

Package ‘rpostgis’

December 6, 2024

Version 1.6.0

Date 2024-12-05

Title R Interface to a 'PostGIS' Database

Description Provides an interface between R and 'PostGIS'-enabled 'PostgreSQL' databases to transparently transfer spatial data. Both vector (points, lines, polygons) and raster data are supported in read and write modes. Also provides convenience functions to execute common procedures in 'PostgreSQL/PostGIS'.

SystemRequirements 'PostgreSQL' with 'PostGIS' extension

Depends R (>= 3.3.0), RPostgreSQL, DBI (>= 0.5)

Imports methods, sf, stats, terra (>= 1.6.7), cli, lifecycle

Suggests RPostgres, testthat (>= 3.0.0), sp, raster

License GPL (>= 3)

URL <https://cidree.github.io/rpostgis/>,
<https://github.com/Cidree/rpostgis>

BugReports <https://github.com/Cidree/rpostgis/issues>

RoxygenNote 7.3.2

Encoding UTF-8

Config/testthat/edition 3

NeedsCompilation no

Author Adrian Cidre Gonzalez [aut, cre]
(<<https://orcid.org/0000-0002-3310-3052>>),
Mathieu Basille [aut] (<<https://orcid.org/0000-0001-9366-7127>>),
David Bucklin [aut]

Maintainer Adrian Cidre Gonzalez <adrian.cidre@gmail.com>

Repository CRAN

Date/Publication 2024-12-06 10:20:02 UTC

Contents

| | |
|----------------------------|-----------|
| dbAddKey | 2 |
| dbAsDate | 4 |
| dbColumn | 5 |
| dbComment | 6 |
| dbDrop | 7 |
| dbIndex | 8 |
| dbSchema | 9 |
| dbTableInfo | 10 |
| dbVacuum | 11 |
| dbWriteDataFrame | 12 |
| pgGetBoundary | 13 |
| pgGetGeom | 14 |
| pgGetRast | 16 |
| pgInsert | 18 |
| pgListGeom | 21 |
| pgMakePts | 21 |
| pgPostGIS | 23 |
| pgSRID | 24 |
| pgWriteGeom | 25 |
| pgWriteRast | 28 |
| Index | 31 |

| | |
|----------|-----------------|
| dbAddKey | <i>Add key.</i> |
|----------|-----------------|

Description

Add a primary or foreign key to a table column.

Usage

```
dbAddKey(
  conn,
  name,
  colname,
  type = c("primary", "foreign"),
  reference,
  colref,
  display = TRUE,
  exec = TRUE
)
```

Arguments

| | |
|-----------|---|
| conn | A connection object. |
| name | A character string, or a character vector, specifying a PostgreSQL table name. |
| colname | A character string specifying the name of the column to which the key will be assign; alternatively, a character vector specifying the name of the columns for keys spanning more than one column. |
| type | The type of the key, either "primary" or "foreign" |
| reference | A character string specifying a foreign table name to which the foreign key will be associated (ignored if type == "primary"). |
| colref | A character string specifying the name of the primary key in the foreign table to which the foreign key will be associated; alternatively, a character vector specifying the name of the columns for keys spanning more than one column (ignored if type == "primary"). |
| display | Logical. Whether to display the query (defaults to TRUE). |
| exec | Logical. Whether to execute the query (defaults to TRUE). |

Value

If exec = TRUE, returns (invisibly) TRUE if the key was successfully added.

Author(s)

Mathieu Basille <mathieu@basille.org>

See Also

The PostgreSQL documentation: <http://www.postgresql.org/docs/current/static/sql-altertable.html>

Examples

```
## Examples use a dummy connection from DBI package
conn <- DBI::ANSI()

## Primary key
dbAddKey(conn, name = c("sch1", "tbl1"), colname = "id1", exec = FALSE)

## Primary key using multiple columns
dbAddKey(conn, name = c("sch1", "tbl1"), colname = c("id1", "id2",
  "id3"), exec = FALSE)

## Foreign key
dbAddKey(conn, name = c("sch1", "tbl1"), colname = "id", type = "foreign",
  reference = c("sch2", "tbl2"), colref = "id", exec = FALSE)

## Foreign key using multiple columns
dbAddKey(conn, name = c("sch1", "tbl1"), colname = c("id1", "id2"),
  type = "foreign", reference = c("sch2", "tbl2"), colref = c("id3",
  "id4"), exec = FALSE)
```

| | |
|----------|-------------------------------|
| dbAsDate | <i>Converts to timestamp.</i> |
|----------|-------------------------------|

Description

Convert a date field to a timestamp with or without time zone.

Usage

```
dbAsDate(conn, name, date = "date", tz = NULL, display = TRUE, exec = TRUE)
```

Arguments

| | |
|---------|---|
| conn | A connection object. |
| name | A character string specifying a PostgreSQL table name. |
| date | A character string specifying the date field. |
| tz | A character string specifying the time zone, in "EST", "America/New_York", "EST5EDT", "-5". |
| display | Logical. Whether to display the query (defaults to TRUE). |
| exec | Logical. Whether to execute the query (defaults to TRUE). |

Value

If exec = TRUE, returns (invisibly) TRUE if the conversion was successful.

Author(s)

Mathieu Basille <mathieu@basille.org>

See Also

The PostgreSQL documentation: <http://www.postgresql.org/docs/current/static/datatype-datetime.html>

Examples

```
## Example uses a dummy connection from DBI package
conn <- DBI::ANSI()
dbAsDate(conn, name = c("schema", "table"), date = "date", tz = "GMT",
  exec = FALSE)
```

| | |
|----------|--------------------------------|
| dbColumn | <i>Add or remove a column.</i> |
|----------|--------------------------------|

Description

Add or remove a column to/from a table.

Usage

```
dbColumn(  
    conn,  
    name,  
    colname,  
    action = c("add", "drop"),  
    coltype = "integer",  
    cascade = FALSE,  
    display = TRUE,  
    exec = TRUE  
)
```

Arguments

| | |
|---------|--|
| conn | A connection object. |
| name | A character string specifying a PostgreSQL table name. |
| colname | A character string specifying the name of the column |
| action | A character string specifying if the column is to be added ("add", default) or removed ("drop"). |
| coltype | A character string indicating the type of the column, if action = "add". |
| cascade | Logical. Whether to drop foreign key constraints of other tables, if action = "drop". |
| display | Logical. Whether to display the query (defaults to TRUE). |
| exec | Logical. Whether to execute the query (defaults to TRUE). |

Value

If exec = TRUE, returns (invisibly) TRUE if the column was successfully added or removed.

Author(s)

Mathieu Basille <mathieu@basille.org>

See Also

The PostgreSQL documentation: <http://www.postgresql.org/docs/current/static/sql-altertable.html>

Examples

```
## examples use a dummy connection from DBI package
conn<-DBI::ANSI()
## Add an integer column
dbColumn(conn, name = c("schema", "table"), colname = "field", exec = FALSE)
## Drop a column (with CASCADE)
dbColumn(conn, name = c("schema", "table"), colname = "field", action = "drop",
          cascade = TRUE, exec = FALSE)
```

dbComment

Comment table/view/schema.

Description

Comment on a table, a view or a schema.

Usage

```
dbComment(
  conn,
  name,
  comment,
  type = c("table", "view", "schema"),
  display = TRUE,
  exec = TRUE
)
```

Arguments

| | |
|---------|--|
| conn | A connection object. |
| name | A character string specifying a PostgreSQL table, view or schema name. |
| comment | A character string specifying the comment. |
| type | The type of the object to comment, either "table", "view", or "schema" |
| display | Logical. Whether to display the query (defaults to TRUE). |
| exec | Logical. Whether to execute the query (defaults to TRUE). |

Value

If exec = TRUE, returns (invisibly) TRUE if the comment was successfully applied.

Author(s)

Mathieu Basille <mathieu@basille.org>

See Also

The PostgreSQL documentation: <http://www.postgresql.org/docs/current/static/sql-comment.html>

Examples

```
## examples use a dummy connection from DBI package
conn <- DBI::ANSI()
dbComment(conn, name = c("schema", "table"), comment = "Comment on a view.",
  type = "view", exec = FALSE)
dbComment(conn, name = "test_schema", comment = "Comment on a schema.", type = "schema",
  exec = FALSE)
```

dbDrop

Drop table/view/schema.

Description

Drop a table, a view or a schema.

Usage

```
dbDrop(
  conn,
  name,
  type = c("table", "schema", "view", "materialized view"),
  ifexists = FALSE,
  cascade = FALSE,
  display = TRUE,
  exec = TRUE
)
```

Arguments

| | |
|----------|---|
| conn | A connection object. |
| name | A character string specifying a PostgreSQL table, schema, or view name. |
| type | The type of the object to drop, either "table", "schema", "view", or "materialized view". |
| ifexists | Do not throw an error if the object does not exist. A notice is issued in this case. |
| cascade | Automatically drop objects that depend on the object (such as views). |
| display | Logical. Whether to display the query (defaults to TRUE). |
| exec | Logical. Whether to execute the query (defaults to TRUE). |

Value

If `exec = TRUE`, returns (invisibly) TRUE if the table/schema/view was successfully dropped.

Author(s)

Mathieu Basille <mathieu@basille.org>

See Also

The PostgreSQL documentation: <http://www.postgresql.org/docs/current/static/sql-droptable.html>, <http://www.postgresql.org/docs/current/static/sql-dropview.html>, <http://www.postgresql.org/docs/current/static/sql-dropschema.html>

Examples

```
## examples use a dummy connection from DBI package
conn <- DBI::ANSI()
dbDrop(conn, name = c("schema", "view_name"), type = "view", exec = FALSE)
dbDrop(conn, name = "test_schema", type = "schema", cascade = "TRUE", exec = FALSE)
```

dbIndex

Create an index.

Description

Defines a new index on a PostgreSQL table.

Usage

```
dbIndex(
  conn,
  name,
  colname,
  idxname,
  unique = FALSE,
  method = c("btree", "hash", "rtree", "gist"),
  display = TRUE,
  exec = TRUE
)
```

Arguments

| | |
|---------|--|
| conn | A connection object. |
| name | A character string specifying a PostgreSQL table name. |
| colname | A character string, or a character vector specifying the name of the column to which the key will be associated; alternatively, a character vector specifying the name of the columns to build the index. |
| idxname | A character string specifying the name of the index to be created. By default, this uses the name of the table (without the schema) and the name of the columns as follows: <table_name>_<column_names>_idx. |
| unique | Logical. Causes the system to check for duplicate values in the table when the index is created (if data already exist) and each time data is added. Attempts to insert or update data which would result in duplicate entries will generate an error. |

| | |
|---------|--|
| method | The name of the method to be used for the index. Choices are "btree", "hash", "rtree", and "gist". The default method is "btree", although "gist" should be the index of choice for PostGIS spatial types (geometry, geography, raster). |
| display | Logical. Whether to display the query (defaults to TRUE). |
| exec | Logical. Whether to execute the query (defaults to TRUE). |

Value

If exec = TRUE, returns (invisibly) TRUE if the index was successfully created.

Author(s)

Mathieu Basille <mathieu@basille.org>

See Also

The PostgreSQL documentation: <http://www.postgresql.org/docs/current/static/sql-createindex.html>; the PostGIS documentation for GiST indexes: http://postgis.net/docs/using_postgis_dbmanagement.html#id541286

Examples

```
## Examples use a dummy connection from DBI package
conn <- DBI::ANSI()

## GIST index
dbIndex(conn, name = c("sch", "tbl"), colname = "geom", method = "gist",
        exec = FALSE)

## Regular BTREE index on multiple columns
dbIndex(conn, name = c("sch", "tbl"), colname = c("col1", "col2",
        "col3"), exec = FALSE)
```

dbSchema

Check and create schema.

Description

Checks the existence, and if necessary, creates a schema.

Usage

```
dbSchema(conn, name, display = TRUE, exec = TRUE)
```

Arguments

| | |
|---------|--|
| conn | A connection object (required, even if exec = FALSE). |
| name | A character string specifying a PostgreSQL schema name. |
| display | Logical. Whether to display the query (defaults to TRUE). |
| exec | Logical. Whether to execute the query (defaults to TRUE). Note: if exec = FALSE, the function still checks the existence of the schema, but does not create it if it does not exist. |

Value

If exec = TRUE, returns (invisible) TRUE if the schema exists (whether it was already available or was just created).

Author(s)

Mathieu Basille <mathieu@basille.org>

See Also

The PostgreSQL documentation: <http://www.postgresql.org/docs/current/static/sql-createschema.html>

Examples

```
## Not run:
dbSchema(conn, name = "schema", exec = FALSE)

## End(Not run)
```

dbTableInfo

Get information about table columns.

Description

Get information about columns in a PostgreSQL table.

Usage

```
dbTableInfo(conn, name, allinfo = FALSE)
```

Arguments

| | |
|---------|---|
| conn | A connection object to a PostgreSQL database. |
| name | A character string specifying a PostgreSQL schema (if necessary), and table or view name (e.g., name = c("schema", "table")). |
| allinfo | Logical, Get all information on table? Default is column names, types, nullable, and maximum length of character columns. |

Value

data frame

Author(s)

David Bucklin <david.bucklin@gmail.com>

Examples

```
## Not run:
dbTableInfo(conn, c("schema", "table"))

## End(Not run)
```

dbVacuum

Vacuum.

Description

Performs a VACUUM (garbage-collect and optionally analyze) on a table.

Usage

```
dbVacuum(
  conn,
  name,
  full = FALSE,
  verbose = FALSE,
  analyze = TRUE,
  display = TRUE,
  exec = TRUE
)
```

Arguments

| | |
|---------|---|
| conn | A connection object. |
| name | A character string specifying a PostgreSQL table name. |
| full | Logical. Whether to perform a "full" vacuum, which can reclaim more space, but takes much longer and exclusively locks the table. |
| verbose | Logical. Whether to print a detailed vacuum activity report for each table. |
| analyze | Logical. Whether to update statistics used by the planner to determine the most efficient way to execute a query (default to TRUE). |
| display | Logical. Whether to display the query (defaults to TRUE). |
| exec | Logical. Whether to execute the query (defaults to TRUE). |

Value

If `exec = TRUE`, returns (invisibly) `TRUE` if query is successfully executed.

Author(s)

Mathieu Basille <mathieu@basille.org>

See Also

The PostgreSQL documentation: <http://www.postgresql.org/docs/current/static/sql-vacuum.html>

Examples

```
## examples use a dummy connection from DBI package
conn <- DBI::ANSI()
dbVacuum(conn, name = c("schema", "table"), full = TRUE, exec = FALSE)
```

dbWriteDataFrame *Write/read in data frame mode to/from database table.*

Description

Write `data.frame` or similar (e.g. `tibble`) to database table, with column definitions, row names, and a new integer primary key column. Read back into R with `dbReadDataFrame`, which recreates original data.

Usage

```
dbWriteDataFrame(conn, name, df, overwrite = FALSE, only.defs = FALSE)
```

```
dbReadDataFrame(conn, name, df = NULL)
```

Arguments

| | |
|------------------------|--|
| <code>conn</code> | A connection object to a PostgreSQL database |
| <code>name</code> | Character, schema and table of the PostgreSQL table |
| <code>df</code> | The data frame to write (for <code>dbReadDataFrame</code> , this allows to update an existing <code>data.frame</code> with definitions stored in the database) |
| <code>overwrite</code> | Logical; if <code>TRUE</code> , a new table (name) will overwrite the existing table (name) in the database. Note: overwriting a view must be done manually (e.g., with <code>dbDrop()</code>). |
| <code>only.defs</code> | Logical; if <code>TRUE</code> , only the table definitions will be written. |

Details

Writing in data frame mode is only for new database tables (or for overwriting an existing one). It will save all column names as they appear in R, along with column data types and attributes. This is done by adding metadata to a lookup table in the table's schema named ".R_df_defs" (will be created if not present). It also adds two fields with fixed names to the database table: ".R_rownames" (storing the row.names of the data frame), and ".db_pkid", which is a new integer primary key. Existing columns in the data.frame matching these names will be automatically changed.

The rpostgis database table read functions dbReadDataFrame and pgGetGeom will use the metadata created in data frame mode to recreate a data.frame in R, if it is available. Otherwise, it will be imported using default RPostgreSQL::dbGetQuery methods.

All spatial objects must be written with pgWriteGeom(). For more flexible writing of data.frames to the database (including all writing into existing database tables), use pgWriteGeom() with df.mode = FALSE.

Value

invisible TRUE for successful write with dbWriteDataFrame, data.frame for dbReadDataFrame

Author(s)

David Bucklin <david.bucklin@gmail.com>

Adrián Cidre González <adrian.cidre@gmail.com>

Examples

```
## Not run:
library(datasets)

## Write the mtcars data.frame to the database:
dbWriteDataFrame(conn, name = "mtcars_data", df = mtcars)

## Reads it back into a different object:
mtcars2 <- dbReadDataFrame(conn, name = "mtcars_data")

## Check equality:
all.equal(mtcars, mtcars2)
## Should return TRUE.

## End(Not run)
```

pgGetBoundary

Retrieve bounding envelope of geometries or rasters.

Description

Retrieve bounding envelope (rectangle) of all geometries or rasters in a PostGIS table as a sfc object.

Usage

```
pgGetBoundary(conn, name, geom = "geom", clauses = NULL, returnclass = "sf")
```

Arguments

| | |
|-------------|--|
| conn | A connection object to a PostgreSQL database |
| name | A character string specifying a PostgreSQL schema and table/view name holding the geometry (e.g., name = c("schema", "table")) |
| geom | A character string specifying the name of the geometry column in the table name (Default = "geom"). Note that for raster objects you will need to change the default value |
| clauses | character, additional SQL to append to modify select query from table. Must begin with an SQL clause (e.g., "WHERE ...", "ORDER BY ...", "LIMIT ..."); same usage as in pgGetGeom. |
| returnclass | 'sf' by default; 'terra' for SpatVector; or 'sp' for sp objects. |

Value

object of class sfc (list-column with geometries); SpatVector or sp object

Author(s)

David Bucklin <david.bucklin@gmail.com> and Adrian Cidre González <adrian.cidre@gmail.com>

Examples

```
## Not run:
pgGetBoundary(conn, c("schema", "polys"), geom = "geom")
pgGetBoundary(conn, c("schema", "rasters"), geom = "rast")

## End(Not run)
```

pgGetGeom

Load a PostGIS geometry from a PostgreSQL table/view/query into R.

Description

Retrieve geometries from a PostGIS table/view/query, and convert it to an R sf object.

Usage

```
pgGetGeom(
  conn,
  name,
  geom = "geom",
  gid = NULL,
```

```

    other.cols = TRUE,
    clauses = NULL,
    boundary = NULL,
    query = NULL,
    returnclass = "sf"
)

```

Arguments

| | |
|-------------|--|
| conn | A connection object to a PostgreSQL database |
| name | A character string specifying a PostgreSQL schema and table/view name holding the geometry (e.g., name = c("schema", "table")) |
| geom | The name of the geometry/(geography) column. (Default = "geom") |
| gid | Name of the column in name holding the IDs. Should be unique for each record to return. gid=NULL (default) automatically creates a new unique ID for each row in the sf object. |
| other.cols | Names of specific columns in the table to retrieve, in a character vector (e.g. other.cols=c("col1", "col2").) The default (other.cols = TRUE) is to attach all columns. Setting other.cols=FALSE will return a Spatial-only object without attributes (no data frame). |
| clauses | character, additional SQL to append to modify select query from table. Must begin with an SQL clause (e.g., "WHERE ...", "ORDER BY ...", "LIMIT ..."); see below for examples. |
| boundary | sf, SpatVector or sp object; or numeric. If a spatial object is provided, its bounding box will be used to select geometries to import. Alternatively, a numeric vector (c([top], [bottom], [right], [left])) indicating the projection-specific limits with which to subset the spatial data. If not value is provided, the default boundary = NULL will not apply any boundary subset. |
| query | character, a full SQL query including a geometry column. For use with query mode only (see details). |
| returnclass | 'sf' by default; 'terra' for SpatVector; or 'sp' for sp objects. |

Details

The features of the table to retrieve must have the same geometry type. The query mode version of pgGetGeom allows the user to enter a complete SQL query (query) that returns a Geometry column, and save the query as a new view (name) if desired. If (name) is not specified, a temporary view with name ".rpostgis_TEMPview" is used only within the function execution. In this mode, the other arguments can be used normally to modify the Spatial* object returned from the query.

Value

sf, SpatVector or sp object

Author(s)

David Bucklin <david.bucklin@gmail.com>
 Mathieu Basille <mathieu@basille.org>
 Adrián Cidre González <adrian.cidre@gmail.com>

Examples

```
## Not run:
## Retrieve a sf with all data from table
## 'schema.tablename', with geometry in the column 'geom'
pgGetGeom(conn, c("schema", "tablename"))
## Return a sf with columns c1 & c2 as data
pgGetGeom(conn, c("schema", "tablename"), other.cols = c("c1","c2"))
## Return a spatial-only (no data frame),
## retaining id from table as rownames
pgGetGeom(conn, c("schema", "tablename"), gid = "table_id",
  other.cols = FALSE)
## Return a spatial-only (no data frame),
## retaining id from table as rownames and with a subset of the data
pgGetGeom(conn, c("schema", "roads"), geom = "roadgeom", gid = "road_ID",
  other.cols = FALSE, clauses = "WHERE road_type = 'highway'")
## Query mode
pgGetGeom(conn, query = "SELECT r.gid as id, ST_Buffer(r.geom, 100) as geom
                        FROM
                        schema.roads r,
                        schema.adm_boundaries b
                        WHERE
                        ST_Intersects(r.geom, b.geom);")

## End(Not run)
```

pgGetRast

Load raster from PostGIS database into R.

Description

Retrieve rasters from a PostGIS table into a terra SpatRaster object

Usage

```
pgGetRast(
  conn,
  name,
  rast = "rast",
  bands = 1,
  boundary = NULL,
  clauses = NULL,
  returnclass = "terra",
```



```

    progress = TRUE
  )

```

Arguments

| | |
|-------------|--|
| conn | A connection object to a PostgreSQL database |
| name | A character string specifying a PostgreSQL schema and table/view name holding the geometry (e.g., name = c("schema", "table")) |
| rast | Name of the column in name holding the raster object. Defaults to "rast". |
| bands | Index number(s) for the band(s) to retrieve (defaults to 1). The special case (bands = TRUE) returns all bands in the raster. See also 'Details' |
| boundary | sf object, SpatVector object, or numeric. If a spatial object is provided, its bounding box will be used to select the part of the raster to import. Alternatively, a numeric vector (c([top], [bottom], [right], [left])) indicating the projection-specific limits with which to clip the raster. If not value is provided, the default boundary = NULL will return the full raster. |
| clauses | character, optional SQL to append to modify select query from table. Must begin with 'WHERE'. |
| returnclass | 'terra' by default; or 'raster' for raster objects. |
| progress | whether to show a progress bar (TRUE by default). The progress bar mark the progress of reading bands from the database. |

Details

Since version 1.5.0, this function retrieve SpatRaster objects from terra package by default. The argument returnclass can be used to return raster objects instead.

The argument bands can take as argument:

- The index of the desirable band (e.g. bands = 2 will fetch the second band of the raster).
- More than one index for several bands (e.g. bands = c(2,4) will return a SpatRaster with two bands).
- All bands in the raster (bands = TRUE).

Value

SpatRaster; raster; or RasterStack object

Author(s)

David Bucklin <david.bucklin@gmail.com> and Adrián Cidre González <adrian.cidre@gmail.com>

Examples

```

## Not run:
pgGetRast(conn, c("schema", "tablename"))
pgGetRast(conn, c("schema", "DEM"), boundary = c(55,
  50, 17, 12))

## End(Not run)

```

 pgInsert

Inserts data into a PostgreSQL table.

Description

[Deprecated]

This function has been deprecated in favour of `pgWriteGeom()` and will be removed in a future release.

This function takes a take an R `sp` object (`Spatial*` or `Spatial*DataFrame`), or a regular `data.frame`, and performs the database insert (and table creation, when the table does not exist) on the database.

If `new.id` is specified, a new sequential integer field is added to the data frame for insert. For `Spatial*`-only objects (no data frame), a new ID column is created by default with name "gid".

This function will use `st_as_text` for geography types, and `st_as_binary` for geometry types.

In the event of function or database error, the database uses ROLLBACK to revert to the previous state.

If the user specifies `return.pgi = TRUE`, and data preparation is successful, the function will return a `pgi` object (see next paragraph), regardless of whether the insert was successful or not. This object can be useful for debugging, or re-used as the `data.obj` in `pgInsert`; (e.g., when data preparation is slow, and the exact same data needs to be inserted into tables in two separate tables or databases). If `return.pgi = FALSE` (default), the function will return TRUE for successful insert and FALSE for failed inserts.

Use this function with `df.mode = TRUE` to save data frames from `Spatial*`-class objects to the database in "data frame mode". Along with normal `dbwriteDataFrame` operation, the `proj4string` of the spatial data will also be saved, and re-attached to the data when using `pgGetGeom` to import the data. Note that other attributes of `Spatial*` objects are **not** saved (e.g., `coords.nrs`, which is used to specify the column index of x/y columns in `SpatialPoints*`).

`pgi` objects are a list containing four character strings: (1) `in.table`, the table name which will be created or inserted into (2) `db.new.table`, the SQL statement to create the new table, (3) `db.cols.insert`, a character string of the database column names to insert into, and (4) `insert.data`, a character string of the data to insert.

Usage

```
pgInsert(
  conn,
  name,
  data.obj,
  geom = "geom",
  df.mode = FALSE,
  partial.match = FALSE,
  overwrite = FALSE,
  new.id = NULL,
  row.names = FALSE,
  upsert.using = NULL,
```

```

    alter.names = FALSE,
    encoding = NULL,
    return.pgi = FALSE,
    df.geom = NULL,
    geog = FALSE
  )

```

Arguments

| | |
|----------------------------|---|
| <code>conn</code> | A connection object to a PostgreSQL database |
| <code>name</code> | A character string specifying a PostgreSQL schema and table name (e.g., <code>name = c("schema", "table")</code>). If not already existing, the table will be created. If the table already exists, the function will check if all R data frame columns match database columns, and if so, do the insert. If not, the insert will be aborted. The argument <code>partial.match</code> allows for inserts with only partial matches of data frame and database column names, and <code>overwrite</code> allows for overwriting the existing database table. |
| <code>data.obj</code> | A <code>Spatial*</code> or <code>Spatial*DataFrame</code> , or <code>data.frame</code> |
| <code>geom</code> | character string. For <code>Spatial*</code> datasets, the name of geometry/(geography) column in the database table. (existing or to be created; defaults to "geom"). The special name "geog" will automatically set <code>geog</code> to TRUE. |
| <code>df.mode</code> | Logical; Whether to write the (Spatial) data frame in data frame mode (preserving data frame column attributes and <code>row.names</code>). A new table must be created with this mode (or <code>overwrite</code> set to TRUE), and the <code>row.names</code> , <code>alter.names</code> , and <code>new.id</code> arguments will be ignored (see dbWriteDataFrame for more information). |
| <code>partial.match</code> | Logical; allow insert on partial column matches between data frame and database table. If TRUE, columns in R data frame will be compared with the existing database table name. Columns in the data frame that exactly match the database table will be inserted into the database table. |
| <code>overwrite</code> | Logical; if true, a new table (name) will overwrite the existing table (name) in the database. Note: overwriting a view must be done manually (e.g., with dbDrop). |
| <code>new.id</code> | Character, name of a new sequential integer ID column to be added to the table for insert (for spatial objects without data frames, this column is created even if left NULL and defaults to the name "gid"). If <code>partial.match = TRUE</code> and the column does not exist in the database table, it will be discarded. |
| <code>row.names</code> | Whether to add the data frame row names to the database table. Column name will be <code>'R_rownames'</code> . |
| <code>upsert.using</code> | Character, name of the column(s) in the database table or constraint name used to identify already-existing rows in the table, which will be updated rather than inserted. The column(s) must have a unique constraint already created in the database table (e.g., a primary key). Requires PostgreSQL 9.5+. |
| <code>alter.names</code> | Logical, whether to make database column names DB-compliant (remove special characters/capitalization). Default is FALSE. (This must be set to FALSE to match with non-standard names in an existing database table.) |

| | |
|------------|--|
| encoding | Character vector of length 2, containing the from/to encodings for the data (as in the function <code>base::iconv()</code>). For example, if the dataset contain certain latin characters (e.g., accent marks), and the database is in UTF-8, use <code>encoding = c("latin1", "UTF-8")</code> . Left NULL, no conversion will be done. |
| return.pgi | Whether to return a formatted list of insert parameters (i.e., a <code>pgi</code> object; see function details.) |
| df.geom | Character vector, name of a character column in an R data.frame storing PostGIS geometries, this argument can be used to insert a geometry stored as character type in a data.frame (do not use with Spatial* data types). If only the column name is used (e.g., <code>df.geom = "geom"</code>), the column type will be a generic (GEOMETRY); use a two-length character vector (e.g., <code>df.geom = c("geom", "(POINT,4326)")</code>) to also specify a specific PostGIS geometry type and SRID for the column. Only recommended for for new tables/overwrites, since this method will change the existing column type. |
| geog | Logical; Whether to write the spatial data as a PostGIS 'GEOGRAPHY' type. By default, FALSE, unless <code>geom = "geog"</code> . |

Value

Returns TRUE if the insertion was successful, FALSE if failed, or a `pgi` object if specified.

Author(s)

David Bucklin <david.bucklin@gmail.com>

Examples

```
## Not run:
library(sp)
data(meuse)
coords <- SpatialPoints(meuse[, c("x", "y")])
spdf <- SpatialPointsDataFrame(coords, meuse)

## Insert data in new database table
pgInsert(conn, name = c("public", "meuse_data"), data.obj = spdf)

## The same command will insert into already created table (if all R
## columns match)
pgInsert(conn, name = c("public", "meuse_data"), data.obj = spdf)

## If not all database columns match, need to use partial.match = TRUE,
## where non-matching columns are not inserted
colnames(spdf@data)[4] <- "cu"
pgInsert(conn, name = c("public", "meuse_data"), data.obj = spdf,
         partial.match = TRUE)

## End(Not run)
```

| | |
|------------|--------------------------------|
| pgListGeom | <i>List geometries/rasters</i> |
|------------|--------------------------------|

Description

List all geometry/(geography) or raster columns available in a PostGIS database.

Usage

```
pgListGeom(conn, geog = TRUE)
```

```
pgListRast(conn)
```

Arguments

| | |
|------|---|
| conn | A PostgreSQL database connection. |
| geog | Logical. For pgListGeom, whether to include PostGIS geography-type columns stored in the database |

Value

If exec = TRUE, a data frame with schema, table, geometry/(geography) or raster (for pgListRast) column, and geometry/(geography) type.

Author(s)

David Bucklin <david.bucklin@gmail.com>

Examples

```
## Not run:  
pgListGeom(conn)  
  
pgListRast(conn)  
  
## End(Not run)
```

| | |
|-----------|--|
| pgMakePts | <i>Add a POINT or LINESTRING geometry field.</i> |
|-----------|--|

Description

Add a new POINT or LINESTRING geometry field.

Usage

```
pgMakePts(
  conn,
  name,
  colname = "geom",
  x = "x",
  y = "y",
  srid,
  index = TRUE,
  display = TRUE,
  exec = TRUE
)
```

```
pgMakeStp(
  conn,
  name,
  colname = "geom",
  x = "x",
  y = "y",
  dx = "dx",
  dy = "dy",
  srid,
  index = TRUE,
  display = TRUE,
  exec = TRUE
)
```

Arguments

| | |
|---------|--|
| conn | A connection object. |
| name | A character string specifying a PostgreSQL schema and table name (e.g., name = c("schema", "table")) |
| colname | A character string specifying the name of the new geometry column. |
| x | The name of the x/longitude field. |
| y | The name of the y/latitude field. |
| srid | A valid SRID for the new geometry. |
| index | Logical. Whether to create an index on the new geometry. |
| display | Logical. Whether to display the query (defaults to TRUE). |
| exec | Logical. Whether to execute the query (defaults to TRUE). |
| dx | The name of the dx field (i.e. increment in x direction). |
| dy | The name of the dy field (i.e. increment in y direction). |

Value

If exec = TRUE, returns TRUE if the geometry field was successfully created.

Author(s)

Mathieu Basille <mathieu@basille.org> and Adrián Cidre González <adrian.cidre@gmail.com>

See Also

The PostGIS documentation for ST_MakePoint: http://postgis.net/docs/ST_MakePoint.html, and for ST_MakeLine: http://postgis.net/docs/ST_MakeLine.html, which are the main functions of the call.

Examples

```
## Examples use a dummy connection from DBI package
conn <- DBI::ANSI()

## Create a new POINT field called 'pts_geom'
pgMakePts(conn, name = c("schema", "table"), colname = "pts_geom",
  x = "longitude", y = "latitude", srid = 4326, exec = FALSE)

## Create a new LINESTRING field called 'stp_geom'
pgMakeStp(conn, name = c("schema", "table"), colname = "stp_geom",
  x = "longitude", y = "latitude", dx = "xdiff", dy = "ydiff",
  srid = 4326, exec = FALSE)
```

pgPostGIS

Check and create PostGIS extension.

Description

The function checks for the availability of the PostGIS extension, and if it is available, but not installed, install it. Additionally, can also install Topology, Tiger Geocoder, SFCGAL and Raster extensions.

Usage

```
pgPostGIS(
  conn,
  topology = FALSE,
  tiger = FALSE,
  sfcgal = FALSE,
  raster = FALSE,
  display = TRUE,
  exec = TRUE
)
```

Arguments

| | |
|----------|--|
| conn | A connection object (required, even if exec = FALSE). |
| topology | Logical. Whether to check/install the Topology extension. |
| tiger | Logical. Whether to check/install the Tiger Geocoder extension. Will also install extensions "fuzzystrmatch", "address_standardizer", and "address_standardizer_data_us" if all are available. |
| sfcgal | Logical. Whether to check/install the SFCGAL extension. |
| raster | Logical. Whether to check/install the Raster extension |
| display | Logical. Whether to display the query (defaults to TRUE). |
| exec | Logical. Whether to execute the query (defaults to TRUE). |

Value

If exec = TRUE, returns (invisibly) TRUE if PostGIS is installed.

Author(s)

Mathieu Basille <mathieu@basille.org> and Adrián Cidre González <adrian.cidre@gmail.com>

Examples

```
## 'exec = FALSE' does not install any extension, but nevertheless
## check for available and installed extensions:
## Not run:
  pgPostGIS(con, topology = TRUE, tiger = TRUE, sfcgal = TRUE,
            exec = FALSE)

## End(Not run)
```

pgSRID

Find (or create) PostGIS SRID based on CRS object.

Description

This function takes `sf::st_crs()`-class object and a PostgreSQL database connection (with PostGIS extension), and returns the matching SRID(s) for that CRS. If a match is not found, a new entry can be created in the PostgreSQL `spatial_ref_sys` table using the parameters specified by the CRS. New entries will be created with `auth_name = 'rpostgis_custom'`, with the default value being the next open value between 880001-889999 (a different SRID value can be entered if desired.)

Usage

```
pgSRID(conn, crs, create.srid = FALSE, new.srid = NULL)
```


Arguments

| | |
|-------------|--|
| conn | A connection object to a PostgreSQL database. |
| crs | crs object, created through a call to <code>sf::st_crs()</code> . |
| create.srid | Logical. If no matching SRID is found, should a new SRID be created? User must have write access on <code>spatial_ref_sys</code> table. |
| new.srid | Integer. Optional SRID to give to a newly created SRID. If left NULL (default), the next open value of <code>srid</code> in <code>spatial_ref_sys</code> between 880001 and 889999 will be used. |

Value

SRID code (integer).

Author(s)

David Bucklin <david.bucklin@gmail.com> and Adrián Cidre González <adrian.cidre@gmail.com>

Examples

```
## Not run:
drv <- dbDriver("PostgreSQL")
conn <- dbConnect(drv, dbname = "dbname", host = "host", port = "5432",
  user = "user", password = "password")
(crs <- sf::st_crs("+proj=longlat"))
pgSRID(conn, crs)
(crs2 <- sf::st_crs(paste("+proj=stere", "+lat_0=52.15616055555555 +lon_0=5.387638888888889",
  "+k=0.999908 +x_0=155000 +y_0=463000", "+ellps=bessel",
  "+towgs84=565.237,50.0087,465.658,-0.406857,0.350733,-1.87035,4.0812",
  "+units=m")))
pgSRID(conn, crs2, create.srid = TRUE)

## End(Not run)
```

pgWriteGeom

Inserts data into a PostgreSQL table.

Description

This function takes a take an R `sf`, a `SpatVector` or `sp` object (`Spatial*` or `Spatial*DataFrame`); or a regular `data.frame`, and performs the database insert (and table creation, when the table does not exist) on the database.

Usage

```

pgWriteGeom(
  conn,
  name,
  data.obj,
  geom = "geom",
  df.mode = FALSE,
  partial.match = FALSE,
  overwrite = FALSE,
  new.id = NULL,
  row.names = FALSE,
  upsert.using = NULL,
  alter.names = FALSE,
  encoding = NULL,
  return.pgi = FALSE,
  df.geom = NULL,
  geog = FALSE
)

## S3 method for class 'pgi'
print(x, ...)

```

Arguments

| | |
|----------------------------|---|
| <code>conn</code> | A connection object to a PostgreSQL database |
| <code>name</code> | A character string specifying a PostgreSQL schema and table name (e.g., <code>name = c("schema", "table")</code>). If not already existing, the table will be created. If the table already exists, the function will check if all R data frame columns match database columns, and if so, do the insert. If not, the insert will be aborted. The argument <code>partial.match</code> allows for inserts with only partial matches of data frame and database column names, and <code>overwrite</code> allows for overwriting the existing database table. |
| <code>data.obj</code> | A <code>sf</code> , <code>SpatVector</code> , <code>sp-class</code> , or <code>data.frame</code> |
| <code>geom</code> | character string. For <code>Spatial*</code> datasets, the name of geometry/(geography) column in the database table. (existing or to be created; defaults to "geom"). The special name "geog" will automatically set <code>geog</code> to TRUE. |
| <code>df.mode</code> | Logical; Whether to write the (Spatial) data frame in data frame mode (preserving data frame column attributes and <code>row.names</code>). A new table must be created with this mode (or <code>overwrite</code> set to TRUE), and the <code>row.names</code> , <code>alter.names</code> , and <code>new.id</code> arguments will be ignored (see dbWriteDataFrame for more information). |
| <code>partial.match</code> | Logical; allow insert on partial column matches between data frame and database table. If TRUE, columns in R data frame will be compared with the existing database table name. Columns in the data frame that exactly match the database table will be inserted into the database table. |
| <code>overwrite</code> | Logical; if true, a new table (name) will overwrite the existing table (name) in the database. Note: overwriting a view must be done manually (e.g., with dbDrop). |

| | |
|---------------------------|---|
| <code>new.id</code> | Character, name of a new sequential integer ID column to be added to the table for insert (for spatial objects without data frames, this column is created even if left NULL and defaults to the name "gid"). If <code>partial.match = TRUE</code> and the column does not exist in the database table, it will be discarded. |
| <code>row.names</code> | Whether to add the data frame row names to the database table. Column name will be <code>'R_rownames'</code> . |
| <code>upsert.using</code> | Character, name of the column(s) in the database table or constraint name used to identify already-existing rows in the table, which will be updated rather than inserted. The column(s) must have a unique constraint already created in the database table (e.g., a primary key). Requires PostgreSQL 9.5+. |
| <code>alter.names</code> | Logical, whether to make database column names DB-compliant (remove special characters/capitalization). Default is FALSE. (This must be set to FALSE to match with non-standard names in an existing database table.) |
| <code>encoding</code> | Character vector of length 2, containing the from/to encodings for the data (as in the function <code>iconv</code>). For example, if the dataset contain certain latin characters (e.g., accent marks), and the database is in UTF-8, use <code>encoding = c("latin1", "UTF-8")</code> . Left NULL, no conversion will be done. |
| <code>return.pgi</code> | Whether to return a formatted list of insert parameters (i.e., a <code>pgi</code> object; see function details.) |
| <code>df.geom</code> | Character vector, name of a character column in an R data.frame storing PostGIS geometries, this argument can be used to insert a geometry stored as character type in a data.frame (do not use with Spatial* data types). If only the column name is used (e.g., <code>df.geom = "geom"</code>), the column type will be a generic (GEOMETRY); use a two-length character vector (e.g., <code>df.geom = c("geom", "(POINT, 4326)")</code>) to also specify a specific PostGIS geometry type and SRID for the column. Only recommended for for new tables/overwrites, since this method will change the existing column type. |
| <code>geog</code> | Logical; Whether to write the spatial data as a PostGIS 'GEOGRAPHY' type. By default, FALSE, unless <code>geom = "geog"</code> . |
| <code>x</code> | A list of class <code>pgi</code> |
| <code>...</code> | Further arguments not used. |

Details

If `new.id` is specified, a new sequential integer field is added to the data frame for insert. For spatial-only objects (no data frame), a new ID column is created by default with name "gid".

This function will use `sf::st_as_text()` for geography types, and `sf::st_as_binary()` for geometry types.

In the event of function or database error, the database uses ROLLBACK to revert to the previous state.

If the user specifies `return.pgi = TRUE`, and data preparation is successful, the function will return a `pgi` object (see next paragraph), regardless of whether the insert was successful or not. This object can be useful for debugging, or re-used as the `data.obj` in `pgWriteGeom`; (e.g., when data preparation is slow, and the exact same data needs to be inserted into tables in two separate tables

or databases). If `return.pgi = FALSE` (default), the function will return `TRUE` for successful insert and `FALSE` for failed inserts.

Use this function with `df.mode = TRUE` to save data frames from spatial-class objects to the database in "data frame mode". Along with normal `dbwriteDataFrame` operation, the `proj4string` of the spatial data will also be saved, and re-attached to the data when using `pgGetGeom` to import the data. Note that other attributes of spatial objects are **not** saved (e.g., `coords.nrs`, which is used to specify the column index of x/y columns in `*POINT` and `SpatialPoints*`).

`pgi` objects are a list containing four character strings: (1) `in.table`, the table name which will be created or inserted into (2) `db.new.table`, the SQL statement to create the new table, (3) `db.cols.insert`, a character string of the database column names to insert into, and (4) `insert.data`, a character string of the data to insert.

Value

Returns `TRUE` if the insertion was successful, `FALSE` if failed, or a `pgi` object if specified.

Author(s)

David Bucklin <david.bucklin@gmail.com> and Adrián Cidre González <adrian.cidre@gmail.com>

Examples

```
## Not run:
library(sf)
pts <- st_sf(a = 1:2, geom = st_sfc(st_point(0:1), st_point(1:2)), crs = 4326)

## Insert data in new database table
pgWriteGeom(conn, name = c("public", "my_pts"), data.obj = pts)

## The same command will insert into already created table (if all R
## columns match)
pgWriteGeom(conn, name = c("public", "my_pts"), data.obj = pts)

## If not all database columns match, need to use partial.match = TRUE,
## where non-matching columns are not inserted
names(pts)[1] <- "b"
pgWriteGeom(conn, name = c("public", "my_pts"), data.obj = pts,
  partial.match = TRUE)

## End(Not run)
```

pgWriteRast

Write raster to PostGIS database table.

Description

Sends R raster to a PostGIS database table.

Usage

```
pgWriteRast(
  conn,
  name,
  raster,
  bit.depth = NULL,
  blocks = NULL,
  constraints = TRUE,
  overwrite = FALSE,
  append = FALSE
)
```

Arguments

| | |
|-------------|--|
| conn | A connection object to a PostgreSQL database. |
| name | A character string specifying a PostgreSQL schema in the database (if necessary) and table name to hold the raster (e.g., name = c("schema", "table")). |
| raster | An terra SpatRaster; objects from the raster package (RasterLayer, RasterBrick, or RasterStack); a SpatialGrid* or SpatialPixels* from sp package. |
| bit.depth | The bit depth of the raster. Will be set to 32-bit (unsigned int, signed int, or float, depending on the data) if left null, but can be specified (as character) as one of the PostGIS pixel types (see http://postgis.net/docs/RT_ST_BandPixelType.html). |
| blocks | Optional desired number of blocks (tiles) to split the raster into in the resulting PostGIS table. This should be specified as a one or two-length (columns, rows) integer vector. See also 'Details'. |
| constraints | Whether to create constraints from raster data. Recommended to leave TRUE unless applying constraints manually (see http://postgis.net/docs/RT_AddRasterConstraints.html). Note that constraint notices may print to the console, depending on the PostgreSQL server settings. |
| overwrite | Whether to overwrite the existing table (name). |
| append | Whether to append to the existing table (name). |

Details

SpatRaster band names will be stored in an array in the column "band_names", which will be restored in R when imported with the function `pgGetRast()`.

Rasters from the sp and raster packages are converted to terra objects prior to insert.

If `blocks = NULL`, the number of block will vary by raster size, with a default value of 100 copies of the data in the memory at any point in time. If a specified number of blocks is desired, set blocks to a one or two-length integer vector. Note that fewer, larger blocks generally results in faster write times.

Value

TRUE (invisibly) for successful import.

Author(s)

David Bucklin <david.bucklin@gmail.com> and Adrián Cidre González <adrian.cidre@gmail.com>

See Also

Function follows process from http://postgis.net/docs/using_raster_dataman.html#RT_Creating_Rasters.

Examples

```
## Not run:
pgWriteRast(conn, c("schema", "tablename"), raster_name)

# basic test
r <- terra::rast(nrows=180, ncols=360, xmin=-180, xmax=180,
  ymin=-90, ymax=90, vals=1)
pgWriteRast(conn, c("schema", "test"), raster = r,
  bit.depth = "2BUI", overwrite = TRUE)

## End(Not run)
```

Index

`base::iconv()`, 20

`dbAddKey`, 2
`dbAsDate`, 4
`dbColumn`, 5
`dbComment`, 6
`dbDrop`, 7, 19, 26
`dbDrop()`, 12
`dbIndex`, 8
`dbReadDataFrame (dbWriteDataFrame)`, 12
`dbReadDF (dbWriteDataFrame)`, 12
`dbSchema`, 9
`dbTableInfo`, 10
`dbVacuum`, 11
`dbWriteDataFrame`, 12, 19, 26
`dbWriteDF (dbWriteDataFrame)`, 12

`iconv`, 27

`pgGetBoundary`, 13
`pgGetGeom`, 14
`pgGetRast`, 16
`pgGetRast()`, 29
`pgInsert`, 18
`pgListGeom`, 21
`pgListRast (pgListGeom)`, 21
`pgMakePts`, 21
`pgMakeStp (pgMakePts)`, 21
`pgPostGIS`, 23
`pgSRID`, 24
`pgWriteGeom`, 25
`pgWriteGeom()`, 13, 18
`pgWriteRast`, 28
`print.pgi (pgWriteGeom)`, 25

`sf::st_as_binary()`, 27
`sf::st_as_text()`, 27
`sf::st_crs()`, 24, 25
`st_as_binary`, 18
`st_as_text`, 18