
intake_sql Documentation

Release 0.1.1

Martin Durant

Jul 25, 2018

Contents:

1	Quickstart	3
1.1	Installation	3
2	API Reference	5
3	Indices and tables	9

Access any SQL data service which *sqlalchemy* can talk to from Intake.

CHAPTER 1

Quickstart

`intake-sql` provides quick and easy access to tabular data stored in sql data sources.

1.1 Installation

To use this plugin for `intake`, install with the following command:

```
conda install -c intake intake-sql
```

In addition, you will also need other packages, depending on the database you wish to talk to. For example, if your database is Hive, you will also need to install `pyhive`.

CHAPTER 2

API Reference

<code>intake_sql.SQLManualPartition</code>	
<code>intake_sql.SQLSource(uri, sql_expr[, ...])</code>	One-shot SQL to dataframe reader (no partitioning)
<code>intake_sql.SQLSourceAutoPartition(uri,</code> <code>...)</code>	SQL table reader with automatic partitioning
<code>intake_sql.SQLSourceManualPartition(uri,</code> <code>...)</code>	SQL expression reader with explicit partitioning
<code>intake_sql.SQLCatalog(uri[, views])</code>	Makes data sources out of known tables in the given SQL service

class `intake_sql.SQLSource(uri, sql_expr, sql_kwargs={}, metadata={})`

One-shot SQL to dataframe reader (no partitioning)

Caches entire dataframe in memory.

Parameters

uri: str or None Full connection string in sqlalchemy syntax

sql_expr: str Query expression to pass to the DB backend

sql_kwargs: dict Further arguments to pass to pandas.read_sql

Attributes

datasource

description

hvplot Returns a hvPlot object to provide a high-level plotting API.

plot Returns a hvPlot object to provide a high-level plotting API.

Methods

<code>close()</code>	Close open resources corresponding to this data source.
<code>discover()</code>	Open resource and populate the source attributes.
<code>read()</code>	Load entire dataset into a container and return it
<code>read_chunked()</code>	Return iterator over container fragments of data source
<code>read_partition(i)</code>	Return a (offset_tuple, container) corresponding to i-th partition.
<code>to_dask()</code>	Return a dask container for this data source
<code>yaml()</code>	Return YAML representation of this data-source

`read()`
Load entire dataset into a container and return it

class `intake_sql.SQLSourceAutoPartition(uri, table, index, sql_kwargs={}, metadata={})`

SQL table reader with automatic partitioning

Only reads existing tables, not arbitrary SQL expressions.

For partitioning, require to provide the column to be used, which should be indexed in the database. Can then provide list of boundaries, number of partitions or target partition size; see `dask.dataframe.read_sql_table` and examples for a list of possibilities.

Parameters

uri: str or None Full connection string in sqlalchemy syntax

table: str Table to read

index: str Column to use for partitioning and as the index of the resulting dataframe

sql_kwargs: dict Further arguments to pass to `dask.dataframe.read_sql`

Attributes

datasource

description

hvplot Returns a hvPlot object to provide a high-level plotting API.

plot Returns a hvPlot object to provide a high-level plotting API.

Methods

<code>close()</code>	Close open resources corresponding to this data source.
<code>discover()</code>	Open resource and populate the source attributes.
<code>read()</code>	Load entire dataset into a container and return it
<code>read_chunked()</code>	Return iterator over container fragments of data source
<code>read_partition(i)</code>	Return a (offset_tuple, container) corresponding to i-th partition.
<code>to_dask()</code>	Return a dask container for this data source
<code>yaml()</code>	Return YAML representation of this data-source

read()

Load entire dataset into a container and return it

to_dask()

Return a dask container for this data source

```
class intake_sql.SQLSourceManualPartition(uri,           sql_expr,           where_values,
                                         where_template=None, sql_kwargs={}, meta-
                                         data={})
```

SQL expression reader with explicit partitioning

Reads any arbitrary SQL expressions into partitioned data-frame, but requires a full specification of the boundaries.

The boundaries are specified as either a set of strings with *WHERE* clauses to be applied to the main SQL expression, or a string to be formatted with a set of values to produce the complete SQL expressions.

Note, if not supplying a *meta* argument, dask will load the first partition in order to determine the schema. If some of the partitions are empty, loading without a meta will likely fail.

Parameters

uri: str or None Full connection string in sqlalchemy syntax

sql_expr: str SQL expression to evaluate

where_values: list of str or list of values/tuples Either a set of explicit partitioning statements (e.g., “*WHERE index_col < 50*”...) or pairs of values to be entered into *where_template*, if using

where_template: str (optional) Template for generating partition selection clauses, using the values from *where_values*, e.g., “*WHERE index_col >= {} AND index_col < {}*”

sql_kwargs: dict Further arguments to pass to pd.read_sql_query

Attributes

datasource

description

hvplot Returns a hvPlot object to provide a high-level plotting API.

plot Returns a hvPlot object to provide a high-level plotting API.

Methods

<code>close()</code>	Close open resources corresponding to this data source.
<code>discover()</code>	Open resource and populate the source attributes.
<code>read()</code>	Load entire dataset into a container and return it
<code>read_chunked()</code>	Return iterator over container fragments of data source
<code>read_partition(i)</code>	Return a (offset_tuple, container) corresponding to i-th partition.
<code>to_dask()</code>	Return a dask container for this data source
<code>yaml()</code>	Return YAML representation of this data-source

read()

Load entire dataset into a container and return it

to_dask()

Return a dask container for this data source

class intake_sql.SQLCatalog(uri, views=False, **kwargs)

Makes data sources out of known tables in the given SQL service

Attributes

datasource

description

hvplot Returns a hvPlot object to provide a high-level plotting API.

plot Returns a hvPlot object to provide a high-level plotting API.

Methods

<code>close()</code>	Close open resources corresponding to this data source.
<code>discover()</code>	Open resource and populate the source attributes.
<code>force_reload()</code>	Imperative reload data now
<code>read()</code>	Load entire dataset into a container and return it
<code>read_chunked()</code>	Return iterator over container fragments of data source
<code>read_partition(i)</code>	Return a (offset_tuple, container) corresponding to i-th partition.
<code>reload()</code>	Reload catalog if sufficient time has passed
<code>to_dask()</code>	Return a dask container for this data source
<code>walk([sofar, prefix, depth])</code>	Get all entries in this catalog and sub-catalogs
<code>yaml()</code>	Return YAML representation of this data-source

CHAPTER 3

Indices and tables

- genindex
- modindex
- search

Index

R

read() (intake_sql.SQLSource method), [6](#)
read() (intake_sql.SQLSourceAutoPartition method), [7](#)
read() (intake_sql.SQLSourceManualPartition method), [7](#)

S

SQLCatalog (class in intake_sql), [8](#)
SQLSource (class in intake_sql), [5](#)
SQLSourceAutoPartition (class in intake_sql), [6](#)
SQLSourceManualPartition (class in intake_sql), [7](#)

T

to_dask() (intake_sql.SQLSourceAutoPartition method),
 [7](#)
to_dask() (intake_sql.SQLSourceManualPartition
method), [7](#)