flight_server`, [71](#)
- `LocalFileSystem (FileSystem)`, [57](#)
- `log()`, [8](#)
- `log10()`, [8](#)
- `log1p()`, [8](#)
- `log2()`, [8](#)
- `logb()`, [8](#)
- `make_date()`, [11](#)
- `make_datetime()`, [11](#)
- `make_difftime()`, [11](#)
- `map_batches`, [71](#)
- `map_of (data-type)`, [43](#)
- `MapArray (Array)`, [14](#)
- `MapType (data-type)`, [43](#)
- `match_arrow`, [72](#)
- `matches()`, [13](#)
- `max()`, [8](#)
- `mday()`, [11](#)
- `mdy()`, [11](#)
- `mdy_h()`, [11](#)
- `mdy_hm()`, [11](#)
- `mdy_hms()`, [11](#)
- `mean()`, [8](#)
- `median()`, [12](#)
- `MemoryMappedFile (InputStream)`, [66](#)
- `Message`, [73](#)
- `MessageReader`, [74](#)
- `min()`, [9](#)
- `minute()`, [11](#)
- `mmap_create`, [74](#)
- `mmap_open`, [75](#)
- `mmap_open()`, [66](#)

- month(), [11](#)
- mutate(), [6](#)
- my(), [11](#)
- myd(), [11](#)
- n(), [9](#)
- n_distinct(), [10](#)
- nchar(), [9](#)
- new_extension_array
 - (new_extension_type), [75](#)
- new_extension_array(), [76](#)
- new_extension_type, [75](#)
- new_extension_type(), [76](#)
- null (data-type), [43](#)
- null(), [96](#), [100](#)
- num_range(), [13](#)
- one_of(), [13](#)
- open_csv_dataset (open_delim_dataset), [82](#)
- open_dataset, [78](#), [82](#)
- open_dataset(), [47](#), [48](#), [50](#), [86](#)
- open_delim_dataset, [82](#)
- open_tsv_dataset (open_delim_dataset), [82](#)
- Other Arrow data types, [51](#)
- output stream, [60](#)
- OutputStream, [31](#), [87](#), [116](#), [120](#), [121](#), [127–130](#)
- ParquetArrowReaderProperties, [87](#), [88](#), [102](#)
- ParquetFileFormat (FileFormat), [55](#)
- ParquetFileReader, [87](#), [88](#), [90](#)
- ParquetFileWriter, [89](#), [90](#), [131](#)
- ParquetFragmentScanOptions
 - (FragmentScanOptions), [63](#)
- ParquetReaderProperties, [88](#), [90](#)
- ParquetWriterProperties, [89](#), [90](#)
- parse_date_time(), [11](#)
- Partitioning, [91](#)
- paste(), [9](#)
- paste0(), [9](#)
- pm(), [11](#)
- pmax(), [9](#)
- pmin(), [9](#)
- prod(), [9](#)
- pull(), [6](#)
- qday(), [11](#)
- quantile(), [12](#)
- quarter(), [11](#)
- R6::R6Class, [76](#)
- RandomAccessFile, [105](#)
- RandomAccessFile (InputStream), [66](#)
- raw, [94](#), [100](#)
- raw(), [53](#), [76](#)
- read.csv(), [35](#), [41](#)
- read_csv2_arrow (read_delim_arrow), [92](#)
- read_csv_arrow (read_delim_arrow), [92](#)
- read_csv_arrow(), [40](#), [42](#), [50](#), [56](#), [80](#), [82](#)
- read_delim_arrow, [92](#)
- read_delim_arrow(), [42](#), [82](#), [86](#)
- read_feather, [97](#)
- read_feather(), [53](#)
- read_ipc_file, [98](#)
- read_ipc_file(), [80](#), [97](#), [99](#), [105](#)
- read_ipc_stream, [99](#)
- read_ipc_stream(), [105](#)
- read_json_arrow, [99](#)
- read_json_arrow(), [40](#), [42](#)
- read_message, [101](#)
- read_parquet, [101](#)
- read_parquet(), [80](#)
- read_schema, [102](#)
- read_tsv_arrow (read_delim_arrow), [92](#)
- read_tsv_arrow(), [82](#)
- ReadableFile (InputStream), [66](#)
- readr::read_csv(), [35](#), [41](#)
- recode_values(), [10](#)
- record batches, [115](#)
- record_batch, [103](#)
- record_batch(), [22](#)
- RecordBatch, [22](#), [32](#), [53](#), [62](#), [89](#), [103](#), [106](#), [120](#), [121](#), [125](#), [127–130](#), [132](#)
- RecordBatchFileReader
 - (RecordBatchReader), [104](#)
- RecordBatchFileWriter
 - (RecordBatchWriter), [106](#)
- RecordBatchReader, [23](#), [98](#), [99](#), [104](#)
- RecordBatchStreamReader
 - (RecordBatchReader), [104](#)
- RecordBatchStreamWriter
 - (RecordBatchWriter), [106](#)
- RecordBatchWriter, [106](#), [129](#)
- register_extension_type
 - (new_extension_type), [75](#)
- register_extension_type(), [76](#)

- register_scalar_function, [107](#)
- register_scalar_function(), [107](#)
- relocate(), [6](#)
- rename(), [6](#)
- rename_with(), [6](#)
- replace_values(), [10](#)
- replace_when(), [10](#)
- reregister_extension_type
 - (new_extension_type), [75](#)
- reregister_extension_type(), [76](#)
- right_join(), [6](#)
- round(), [9](#)
- round_date(), [11](#)
- s3_bucket, [109](#), [109](#)
- s3_bucket(), [79](#), [84](#)
- S3FileSystem, [109](#)
- S3FileSystem (FileSystem), [57](#)
- Scalar, [110](#)
- scalar, [110](#)
- Scanner, [52](#), [111](#)
- ScannerBuilder, [48](#)
- ScannerBuilder (Scanner), [111](#)
- Schema, [18](#), [20](#), [22](#), [24](#), [47](#), [48](#), [56](#), [79](#), [84](#), [85](#), [89](#), [91](#), [92](#), [94](#), [95](#), [100](#), [102](#), [103](#), [105](#), [106](#), [112](#), [113](#), [113](#), [118](#), [129](#)
- schema, [112](#)
- schema(), [23](#), [24](#), [43](#), [64](#), [72](#), [108](#)
- Schemas, [55](#)
- sd(), [12](#)
- second(), [11](#)
- select(), [6](#)
- semester(), [11](#)
- semi_join(), [6](#)
- serialize(), [119](#)
- set_cpu_count (cpu_count), [33](#)
- set_io_thread_count (io_thread_count), [68](#)
- show_exec_plan, [114](#)
- show_query(), [6](#)
- sign(), [9](#)
- sin(), [9](#)
- sinh(), [9](#)
- slice_head(), [6](#)
- slice_max(), [6](#)
- slice_min(), [6](#)
- slice_sample(), [6](#)
- slice_tail(), [6](#)
- sqrt(), [9](#)
- starts_with(), [13](#)
- startsWith(), [9](#)
- str_c(), [13](#)
- str_count(), [13](#)
- str_detect(), [13](#)
- str_dup(), [13](#)
- str_ends(), [13](#)
- str_ilike(), [13](#)
- str_length(), [13](#)
- str_like(), [13](#)
- str_pad(), [13](#)
- str_remove(), [13](#)
- str_remove_all(), [13](#)
- str_replace(), [13](#)
- str_replace_all(), [13](#)
- str_replace_na(), [13](#)
- str_split(), [13](#)
- str_starts(), [13](#)
- str_sub(), [13](#)
- str_to_lower(), [13](#)
- str_to_title(), [13](#)
- str_to_upper(), [13](#)
- str_trim(), [13](#)
- strftime(), [9](#)
- stri_reverse(), [12](#)
- string (data-type), [43](#)
- strptime, [36](#), [42](#), [86](#), [95](#)
- strptime(), [9](#), [42](#)
- strrep(), [9](#)
- strsplit(), [9](#)
- struct (data-type), [43](#)
- struct(), [100](#)
- StructArray (Array), [14](#)
- StructScalar (scalar), [110](#)
- sub(), [9](#)
- substr(), [9](#)
- substring(), [9](#)
- SubTreeFileSystem (FileSystem), [57](#)
- sum(), [9](#)
- summarise(), [6](#)
- Table, [5](#), [18–20](#), [28](#), [32](#), [53](#), [62](#), [89](#), [95](#), [97–100](#), [102](#), [105–107](#), [112](#), [113](#), [115](#), [120](#), [121](#), [125](#), [127–130](#), [132](#)
- tally(), [6](#)
- tan(), [9](#)
- tanh(), [9](#)
- tibble(), [13](#)

tidy selection specification, [94](#), [97](#), [98](#),
 [100](#), [102](#)
 time32 (data-type), [43](#)
 time32(), [96](#)
 time64 (data-type), [43](#)
 timestamp (data-type), [43](#)
 timestamp(), [36](#), [86](#), [95](#)
 TimestampParser, [36](#), [42](#), [86](#), [95](#)
 TimestampParser (CsvReadOptions), [40](#)
 to_arrow, [116](#)
 to_duckdb, [117](#)
 tolower(), [9](#)
 toupper(), [9](#)
 transmute(), [6](#)
 trunc(), [9](#)
 Type, [53](#)
 type, [19](#), [20](#)
 tz(), [11](#)

uint16 (data-type), [43](#)
 uint32 (data-type), [43](#)
 uint64 (data-type), [43](#)
 uint8 (data-type), [43](#)
 ungroup(), [6](#)
 unify_schemas, [118](#)
 union(), [6](#)
 union_all(), [6](#)
 UnionDataset (Dataset), [47](#)
 unregister_extension_type
 (new_extension_type), [75](#)
 utf8 (data-type), [43](#)
 utf8(), [51](#), [52](#), [96](#), [100](#)

value_counts, [119](#)
 var(), [12](#)
 vctrs extension type, [75](#)
 vctrs::vec_data(), [119](#)
 vctrs::vec_is(), [119](#)
 vctrs::vec_ptype(), [119](#)
 vctrs::vec_restore(), [119](#)
 vctrs_extension_array, [119](#)
 vctrs_extension_array(), [119](#)
 vctrs_extension_type
 (vctrs_extension_array), [119](#)
 vctrs_extension_type(), [76](#)

wday(), [11](#)
 week(), [12](#)
 when_all(), [10](#)
 when_any(), [10](#)
 with_tz(), [12](#)
 write_csv_arrow, [120](#)
 write_csv_dataset
 (write_delim_dataset), [124](#)
 write_dataset, [121](#), [124](#)
 write_dataset(), [127](#)
 write_delim_dataset, [124](#)
 write_feather, [127](#)
 write_ipc_file, [128](#)
 write_ipc_file(), [99](#), [106](#), [127–129](#), [132](#)
 write_ipc_stream, [129](#)
 write_ipc_stream(), [106](#), [132](#)
 write_parquet, [90](#), [91](#), [130](#)
 write_parquet(), [122](#)
 write_to_raw, [132](#)
 write_to_raw(), [106](#), [129](#)
 write_tsv_dataset
 (write_delim_dataset), [124](#)

yday(), [12](#)
 ydm(), [12](#)
 ydm_h(), [12](#)
 ydm_hm(), [12](#)
 ydm_hms(), [12](#)
 year(), [12](#)
 ym(), [12](#)
 ymd(), [12](#)
 ymd_h(), [12](#)
 ymd_hm(), [12](#)
 ymd_hms(), [12](#)
 yq(), [12](#)