

Package ‘RJDBC’

August 17, 2023

Version 0.3-0

Title Provides Access to Databases Through the JDBC Interface

Author Simon Urbanek <Simon.Urbanek@r-project.org>

Maintainer Simon Urbanek <Simon.Urbanek@r-project.org>

Depends methods, DBI, rJava (>= 0.4-15), R (>= 2.4.0)

Description The RJDBC package is an implementation of R's DBI interface using JDBC as a back-end. This allows R to connect to any DBMS that has a JDBC driver.

License MIT + file LICENSE

URL <http://www.rforge.net/RJDBC/>

NeedsCompilation no

R topics documented:

dbOption	1
dbSetTypeMaps	2
JDBC	3
JDBCConnection-class	5
JDBCConnection-methods	6
JDBCDriver-class	7
JDBCDriver-methods	8
JDBCResult-class	9
JDBCResult-methods	9

Index	11
--------------	-----------

dbOption	<i>Retrieves DMBS options.</i>
----------	--------------------------------

Description

dbOption is the equivalent of `getOption`, but for options related JDBC.

At each level (driver and connection), you can set options that will act as defaults for the DBI methods. Options are inherited, so setting an option in a call to `JDBC` will automatically propagate that option to the `JDBCConnection` object created by `dbConnect`. You can override an option by specifying it downstream, so, setting `dbOptions=` parameter in `dbConnect` will replace the options from the driver with the same name, but will keep all others.

Usage

```
dbOption(dbo, name, default)
```

Arguments

dbo	JDBC database object which the options relate to.
name	string, name of the option (must not be an empty string).
default	default value to return in case the option is not set.

Value

Returns either the value of the option (if set) or the ‘default’ value otherwise.

Author(s)

Simon Urbanek

See Also

[JDBC](#), [dbConnect](#)

dbSetTypeMaps *Manage type mapping functions.*

Description

Type mapping functions are functions that perform conversions between types as they are received from JDBC to R objects. For any JDBC type it is possible to define a function which is called before returning results from `fetch()`.

`dbSetTypeMaps()` function sets, removes or lists the mapping functions.

Usage

```
dbSetTypeMaps(...)
```

Arguments

... one or more arguments of the form `TYPE=function(x)` ... (set a new conversion function) or `TYPE=NULL` (remove conversion). If empty, no changes are made.

Details

The type names correspond to the column type names as defined by JDBC (see `java.sql.Types` Java documentation). For example, if you want to convert `NUMERIC` columns to numeric values in R (regardless of precision)

Value

Named list of currently registered type mapping functions. If ... is empty then visibly, otherwise invisibly.

Examples

```
str(dbSetTypeMaps())

## use bit64::as.integer64 (if available) to convert BIGINT to integer64
if (is.environment(bit64 <- tryCatch(getNamespace("bit64"), error=function(e) FALSE)))
  dbSetTypeMaps(BIGINT=bit64$as.integer64)
```

JDBC

JDBC engine

Description

JDBC creates a new DBI driver that can be used to start JDBC connections.

`findDrivers` attempts to find and load all JDBC 4 drivers on the class path.

Usage

```
JDBC (driverClass = "", classPath = "", identifier.quote = NA, ...)
findDrivers(classPath = "", service = "java.sql.Driver", loader = NULL)
```

Arguments

<code>driverClass</code>	name of the Java class of the JDBC driver to load or a one of Java driver instances from <code>findDrivers</code> .
<code>classPath</code>	character vector, class path that needs to be appended in order to load the desired JDBC driver. Usually it is the path to the JAR file containing the driver and all necessary dependencies. It can be a vector and all paths are expanded.
<code>identifier.quote</code>	character to use for quoting identifiers in automatically generated SQL statements or <code>NA</code> if the back-end doesn't support quoted identifiers. See details section below.
<code>service</code>	string, name of the services class (for JDBC always "java.sql.Driver")
<code>loader</code>	Java class loader to use during the look-up or <code>NULL</code> for the default one
<code>...</code>	optional list, with options to associate with this driver. See also dbOption .

Details

JDBC function has two purposes. One is to initialize the Java VM and load a Java JDBC driver (not to be confused with the `JDBCdriver` R object which is actually a DBI driver). The second purpose is to create a proxy R object which can be used to call `dbConnect` which actually creates a connection.

JDBC requires a JDBC driver for a database-backend to be loaded. Usually a JDBC driver is supplied in a Java Archive (jar) file. The path to such a file can be specified in `classPath`. The driver itself has a Java class name that is used to load the driver (for example the MySQL driver uses `com.mysql.jdbc.Driver`), this has to be specified in `driverClass`.

Modern drivers (those supporting JDBC 4) may use Java Service Provider interface for discovery and those can be found using the `findDrivers()` function which returns a list of drivers. You can pass any of the returned elements as `classDriver`. Note that the discovery is dynamic, so you can use `rJava::.jaddClassPath(...)` to add new locations in which Java will look

for driver JAR files. However, only drivers providing JSP metadata in their JAR files can be found. JSP was introduced in Java 1.6 so `findDrivers()` only works on Java 1.6 or higher.

There are currently three different ways to specify drivers:

1. `dbConnect(JDBC("my.Driver.Class", "driver.jar"), ...)` is the most explicit way where the specified driver class is used and expected to be found on the class path. This always works, but the user has to know the full name of the driver class.
2. `dbConnect(JDBC(, "driver.jar"), ...)` omits the driver class which means JDBC will try to find the driver using the `DriverManager`. This *only* works if the JVM has been loaded with the driver when initialized, so this method is discouraged as it is in general very unreliable. The `DriverManager` never updates the list of drivers, so once your driver is not found, there is nothing you can do about it.
3. `dbConnect(JDBC(findDrivers("driver.jar")[[1]]), ...)` uses `findDrivers()` (see details above) to find all available drivers and then passes the needed driver (in this example the first one) to `JDBC()`. You don't need to repeat the class path in this case as it is already set by `findDrivers()`. It is best to look at the output to see which drivers have been found, but if you pass the list, the first driver is used. Note that if you print the driver you will see the class name so you can also use this information in the first method above instead.

If you have issues loading your driver (e.g., you get `ClassNotFoundException` errors), make sure you specify *all* dependencies of your driver, not just the main JAR file. They *all* must be listed on the class path. Also make sure your JVM is supported by the driver, trying to load drivers with too old JVM versions also leads to `ClassNotFoundException` errors (as the loader will ignore classes it cannot load). You can always enable debugging information in the `rJava` class loader using `.jClassLoader()$setDebug(1L)` for more verbose output that may help in your troubleshooting.

Due to the fact that JDBC can talk to a wide variety of databases, the SQL dialect understood by the database is not known in advance. Therefore the `RJDBC` implementation tries to adhere to the SQL92 standard, but not all databases are compliant. This affects mainly functions such as `dbWriteTable` that have to automatically generate SQL code. One major ability is the support for quoted identifiers. The SQL92 standard uses double-quotes, but many database engines either don't support it or use other character. The `identifier.quote` parameter allows you to set the proper quote character for the database used. For example MySQL would require `identifier.quote="`"`. If set to `NA`, the ability to quote identifiers is disabled, which poses restrictions on the names that can be used for tables and fields. Other functionality is not affected.

As of `RJDBC 0.2-2` JDBC-specific stored procedure calls starting with `{call` and `{?= call` are supported in the statements.

Value

JDBC returns a `JDBCdriver` object that can be used in calls to `dbConnect`.

`findDrivers` returns a list of Java object references to instances of JDBC drivers that were found. The list can be empty if no drivers were found. Elements can be used as the `driverClass` in calls to JDBC.

See Also

[dbConnect](#)

Examples

```
## Not run:
drv <- JDBC("com.mysql.jdbc.Driver",
```

```
"/etc/jdbc/mysql-connector-java-3.1.14-bin.jar", "`")
conn <- dbConnect(drv, "jdbc:mysql://localhost/test")
dbListTables(conn)
data(iris)
dbWriteTable(conn, "iris", iris)
dbGetQuery(conn, "select count(*) from iris")
d <- dbReadTable(conn, "iris")

## End(Not run)
```

JDBCConnection-class

Class "JDBCConnection"

Description

Class representing a (DBI) database connection which uses JDBC to connect to a database.

Generators

Objects can be created by call to `dbConnect` of a `JDBC` driver.

Slots

`jc`: Java object representing the connection.

`identifier.quote`: Quote character to use for quoting identifiers it automatically generated SQL statements or `NA` if the back-end doesn't support quoted identifiers. Usually the value is inherited from the `"JDBCdriver"`.

`options`: list of options saved with this connection - see `dbOption` and `dbOptions=` parameter to `dbConnect`.

`auto.commit`: logical, the auto-commit mode to maintain in the connection while not in transaction mode. It corresponds to the auto-commit status as the end of `dbConnect()`.

Extends

Class `"DBIConnection-class"`, directly. Class `"DBIObject-class"`, by class `"DBIConnection"`, distance 2.

Methods

No methods defined with class `"JDBCConnection"` in the signature.

See Also

`JDBC`, `"JDBCdriver"`

 JDBCConnection-methods

Methods for the class 'JDBCConnect' in Package 'RJDBC'

Description

Methods for the class 'JDBCConnection' in Package 'RJDBC'.

`dbSendQuery` and `dbSendUpdate` submit a SQL query to the database. The difference between the two is only that `dbSendUpdate` is used with DBML queries and thus doesn't return any result set.

`dbGetTables` and `dbGetFields` are similar to `dbListTables` and `dbListFields` but the result is a data frame with all available details (whereas the latter return only a character vector of the names).

Usage

```
dbSendUpdate (conn, statement, ...)
dbGetTables (conn, ...)
dbGetFields (conn, ...)
```

Arguments

<code>conn</code>	connection object
<code>statement</code>	SQL statement to execute
<code>...</code>	additional arguments to prepared statement substituted for "?"

Details

Some notable enhancements to the DBI API:

`dbSendUpdate` supports vectorized arguments which is far more efficient than using scalar updates. Example: `dbSendUpdate(c, "INSERT INTO myTable VALUES (?, ?)", rnorm(1000), runif(1000))` performs a single `JDBC batchExecute()` call. Additional parameter `max.batch=10000L` is an integer that specifies the maximum batch size supported by the DBMS.

`dbSendQuery` and `dbSendUpdate` accept both `...` (populated first) as well as `list=` (populated as second). Only unnamed arguments are used from `...` (assuming that those are function arguments and no data) while all elements are used from `list=`.

`dbGetQuery` is a shorthand for `sendQuery + fetch`. Parameters `n=-1`, `block=2048L` and `use.label=TRUE` are passed through to `fetch()` others to `dbSendQuery`.

`dbListTables` and `dbGetTables` have the arguments `(conn, pattern="%", schema=NULL)`. `dbExistsTable` is just a wrapper for `dbGetTables`.

`dbWriteTable` is defined as `(conn, name, value, overwrite=FALSE, append=FALSE, force=FALSE, field.types=NULL, ..., max.batch=10000L)` and is just a short-hand for the corresponding `dbSendUpdate()` statements. Since it is only a convenience wrapper, it is strongly recommended to use `dbSendUpdate()` in any real use-cases as you have far more control over the shape and properties of the table if you issue the `CREATE TABLE` statement according to your DBMS' capabilities. It is also possible to create the table with `dbSendUpdate` and then use `dbWriteTable(..., append=TRUE)` to store the data which simply constructs `INSERT INTO batch-update statement run with dbSendQuery`.

`dbReadTable` is just a shorthand for `dbGetQuery(c, "SELECT * from <table>")`

Methods

dbBegin signature (conn = "JDBCConnection", ...)
dbCommit signature (conn = "JDBCConnection", ...)
dbDataType signature (dbObj = "JDBCConnection", obj = "ANY", ...)
dbDisconnect signature (conn = "JDBCConnection", ...)
dbExistsTable signature (conn = "JDBCConnection", name = "character", ...)
dbGetException signature (conn = "JDBCConnection", ...)
dbGetFields signature (conn = "JDBCConnection", ...)
dbGetInfo signature (conn = "JDBCConnection", ...)
dbGetQuery signature (conn = "JDBCConnection", statement = "character", ...)
dbGetTables signature (conn = "JDBCConnection", ...)
dbListFields signature (conn = "JDBCConnection", ...)
dbListResults signature (conn = "JDBCConnection", ...)
dbListTables signature (conn = "JDBCConnection", ...)
dbReadTable signature (conn = "JDBCConnection", ...)
dbRemoveTable signature (conn = "JDBCConnection", ...)
dbRollback signature (conn = "JDBCConnection", ...)
dbSendQuery signature (conn = "JDBCConnection", statement = "character", ...)
dbSendUpdate signature (conn = "JDBCConnection", statement = "character", ...)
dbWriteTable signature (conn = "JDBCConnection", ...)

JDBCDriver-class *Class "JDBCDriver"*

Description

A DBI driver that uses any JDBC driver to access databases.

Generators

Objects can be created by calls to [JDBC](#) or [dbDriver](#).

Slots

identifier.quote: Quote character to use for identifiers in automatically generated SQL statements or NA if quoted identifiers are not supported by the back-end.

jdrrv: Java object reference to an instance of the driver if the driver can be instantiated by a default constructor. This object is only used as a fall-back when the driver manager fails to find a driver.

options: list of options saved with this driver - see [dbOption](#).

Extends

Class "[DBIDriver-class](#)", directly. Class "[DBIObject-class](#)", by class "DBIDriver", distance 2.

Methods

No methods defined with class "JDBCdriver" in the signature.

See Also

[JDBC](#)

JDBCdriver-methods *Methods for the class JDBCdriver in Package 'RJDBC'*

Description

Methods for the class 'JDBCdriver' in Package 'RJDBC'.

Most prominent method is `dbConnect`, it creates a new JDBC connection using the specified driver. From RJDBC 0.2-9 on the driver takes precedence over `DriverManager`, because `DriverManager` is static and is not capable of finding drivers dynamically. `DriverManager` is now only used if the driver is a NULL-driver, i.e., JDBC (NULL).

There are only two positional properties `user=' '` and `password=' '` neither of which will be set if empty. All other arguments are treated as additional properties passed to the connection (except when `DriverManager` is used).

Additional arguments to `dbConnect ()` properties are set

`dbListConnections` always return NULL with a warning, because JDBC connections are not tracked.

`dbGetInfo` returns very basic information, because the JDBC driver is not loaded until a connection is created.

`dbUnloadDriver` is a no-op in the current implementation, because drivers are never removed from the JVM.

Methods

dbConnect signature (drv = "JDBCdriver", ...)

dbListConnections signature (drv = "JDBCdriver", ...)

dbGetInfo signature (drv = "JDBCdriver", ...)

dbUnloadDriver signature (drv = "JDBCdriver", ...)

JDBCResult-class *Class "JDBCResult"*

Description

Representation of a DBI result set returned from a JDBC connection.

Generators

Objects can be created by call to [dbSendQuery](#).

Slots

`jr`: Java reference to the JDBC result set

`md`: Java reference to the JDBC result set meta data

`env`: Environment holding cached objects (currently the result helper object used by `fetch()`)

`stat`: Java reference to the JDBC statement which generated this result

`conn`: Connection which created this result set

Extends

Class "[DBIResult-class](#)", directly. Class "[DBIObject-class](#)", by class "DBIResult", distance 2.

Methods

No methods defined with class "JDBCResult" in the signature.

See Also

[JDBC](#), [dbSendQuery](#)

JDBCResult-methods *Methods for the class JDBCResult in Package 'RJDBC' ~~*

Description

Methods for the class 'JDBCResult' in Package 'RJDBC'.

`fetch` retrieves the content of the result set in the form of a data frame. If `n` is -1 then the current implementation fetches 32k rows first and then (if not sufficient) continues with chunks of 512k rows, appending them. If the size of the result set is known in advance, it is most efficient to set `n` to that size.

Additional argument `block` can be used to inform the driver that pre-fetching of a certain block of records is desirable (see `setFetchSize()` in JDBC) leading to possibly faster pulls of large queries. The default is to pre-fetch 2048 records. Note that some databases (like Oracle) don't support a fetch size of more than 32767. If set to `NA` or anything less than 1 then the fetch size is not changed. This is only a hint, drivers are free to ignore it.

Finally, `useLabel` logical argument determines whether column labels are used for naming (`TRUE`, default) or column names should be used (`FALSE`) since some database drivers do not implement labels correctly.

`dbClearResult` releases the result set.

`dbColumnInfo` returns basic information about the columns (fields) in the result set.

`dbGetInfo` returns an empty list.

`dbListResults` returns an empty list and warns that JDBC doesn't track results.

Methods

fetch signature (res = "JDBCResult", ...)

dbClearResult signature (res = "JDBCResult", ...)

dbColumnInfo signature (res = "JDBCResult", ...)

dbGetInfo signature (res = "JDBCResult", ...)

dbListResults signature (res = "JDBCResult", ...)

Index

* classes

JDBCConnection-class, 5
JDBCDriver-class, 7
JDBCResult-class, 9

* interface

dbOption, 1
dbSetTypeMaps, 2
JDBC, 3

* methods

JDBCConnection-methods, 6
JDBCDriver-methods, 8
JDBCResult-methods, 9

dbBegin, JDBCConnection-method
(*JDBCConnection-methods*), 6
dbClearResult, JDBCResult-method
(*JDBCResult-methods*), 9
dbColumnInfo, JDBCResult-method
(*JDBCResult-methods*), 9
dbCommit, JDBCConnection-method
(*JDBCConnection-methods*), 6
dbConnect, 1-5
dbConnect, JDBCDriver-method
(*JDBCDriver-methods*), 8
dbDataType, JDBCConnection-method
(*JDBCConnection-methods*), 6
dbDisconnect, JDBCConnection-method
(*JDBCConnection-methods*), 6
dbDriver, 7
dbExistsTable, JDBCConnection, ANY-method
(*JDBCConnection-methods*), 6
dbExistsTable, JDBCConnection-method
(*JDBCConnection-methods*), 6
dbGetException, JDBCConnection-method
(*JDBCConnection-methods*), 6
dbGetFields
(*JDBCConnection-methods*), 6
dbGetFields, JDBCConnection-method
(*JDBCConnection-methods*), 6
dbGetInfo, JDBCConnection-method
(*JDBCConnection-methods*), 6
dbGetInfo, JDBCDriver-method
(*JDBCDriver-methods*), 8

dbGetInfo, JDBCResult-method
(*JDBCResult-methods*), 9
dbGetQuery, JDBCConnection, character-method
(*JDBCConnection-methods*), 6
dbGetTables
(*JDBCConnection-methods*), 6
dbGetTables, JDBCConnection-method
(*JDBCConnection-methods*), 6
dbListConnections, JDBCDriver-method
(*JDBCDriver-methods*), 8
dbListFields, JDBCConnection, ANY-method
(*JDBCConnection-methods*), 6
dbListFields, JDBCConnection-method
(*JDBCConnection-methods*), 6
dbListResults, JDBCConnection-method
(*JDBCConnection-methods*), 6
dbListTables, JDBCConnection-method
(*JDBCConnection-methods*), 6
dbOption, 1, 3, 5, 7
dbOption, JDBCConnection-method
(*dbOption*), 1
dbOption, JDBCDriver-method
(*dbOption*), 1
dbOption, JDBCResult-method
(*dbOption*), 1
dbReadTable, JDBCConnection, ANY-method
(*JDBCConnection-methods*), 6
dbReadTable, JDBCConnection, character-method
(*JDBCConnection-methods*), 6
dbReadTable, JDBCConnection-method
(*JDBCConnection-methods*), 6
dbRemoveTable, JDBCConnection, ANY-method
(*JDBCConnection-methods*), 6
dbRemoveTable, JDBCConnection-method
(*JDBCConnection-methods*), 6
dbRollback, JDBCConnection-method
(*JDBCConnection-methods*), 6
dbSendQuery, 9
dbSendQuery, JDBCConnection, character-method
(*JDBCConnection-methods*), 6
dbSendUpdate
(*JDBCConnection-methods*), 6
dbSendUpdate, JDBCConnection, character-method

(*JDBCConnection-methods*), 6
dbSetTypeMaps, 2
dbUnloadDriver, JDBCDriver-method
(*JDBCDriver-methods*), 8
dbWriteTable, 4
dbWriteTable, JDBCConnection, ANY-method
(*JDBCConnection-methods*), 6
dbWriteTable, JDBCConnection-method
(*JDBCConnection-methods*), 6

fetch, 2
fetch, JDBCResult, numeric-method
(*JDBCResult-methods*), 9
findDrivers (*JDBC*), 3

getOption, 1

JDBC, 1, 2, 3, 5, 7–9
JDBCConnection-class, 5
JDBCConnection-methods, 6
JDBCDriver, 5
JDBCDriver-class, 7
JDBCDriver-methods, 8
JDBCResult-class, 9
JDBCResult-methods, 9