

Package ‘dbi.table’

December 10, 2024

Type Package

Title Database Queries Using 'data.table' Syntax

Version 1.0.1

Depends data.table

Imports DBI, bit64, dbplyr, methods, rlang, stringi, utils

Suggests R.utils, RSQLite, duckdb, testthat (>= 3.0.0), knitr,
rmarkdown

Description Query database tables over a 'DBI' connection using 'data.table' syntax.
Attach database schemas to the search path. Automatically merge using foreign
key constraints.

License MPL-2.0

URL <https://github.com/kjellpk/dbi.table>

BugReports <https://github.com/kjellpk/dbi.table/issues>

VignetteBuilder knitr

Encoding UTF-8

Config/testthat/edition 3

RoxygenNote 7.3.2

NeedsCompilation no

Author Kjell P. Konis [aut, cre]

Maintainer Kjell P. Konis <kjellk@gmail.com>

Repository CRAN

Date/Publication 2024-12-10 21:30:05 UTC

Contents

dbi.table-package	2
as.data.table	4
as.dbi.table	5
csql	6

dbExecute,dbi.catalog,SQL-method	6
dbi.attach	7
dbi.catalog	8
example_databases	9
merge	9
reference.test	11
sql.join	12

Index	14
--------------	-----------

dbi.table-package	<i>DBI Table</i>
-------------------	------------------

Description

A `dbi.table` is a data structure that describes a SQL query (called the `dbi.table`'s *underlying SQL query*). This query can be manipulated using `data.table`'s `[i, j, by]` syntax.

Usage

```
dbi.table(conn, id)
```

```
## S3 method for class 'dbi.table'
x[i, j, by, nomatch = NA, on = NULL]
```

Arguments

- | | |
|-------------------|--|
| <code>conn</code> | A <code>DBIConnection</code> object, as returned by <code>dbConnect</code> . Alternatively, a <code>dbi.catalog</code> or a <code>dbi.table</code> , in which case the new <code>dbi.table</code> will use the connection embedded in the provided object. |
| <code>id</code> | An <code>Id</code> , a character string (which will be converted to an <code>Id</code> by <code>Id</code>), or a <code>SQL</code> object (advanced) identifying a database object (e.g., table or view) on <code>conn</code> . |
| <code>x</code> | A <code>dbi.table</code> . |
| <code>i</code> | A logical expression of the columns of <code>x</code> , a <code>dbi.table</code> , or a <code>data.frame</code> . Use <code>i</code> to select a subset of the rows of <code>x</code> . Note: unlike <code>data.table</code> , <code>i</code> <i>cannot</i> be a vector.
When <code>i</code> is a logical expression, the rows where the expression is <code>TRUE</code> are returned. If the expression contains a symbol <code>foo</code> that is not a column name of <code>x</code> but that is present in the calling scope, then the value of <code>foo</code> will be substituted into the expression if <code>foo</code> is a scalar, or if <code>foo</code> is a vector and is the right-hand-side argument to <code>%in%</code> or <code>%chin%</code> (substitution occurs when the <code>extract()</code> method is evaluated).
When <code>i</code> inherits from <code>data.frame</code> , it is coerced to a <code>dbi.table</code> .
When <code>i</code> is a <code>dbi.table</code> , the rows of <code>x</code> that match (according to the condition specified in <code>on</code>) the rows of <code>i</code> are returned. When <code>nomatch == NA</code> , all rows of <code>i</code> are returned (right outer join); when <code>nomatch == NULL</code> , only the rows of <code>i</code> that match a row of <code>x</code> are returned (inner join). |

j	A list of expressions, a literal character vector of column names of <code>x</code> , an expression of the form <code>start_name:end_name</code> , or a literal numeric vector of integer values indexing the columns of <code>x</code> . Use <code>j</code> to select (and optionally, transform) the columns of <code>x</code> .
by	A list of expressions, a literal character vector of column names of <code>x</code> , an expression of the form <code>start_name:end_name</code> , or a literal numeric vector of integer values indexing the columns of <code>x</code> . Use <code>by</code> to control grouping when evaluating <code>j</code> .
nomatch	Either NA or NULL.
on	<ul style="list-style-type: none"> • An unnamed character vector, e.g., <code>x[i, on = c("a", "b")]</code>, used when columns <code>a</code> and <code>b</code> are common to both <code>x</code> and <code>i</code>. • Foreign key joins: As a named character vector when the join columns have different names in <code>x</code> and <code>i</code>. For example, <code>x[i, on = c(x1 = "i1", x2 = "i2")]</code> joins <code>x</code> and <code>i</code> by matching columns <code>x1</code> and <code>x2</code> in <code>x</code> with columns <code>i1</code> and <code>i2</code> in <code>i</code>, respectively. • Foreign key joins can also use the binary operator <code>==</code>, e.g., <code>x[i, on = c("x1 == i1", "x2 == i2")]</code>. • It is also possible to use <code>.</code> syntax as <code>x[i, on = .(a, b)]</code>. • Non-equi joins using binary operators <code>>=</code>, <code>></code>, <code><=</code>, <code><</code> are also possible, e.g., <code>x[i, on = c("x >= a", "y <= b")]</code>, or <code>x[i, on = .(x >= a, y <= b)]</code>.

Value

A `dbi.table`.

See Also

- [as.data.table](#) to retrieve the *results set* as a `data.table`,
- [csql](#) to see the underlying SQL query.

Examples

```
# open a connection to the Chinook example database using duckdb
duck <- chinook.duckdb()

# create a dbi.table corresponding to the Album table on duck
Album <- dbi.table(duck, DBI::Id(table_name = "Album"))

# the print method displays a 5 row preview
# print(Album)
Album

# 'id' can also be 'SQL'; use the same DBI connection as Album
Genre <- dbi.table(Album, DBI::SQL("chinook_duckdb.main.Genre"))

# use the extract (\code{[]}) method to subset the dbi.table
Album[AlbumId < 5, .(Title, nchar = paste(nchar(Title), "characters"))]

# use csql to see the underlying SQL query
```

```
csql(Album[AlbumId < 5, #WHERE
      .(Title, #SELECT
        nchar = paste(nchar(Title), "characters"))])
```

as.data.table	<i>Coerce to data.table</i>
---------------	-----------------------------

Description

Execute a [dbi.table](#)'s underlying SQL query and return the result set as a [data.table](#). By default, the result set is limited to 10,000 rows. See Details.

Usage

```
## S3 method for class 'dbi.table'
as.data.table(
  x,
  keep.rownames = FALSE,
  ...,
  n = getOption("dbi_table_max_fetch", 10000L)
)
```

Arguments

x	a dbi.table .
keep.rownames	a logical value. This argument is not used.
...	additional arguments are ignored.
n	an integer value. When nonnegative, the underlying SQL query includes a 'LIMIT n' clause and n is also passed to dbFetch . When negative, the underlying SQL query does not include a LIMIT clause and all rows in the result set are returned.

Details

By default, `as.data.table` returns up to 10,000 rows (see the `n` argument). To override this limit, either call `as.data.table` and provide the `n` argument (e.g., `n = -1` to return the entire result set), or set the option `dbi_table_max_fetch` to the desired default value of `n`.

Value

a `dbi.table`.

See Also

[as.data.table](#) (the generic method in the **data.table** package).

Examples

```
duck <- chinook.duckdb()
Artist <- dbi.table(duck, DBI::Id("Artist"))

as.data.table(Artist, n = 7)[[]]
```

as.dbi.table

Coerce to DBI Table

Description

Test whether an object is a `dbi.table`, or coerce it if possible.

Usage

```
is.dbi.table(x)

as.dbi.table(conn, x, type = c("auto", "query", "temporary"))
```

Arguments

<code>x</code>	any R object.
<code>conn</code>	a connection handle returned by <code>dbConnect</code> . Alternatively, <code>conn</code> may be a <code>dbi.table</code> or a <code>dbi.catalog</code> ; in these cases, the connection handle is extracted from the provided object.
<code>type</code>	a character string. Possible choices are "auto", "query", and "temporary". See Details. The default "auto" uses <i>In Query</i> tables when <code>x</code> has 500 or fewer rows or when creating a temporary table on the database fails.

Details

Two types of tables are provided: *Temporary* (when `type == "temporary"`) and *In Query* (when `type == "query"`). For *Temporary*, the data are written to a SQL temporary table and the associated `dbi.table` is returned. For *In Query*, the data are written into a CTE as part of the query itself - useful when the connection does not permit creating temporary tables.

Value

a `dbi.table`.

Note

The temporary tables created by this function are dropped (by calling `dbRemoveTable`) during garbage collection when they are no longer referenced.

Examples

```
duck <- dbi.catalog(chinook.duckdb)
csql(as.dbi.table(duck, iris[1:4, 1:3], type = "query"))
```

csql	<i>See SQL</i>
------	----------------

Description

View a `dbi.table`'s underlying SQL query.

Usage

```
csql(x, n = getOption("dbi_table_max_fetch", 10000L))
```

Arguments

x	a <code>dbi.table</code> .
n	a single integer value. When nonnegative, limits the number of rows returned by the query to n.

Value

none (invisible NULL).

dbExecute,dbi.catalog,SQL-method
DBI Methods for dbi.tables

Description

Call DBI methods using the underlying DBI connection.

Usage

```
## S4 method for signature 'dbi.catalog,SQL'
dbExecute(conn, statement, ...)
```

```
## S4 method for signature 'dbi.schema,SQL'
dbExecute(conn, statement, ...)
```

```
## S4 method for signature 'dbi.table,SQL'
dbExecute(conn, statement, ...)
```

```

## S4 method for signature 'dbi.table,missing'
dbSendStatement(
  conn,
  statement,
  ...,
  n = getOption("dbi_table_max_fetch", 10000L)
)

## S4 method for signature 'dbi.catalog'
dbGetInfo(dbObj, ...)

## S4 method for signature 'dbi.schema'
dbGetInfo(dbObj, ...)

## S4 method for signature 'dbi.table'
dbGetInfo(dbObj, ...)

```

Arguments

conn	a dbi.catalog , dbi.schema , or dbi.table .
statement	a SQL object.
...	other parameters passed on to methods.
n	an integer value. A nonnegative value limits the number of records returned by the query. A negative value omits the LIMIT (or TOP) clause entirely.
dbObj	a dbi.catalog , dbi.schema , or dbi.table .

See Also

[dbExecute](#), [dbGetInfo](#), [dbSendStatement](#)

dbi.attach

Attach a Database Schema to the Search Path

Description

Create a [dbi.table](#) for each database object in a schema and place them on the search path.

Usage

```

dbi.attach(
  what,
  pos = 2L,
  name = NULL,
  warn.conflicts = FALSE,
  schema = NULL,
  graphics = TRUE
)

```

Arguments

what	a connection handle returned by dbConnect or a zero-argument function that returns a connection handle.
pos	an integer specifying position in search() where to attach.
name	a character string specifying the name to use for the attached database.
warn.conflicts	a logical value. If TRUE, warnings are printed about conflicts from attaching the database, unless that database contains an object <code>.conflicts.OK</code> . A conflict is a function masking a function, or a non-function masking a non-function.
schema	a character string specifying the name of the schema to attach.
graphics	a logical value; passed to menu . In interactive sessions, when schema is NULL and multiple schemas are found on what, a menu is displayed to select a schema.

Value

an [environment](#), the attached schema is invisibly returned.

See Also

[attach](#)

dbi.catalog

Create a dbi.catalog

Description

A `dbi.catalog` represents a database catalog.

Usage

```
dbi.catalog(conn)
```

Arguments

conn	a connection handle returned by dbConnect or a zero-argument function that returns a connection handle.
------	---

Value

`dbi.catalog` returns a `dbi.catalog` (internally an [environment](#) with the class attribute set to `"dbi.catalog"`).

Examples

```
# chinook.duckdb is a zero-argument function that returns a DBI handle
(db <- dbi.catalog(chinook.duckdb))

# a dbi.catalog corresponds to a catalog - list the schemas
ls(db)

# list the tables in the schema 'main'
ls(db$main)
```

example_databases	<i>Example Databases</i>
-------------------	--------------------------

Description

These zero-argument functions return connections to the example databases included in the **dbi.table** package.

Usage

```
chinook.sqlite()
chinook.duckdb()
```

Value

a [DBIConnection](#) object, as returned by [dbConnect](#).

merge	<i>Merge two dbi.tables</i>
-------	-----------------------------

Description

Merge two [dbi.tables](#). The `dbi.table` method is similar to the [data.table](#) method except that the result set is only determined up to row order and is not sorted by default.

Default merge columns: if `x` has a foreign key constraint that references `y` then the columns comprising this key are used; see details. When a foreign key cannot be found, then the common columns between the two `dbi.tables` are used.

Use the `by`, `by.x`, and `by.y` arguments explicitly to override this default.

Usage

```
## S3 method for class 'dbi.table'
merge(
  x,
  y,
  by = NULL,
  by.x = NULL,
  by.y = NULL,
  all = FALSE,
  all.x = all,
  all.y = all,
  sort = FALSE,
  suffixes = c(".x", ".y"),
  no.dups = TRUE,
  recursive = FALSE,
  ...
)
```

Arguments

<code>x, y</code>	<code>dbi.tables</code> sharing the same DBI connection.
<code>by</code>	A vector of shared column names in <code>x</code> and <code>y</code> to merge on.
<code>by.x, by.y</code>	character vectors of column names in <code>x</code> and <code>y</code> to merge on.
<code>all</code>	a logical value. <code>all = TRUE</code> is shorthand to save setting both <code>all.x = TRUE</code> and <code>all.y = TRUE</code> .
<code>all.x</code>	a logical value. When <code>TRUE</code> , rows from <code>x</code> that do not have a matching row in <code>y</code> are included. These rows will have NAs in the columns that are filled with values from <code>y</code> . The default is <code>FALSE</code> so that only rows with data from both <code>x</code> and <code>y</code> are included in the output.
<code>all.y</code>	a logical value. Analogous to <code>all.x</code> above.
<code>sort</code>	a logical value. Currently ignored.
<code>suffixes</code>	a length-2 character vector. The suffixes to be used for making non-by column names unique. The suffix behavior works in a similar fashion to the <code>merge.data.frame</code> method.
<code>no.dups</code>	a logical value. When <code>TRUE</code> , suffixes are also appended to non-by <code>y</code> column names in <code>y</code> when they have the same column name as any <code>by.x</code> .
<code>recursive</code>	a logical value. Only used when <code>y</code> is missing. When <code>TRUE</code> , <code>merge</code> is called recursively on each of the just-merged <code>dbi.tables</code> . See examples.
<code>...</code>	additional arguments are ignored.

Details

Foreign key constraints. Foreign keys can only be queried when (1) the `dbi.table`'s schema is loaded, and (2) `dbi.table` understands the underlying database's information schema. The merge is done at the SQL level and can merge on columns that are not in the `dbi.tables`.

This function uses `sql.join` to join x and y then formats the result set to match the typical merge result.

Value

a `dbi.table`.

Examples

```
dbi.attach(chinook.duckdb, name = "merge-example")

#The Album table has a foreign key constraint that references Artist
merge(Album, Artist)

#When y is omitted, x's foreign key relationship is used to determine y
merge(Album)

#Multiple foreign keys are supported
csql(merge(Track))

#Track reference Album but not Artist, Album does reference Artist
#This dbi.table includes Artist.Name as well
csql(merge(Track, recursive = TRUE))
```

reference.test

Test dbi.table vs. Reference Implementation

Description

Evaluate an expression including at least one `dbi.table` and compare the result with the *Reference Implementation*. This function is primarily for testing and is potentially very slow for large tables.

Usage

```
reference.test(
  expr,
  envir = parent.frame(),
  ignore.row.order = TRUE,
  verbose = TRUE
)
```

Arguments

`expr` an expression involving at least one `dbi.table` and whose result can be coerced into a `data.table`.

`envir` an environment. Where to evaluate `expr`.

`ignore.row.order` a logical value. This argument is passed to `all.equal`.

`verbose` a logical value. When TRUE, the output from `all.equal` is displayed in a message when `all.equal` returns anything other than TRUE.

Value

a logical value.

Reference Implementation

Suppose that `id1` identifies a table in a SQL database and that `[i, j, by]` describes a subset/select/summarize operation using `data.table` syntax. The *Reference Implementation* for this operation is:

```
setDT(dbReadTable(conn, id1))[i, j, by]
```

More generally, for an expression involving multiple SQL database objects and using `data.table` syntax, the *Reference Implementation* would be to download each of these objects in their entirety, convert them to `data.tables`, then evaluate the expression.

The goal of the **dbi.table** is to generate an SQL query that produces the same results set as the Reference Implementation up to row ordering.

Examples

```
duck <- dbi.catalog(chinook.duckdb)
Album <- duck$main$Album
Artist <- duck$main$Artist

reference.test(merge(Album, Artist, by = "ArtistId"))
```

<code>sql.join</code>	<i>Join dbi.tables</i>
-----------------------	------------------------

Description

A SQL-like join of two `dbi.tables` that share the same underlying `DBI` connection. All columns of both `dbi.tables` are returned.

Usage

```
sql.join(x, y, type = "inner", on = NULL, prefixes = c("x.", "y."))
```

Arguments

x, y	dbi.tables to join. x and y must share the same DBI connection.
type	a character string specifying the join type. Valid choices are "inner", "left", "right", "outer", and "cross".
on	a call specifying the join predicate. The symbols in on should be column names of x or column names of y, use prefixes as necessary.
prefixes	a 2-element character vector of distinct values. When x and y both have a column with the same name (e.g., common_name) then, when specifying the join predicate in on, use `prefixes[1]`common_name to refer to the common_name column in x and `prefixes[2]`common_name to refer to the common_name column in y. prefixes are also used to disambiguate the output column names.

Value

a dbi.table.

Examples

```
chinook <- dbi.catalog(chinook.duckdb)
Album <- chinook$main$Album
Artist <- chinook$main$Artist

sql.join(Album, Artist, type = "inner",
         on = Album.ArtistId == Artist.ArtistId,
         prefixes = c("Album.", "Artist."))
```

Index

- [.dbi.table (dbi.table-package), 2
- all.equal, 12
- as.data.table, 3, 4, 4
- as.dbi.table, 5
- attach, 8

- chinook.duckdb (example_databases), 9
- chinook.sqlite (example_databases), 9
- conflicts, 8
- csql, 3, 6

- data.table, 2, 4, 9
- dbConnect, 2, 5, 8, 9
- dbExecute, 7
- dbExecute, dbi.catalog, SQL-method, 6
- dbExecute, dbi.schema, SQL-method
(dbExecute, dbi.catalog, SQL-method),
6
- dbExecute, dbi.table, SQL-method
(dbExecute, dbi.catalog, SQL-method),
6
- dbFetch, 4
- dbGetInfo, 7
- dbGetInfo, dbi.catalog-method
(dbExecute, dbi.catalog, SQL-method),
6
- dbGetInfo, dbi.schema-method
(dbExecute, dbi.catalog, SQL-method),
6
- dbGetInfo, dbi.table
(dbExecute, dbi.catalog, SQL-method),
6
- dbGetInfo, dbi.table-method
(dbExecute, dbi.catalog, SQL-method),
6
- DBI, 12, 13
- dbi.attach, 7
- dbi.catalog, 2, 5, 7, 8
- dbi.table, 4-7, 9-13

- dbi.table (dbi.table-package), 2
- dbi.table-package, 2
- DBIConnection, 2, 9
- dbRemoveTable, 5
- dbSendStatement, 7
- dbSendStatement, dbi.table, missing-method
(dbExecute, dbi.catalog, SQL-method),
6

- environment, 8
- example_databases, 9

- Id, 2
- is.dbi.table (as.dbi.table), 5

- menu, 8
- merge, 9
- merge.data.frame, 10

- reference.test, 11

- search, 8
- SQL, 2, 7
- sql.join, 11, 12