

# Development notes for GAVO DaCHS

**Author:** Markus Demleitner  
**Email:** [gavo@ari.uni-heidelberg.de](mailto:gavo@ari.uni-heidelberg.de)  
**Date:** 2023-02-09  
**Copyright:** Waived under [CC-0](https://creativecommons.org/licenses/by/4.0/)

## Contents

<b>Package Layout</b>	<b>2</b>
<b>Getting Table Metadata and Querying Tables</b>	<b>3</b>
<b>Versioning Issues</b>	<b>4</b>
<b>Error handling, logging</b>	<b>4</b>
Exception classes . . . . .	4
The events subsystem . . . . .	5
Catching exceptions . . . . .	5
<b>Testing</b>	<b>6</b>
XSD validation . . . . .	6
Setting package installs up for testing . . . . .	7
Test framework . . . . .	7
Regression testing of data . . . . .	7
pyflakes . . . . .	7
Coverage . . . . .	8
Test Plan . . . . .	8
Testing in Chroots . . . . .	9
<b>Configuration</b>	<b>10</b>
Future . . . . .	10
<b>Structures</b>	<b>10</b>
<b>Metadata</b>	<b>11</b>
Getting Metadata . . . . .	11
Setting Metadata . . . . .	12
Managed Date-like Metadata . . . . .	12

<b>Memoization</b>	<b>13</b>
<b>Profiling</b>	<b>14</b>
<b>Debugging</b>	<b>14</b>
Debugging memory leaks . . . . .	15
Core dumps . . . . .	15
<b>Delimited SQL identifiers</b>	<b>15</b>
<b>Grammars</b>	<b>16</b>
<b>Procedures</b>	<b>16</b>
<b>ADQL User Defined Functions</b>	<b>18</b>
<b>Schema updates</b>	<b>19</b>
<b>XML Schema updates</b>	<b>19</b>
Schema Evolution . . . . .	20
<b>Javascript</b>	<b>20</b>
gavo.js . . . . .	20
samp.js . . . . .	22
jquery and flot . . . . .	22
<b>Stuff in gavo.imp</b>	<b>22</b>
<b>Different Database Backends</b>	<b>22</b>
<b>Writing Documentation</b>	<b>23</b>
<b>Random Stuff</b>	<b>24</b>
Tracing imports . . . . .	24
matplotlib . . . . .	24
Making a New Version of VOTable the Default . . . . .	24
Parsing Text Files . . . . .	25
Licensing . . . . .	25

Some of this is severely out of date.

## Package Layout

The following rules should be followed as regards subpackages of gavo in order to keep the modules dependency graph manageable (and facilitate factoring out libraries).

- Each functionality block is in a subpackage, the `__init__` for which contains the main functions, classes, etc., of the sub-package interface most clients will be concerned with. Clients needing special tricks may still import individual modules (but then they're in much larger danger of breaking). What's in `__init__` should be considered "public interface" and hence changed very carefully if at all.

- Within each subpackage, *no* module imports the sub-package, i.e., a module in base never says "from gavo import base"
- A subpackage may have a module common, containing objects that multiple modules within that subpackage requires. common may *not* import any module from the subpackage, but may be imported from all of them. No rules wrt importing modules from the same subpackage exist of other modules. Just apply common sense here to avoid circular imports.
- Don't use `import *`. It interferes with our static checking. For relative imports, we will probably be slowly migrating towards `from . import` over the current absolute imports (`from gavo.base import...`).
- There is a hierarchy of subpackages, where subpackages lower in the hierarchy may not import anything from the higher or equal levels, but only from lower levels. This hierarchy currently looks like this: `imp [ < ] utils < stc < (votable, adql) < base < dm < rscdef < grammars < formats < rsc < svcs < registry < protocols < web < rscdesc < (helpers, user) utils` should never assume anything from `imp` is present, i.e., it may *attempt* to import from there, but it should not fail hard if the import doesn't work. Of course, concrete functions (e.g., `from utils.fitstools`) won't work if the base libraries are not present.

## Getting Table Metadata and Querying Tables

The preferred way to do simple queries against tables in DaCHS these days is:

- (a) get the table metadata:

```
td = base.resolveCrossId("resdir/q#mytable")
```

- (b) use the `td`'s `doSimpleQuery(selectClause, fragments, params)` method to get dicts of rows; all arguments are optional and default to pulling all; `selectClause` is just a list of column names:

```
for row in td.doSimpleQuery(["col1", "col2"],
    "col1<%(lim1)s AND col2=%(foo)s",
    {'lim1': 23, 'foo': 42}):
    print row
```

When you need explicit connection management or want to do more complex operations, use the context managers (typically `getTableConn` when querying, `getWritableAdminConn` when writing) using a context manager:

```
with base.getTableConn() as conn:
    for row in conn.queryToDicts(myComplexQuery,
        {'arg1': 23, 'pat': 'M32'}):
        ...
```

Besides `queryToDicts`, there's also `query`, which yields tuples. Both are iterators, which means that the queries contained *will not be executed* unless you fetch at least one row. Hence, for queries that don't return anything (DDL, inserts, etc), use `conn.execute`.

In contrast to `doSimpleQuery`, all these will do case folding of the select list items.

In "user" code, get these symbols from `api` instead of from `base`.

Another feature you might like to know about is the connections's `parameters` context manager. It takes postgres settings keys and values and will reset them at the end of the controlled block. This is useful for things like timeouts and the like, e.g.,:

```
with conn.parameters([
    ("statement_timeout", "%s ms"%int(timeout*1000))]:
    whatever
```

## Versioning Issues

DaCHS itself is versioned such that minor versions (e.g., 2.6) are releases, which technically means they get a bit more pre-release testing, and they get properly announced on `dachs-users`. Micro versions on to of that are then beta releases including new features up to the next stable release. Hence, 2.6.1 will in general be less stable than 2.6. Perhaps we should change this to something less flamboyant one of these days.

When we support protocols, we treat major versions as separate standards, i.e., there are separate RDs for, say, `siap` and `siap2`. Within each such RD, mixins or other evolving material may be tagged with the minor version. For instance a mixin `table-0` would correspond to version 1.0 (or 2.0) of a standard. This hasn't been done consistently in DaCHS' past, so you'll see all kinds of other experiments. But the minor-version tagging is what should happen with future developments.

## Error handling, logging

### Exception classes

It is the goal that all errors that can be triggered from the web or from within resource descriptors yield sensible error messages with, if possible, information on the location of the error. Also, major operations changing the content of the database should be loggable with time and, probably, user information.

The core of error processing is `utils.excs`. All "sensible" exceptions (i.e., `MemoryErrors` and software bugs excepted) should be instances of `gavo.excs.Error`. However, upwards from `base` you should always raise exceptions from `base`; all ("public") exception types from `utils.excs` are available there (i.e., raise `base.NotFoundError(...)` rather than `utils.excs.NotFoundError(...)`).

The base class takes a `hint` argument at construction that should give additional information on how to fix the problem that gave rise to the exception. All exception constructor arguments except the first one must always be keyword arguments as a simple hack to allow pickling the exceptions.

When defining new exceptions, if there is structured information (e.g., line numbers, keys, and the like), always keep the information separate and use the `__str__` method of the exception to construct something humans want to see. All built-in exceptions should accept a `hint` keyword.

## The events subsystem

All proper DC code (i.e. above base) should do user interaction through `base.ui.notify<something>`. In base and below, you can use `utils.sendUIEvent`, but this should be reserved for weird circumstances; code so far down should't normally need to do user interaction or similar.

The `<something>` can be various things. `base.events` defines a class `EventDispatcher` (an instance of which then becomes `base.ui`) that defines the `notify<something>` methods. The docstrings there explain what you're supposed to pass, and they explain what observers get.

`base.events` itself does very little with the events, and in particular it does not do any user interaction -- the idea is that I may yet want to have Tkinter interfaces or whatever, and they should have a fair chance to control the user interaction of a program.

The actual action on events is done by observers; these are ususally defined in `user`, and some can be selected from the `dachs` command line. For convenience, you should derive your `Observer` classes from `base.ObserverBase`. This lets you stuff like:

```
from gavo.base import ObserverBase, listensTo

class PlainUI(ObserverBase):
    @listensTo("NewSource")
    def announceNewSource(self, srcString):
        print "Starting %s"%srcString
```

However, you can also just handle single events by saying things like:

```
from gavo import base

def handleNewSource(srcToken):
    pass

base.ui.subscribeNewSource(handleNewSource)
```

Most logging is done in `user.logui`; if you want logging, say:

```
from gavo.user import logui
logui.LoggingUI(base.ui)
```

## Catching exceptions

In the DC software, it is frequently desirable to ignore the first rule of exception handling, viz., leave them alone as much as possible. Instead, we often map exceptions to DC-internal exceptions (this is very relevant for everything leading up to `ValidationErrors`, since they are used in user interaction on the web interface). However, to make the original exception information available for debugging or problem fixing, whenever you "translate" an exception, have `base.ui.notifyExceptionMutation(newException)` called. This should arrange logging the exception to the error log (although of course that's up to the observer selected).

The convenient way to do this is to call `ui.logOldExc(exc)`:

```
raise base.ui.logOldExc(GavoError(...))
```

`LoggingUI` only logs the information on old exceptions when `base.DEBUG` is true. You can set this from your code, or by passing the `--debug` option to `gavo`.

This should probably be phased out now that `python3` monitors and exposes exception mutation itself.

## Testing

In an installed checkout of DaCHS, you can go to the `tests` subdirectory and run:

```
python3 runAllTests.py
```

for a fairly extensive set of unit tests.

This needs to create a test database, and that will only work if whoever runs this is postgres superuser. `dachsroot` from the Debian package already is. If you want to run tests as another user, you'll have to make yourself a suitable account, typically with:

```
sudo -u postgres createuser -s 'id -nu'
```

The suite will not tear down and build up everything each time it's called. To make it rebuild everything, remove `~/_gavo_test` and `dropdb dachstest`.

Also, you'll need the extra package `python3-testresource` (which the `dachs` packages don't declare as a dependency), and you'll need `build-essentials` as well as `libcfitsio-dev` I'm testing against concrete error messages, and DaCHS sometimes hands through messages from the database. Hence, some tests will fail when `lc_messages` in `postgresql.conf` isn't `c`.

This uses some management of test scaffolds; when something is severely wrong, generating these scaffolds can fail and the execution of the suite will stop. I'm not decided whether to regard that as a bug or a feature, but I'll not fix it any time soon. So, if this bites you, find out why resource generation fails and fix it.

### XSD validation

XML Schema is a pain all around, and given that we don't want to hit W3C and IVOA with requests for schema files every time someone needs schema validation (which includes RD validation and unit tests), DaCHS goes to some lengths to use its own schema files.

The main engine here is the LXML-based validator from `gavo.helpers.testtricks`; the rocket science part of this is to make LXML use the plethora of schema files we have locally.

"Locally" here means in the `schemata` subdirectory of the distribution. When you add a schema there that should be available in validation, you also need to add the filename to `gavo.testtricks.VO_SCHEMATA` (background: we keep some schema files in `gavo/schema` that the validator should not be bothered with; still, we should probably just pull in `*.xsd` at some point).

With this, run `dachs admin xsdVal` to XSD-validate a VO file.

In case you'd like some external truth, here's how you can run `xerces` as a validating parser on a Debian system:

```
export CLASSPATH=/usr/share/doc/libxerces2-java-doc/examples/xercesSamples.jar:/usr/share/java/xerces
exec java dom.Counter -n -v -s -f $@
```

This will only work if the `schemaLocation` attributes are present wherever new namespaces are introduced. For DaCHS' VOResource output, that is the case. We don't do that in VOTables and several other places.

## Setting package installs up for testing

The debian package does not contain unit tests. If you want to nevertheless run them, check out the release corresponding to your package from <http://svn.ari.uni-heidelberg.de/svn/gavo/python/tags/>. Again, see the `tests` subdirectory in your checkout.

## Test framework

All unit tests must import `gavo.helpers.testhelpers` before importing anything else from the `gavo` namespace. This is because `testhelpers` sets up a test environment in `~/gavo_test` (set in `tests/test_data/test-gavorc`). To make this work reliably, it must manipulate the normal way configuration files are read.

`helpers.testhelpers` needs a `dachstest` database for which the current user is a superuser. It will create it provided you're a DB superuser with `ident` authentication (see `install` to figure out how to set this up).

There are `doctests` in modules (though fewer than I'd like), and `pyunit`- and `trial`-based tests in `<project root>/tests`. `tests/runAllTests.py` takes care of locating and executing them all.

In addition to setting up the test environment, `testhelpers` provides (check out the source) some useful helper functions (like `getTestRD`), the `VerboseTest` class adding test resources and some assertions to the normal `unittest.TestCase`. Do *not* import it in production code. Test-like functionality interesting to production code should go to `helpers.testtricks`.

`testhelpers.main` is useful after an `if __name__=='__main__'` in test modules. Pass a default test class, and you can call the module without arguments (in which case it will run all tests), with a single argument (that will be interpreted as a method prefix to locate tests on the default `TestCase`) or with two arguments (a `TestCase` name and a method prefix to find the methods to be run). All `pyunit`-based tests use this `main`.

`testhelpers.main` evaluates the `TEST_VERBOSEITY` environment variable. With `TEST_VERBOSEITY=2`, you'll see the test names as they are executed.

## Regression testing of data

For certain kinds of data, unit testing is useful, too. Since it's always possible that server code changes may break such tests, it makes sense to run those unit tests at each commit. Therefore, `tests/runAllTests.py` has a facility to pick up such tests from directories named in `$GAVO_INPUTS` (the "real" one, not the fake test one) in the `__tests/__unitpaths`. It will pick up tests from there just as it picks them up from `tests`. Such data-based tests (typically) must run "out of tree", i.e., in the actual server environment where the resources expected by the service tested are. To keep `testhelper` from fudging the environment, set the environment variable `GAVO_OOTTEST` to anything before importing `testhelpers`. This is conveniently done in `python`, like this:

```
import os
os.environ["GAVO_OOTTEST"] = "dontcare"

from gavo.helpers import testhelpers
```

## pyflakes

Not really testing, but static code checking using `pyflakes` should regularly be done, and result in no warnings eventually (right now, more annotations are required).

We have added a simple ignoring facility in our `pyflakes` driver, `tests/flake_all.py`:

- To ignore (not check) an entire file, add, preferably near the top, a line like:

```
# Not checked by pyflakes: (reason)
```

Please always give a reason so people can tell whether it has gone away and the file should now be included in the checks.

- To ignore a single error, add a comment like:

```
#noflake: (rationale)
```

to the line reported by pyflakes.

Also note that `flake_all` hardcodes that modules from `imp` are not checked.

## Coverage

There's a shell script `genCov.sh` in `tests` that runs all unit tests and all regression tests. It then combines the coverage from all these runs to `.coverage`. So, after running this, just running `python3-coverage report -i` or `python3-coverage html -i` should do the right thing.

You need the `-i` flag because during testing a lot of generated or extracted code will be executed, and there's no sane way I can think of to include that with the testing. For RD-embedded code it might still work, but I won't tackle this any time soon (though, sure, there's probably a severe lack of testing for all the code in the system RDs).

To exclude code from coverage computation, use:

```
# pragma: no cover
```

## Test Plan

(This is somewhat specific to Markus' setup; something similar is recommended for everyone, though)

*Before every commit*, do:

- start a local server
- go to `$checkout/tests`
- `python flake_all.py` (which does some static code checking)
- `python runAllTests.py` (which arranges for doctests, pyunit tests, trial tests, and data unit tests to be run)
- run `dachs val -tv ALL` (which, apart from validating the RDs, also runs the RD-defined regression tests against the server running locally)
- go to `$checkout`
- run `svn status` to make sure no files are left not in version control or explicitly ignored

*After a checkout on the production server*, do:

- `dachs test -t bigserver -u http://dc.g-vo.org/ ALL` (which runs all tests defined in the local RDs, even for the production server, against the production server; this does what it's supposed to do as the repo for the RDs is the same on development and production).



## Testing in Chroots

When we release to Debian, we release to unstable, and hence we need to test in a chroot. Here's my current practice:

- (a) create a chroot environment; I've been using `DIST=unstable git-pbuilder create`
- (b) get yourself a shell in there and make sure the dependencies for our unit tests are in (the path will probably be different for you):

```
sudo chroot /home/pbuilder/unstable-i386/base.cow/ bash

apt install subversion vim
apt install libjs-jquery libjs-jquery-flot python3-cryptography python3-docutils python3-lxml
```

- (c) Make a user to run the things, and make it identical to what you're using to run the unit tests outside of the chroot:

```
adduser msdemlei (well, your user name)
set uid to 135 (well, whatever you have on the host)
addgroup gavo
adduser msdemlei gavo
```

- (d) Become the user and get create the environment to run the unit tests:

```
su - msdemlei
mkdir -p gavo/logs
svn co https://svn.ari.uni-heidelberg.de/svn/gavo/python/trunk
```

As root in the container:

```
python3 /home/msdemlei/trunk/setup.py develop
```

Back as msdemlei:

```
cd trunk/tests
python3 runAllTests.py
```

This will fail because it won't find the database (or createdb). Just copy the data from your external test harness, i.e., from the host system:

```
cd ~/_gavo_test/etc
cp * /home/pbuilder/unstable-i386/base.cow/home/msdemlei/_gavo_test/etc)
```

Back in the chroot, this should now work::

```
python3 runAllTests.py
```

To run DaCHS interactively, an alias like the following comes handy:

```
alias testgavo='env GAVOSETTINGS=/home/msdemlei/trunk/tests/test_data/test-gavorc GAVOCUSTOM=/i
```

Due to the uid magic, you can edit the in-chroot files in the host system; just go to the `.../base.cow/home/msdemlei/trunk/...`

## Configuration

DaCHS has far too many configuration hooks: `gavo.rc`, `defaultmeta.txt`, the database profiles, `vanitynames.txt`, `userconfig.rd`, as well as locally-overridden system RDs and templates. At least `defaultmeta.txt` was a mistake, as was probably `vanitynames.txt`. We should be working on getting rid of it.

New configuration should preferably go into `userconfig.rd`, while there's always going to be room for `gavo.rc`, too.

Configuration items for `userconfig.rd` typically are going to be STREAMs. To provide fallbacks for those if the user hasn't defined any, there's `//userconfig`, which also serves as built-in documentation for what's there. As an identifier is resolved in `//userconfig`, the system first looks in a `etc/userconfig.rd` and then, even if that file exists (but has no element with the id in question), in `//userconfig`.

When using the elements, always use the canonical abbreviation for `userconfig`, `%`, as in `<FEED source="%#registry-interfacerecords"/>`.

## Future

If you add features that will make DaCHS produce responses that may break legacy components (e.g., new Registry features), make them conditional on future entries. That's done by choosing a suitable string (e.g., `dali-interface-in-tap-1`) and then protecting the generation of the new elements with something like:

```
if "dali-interface-in-tap-1" in base.getConfig("future"):
    ...
```

To try out the change, write:

```
future: dali-interface-in-tap-1, some-other-new-feature-if-necessary
```

into your `gavo.rc`.

Once the change is sufficiently widely accepted, remove the condition, and all DaCHSes will produce new-style responses.

If a change is suitably safe so it can be enabled by default, invert the logic and use `no-new-feature` to let people turn things off if they're causing trouble.

Don't forget to add the future keys in `tests/test_data/test-gavorc` for while you're testing the experimental features.

## Structures

Resource description within the DC works via instances of `base.Structure`. These parse themselves from XML strings, do validation, etc. All compound RD elements correspond to a structure class (well, almost; meta is an exception).

A structure instance has the following callbacks:

- `completeElement(ctx)` -- called when the element's closing tag is encountered, used to fill in computed defaults. `ctx` is a parse context that you can use to, e.g. resolve XML ids.
- `validate()` -- called after `completeElement`, used to raise errors if some gross ("syntactic") mistakes are in the element

- `onElementComplete()` -- called after `validate`, i.e., `elementCompleted` can rely on seeing a "valid" structure

In addition, attributes can define `onParentCompleted` methods. These are called after `onElementCompleted` of the parent element is run when the attribute value is different from its default. They receive the new attribute value as the single argument. Maybe this is bad design.

This processing is done automatically when parsing elements from XML. When building elements manually, you must call the structure's `finishElement` method when done to arrange for these methods being called; to make sure this happens, you usually want to construct Structures using `base.makeStruct`.

If you override these methods, you (almost) always want to call the corresponding super-classes' methods using `super().methname([ctx])`. Structures in DaCHS sometimes use multiple inheritance, and hence there's really no alternative to using `super` here. To make sure this works as expected, any (python) mixin for structures must inherit from `base.StructCallbacks`.

The `user.docgen` module makes documentation out of these structures. There are several catches. One of the more striking is that element names in the *entire* DaCHS code must be unique, since `docgen` generates section heading from those names and actually checks that these headings are unique; hence, only one (essentially randomly selected) of two identically-named elements would be documented, and parent links would both point there.

Since there are cases when that limitation is a real pain (e.g., the `publish` element of `services` and `data`), there's a workaround: you can set a `docName_` class attribute on a structure that contains the name used for the documentation. See `rsedef.common.Registration` for an example.

## Metadata

"Open" metadata (as opposed to the attributes of columns and the like) is kept in a `meta_` structure added by `base.meta.MetaMixin`. You should probably not access that attribute directly if at all possible since the current implementation is incredibly messy and liable to change.

For this kind of metadata, a simple inheritance exists. `MetaMixins` have a `setMetaParent` method that declares another structure as the current's meta parent. Any request for metadata that cannot be satisfied from self will then be propagated up to this parent (unless propagation is suppressed). Usually, parents will call their children's `setMetaParent` methods.

The metadata is organized in a tree with `MetaItem`'s as nodes. Each `MetaItem` contains one or more children that are instances of `MetaValue` (or more specialized classes). A `MetaValue` in turn can have more `MetaItem` children.

## Getting Metadata

Metadata are accessed by name (or "key", if you will).

The `getMeta(key, ...)->MetaItem` method usually follows the inheritance hierarchy up, meaning that if a meta item is not found in the current instance, it will ask its parent for that item, and so on. If no parent is known, the meta information contained in the configuration will be consulted. If all fails, a default is returned (which is set via a

keyword argument that again defaults to None) or, if the `raiseOnFail` keyword argument evaluates to true, a `gavo.NoMetaKey` exception is raised.

If you require metadata exactly for the item you are querying, call `getMeta(key, propagate=False)`.

`getMeta` will raise a `gavo.MetaCardError` when there is more than one matching meta item. For these, you will usually use a builder, which will usually be a subclass of `meta.metaBuilder`. `web.common.HtmlMetaBuilder` is an example of how such a thing may look like, for simple cases you may get by using `ModelBasedBulder` (see the registry code for examples). This really is too messy and needs to be replaced by something smarter.

The builders are passed to a `MetaMixin`'s `buildRepr(metakey, builder)` method that returns whatever the builder's `getResult` method returns.

## Setting Metadata

You can programmatically set metadata on any metadata container by calling its method `addMeta(key, value)`, where both key and value are (unicode-compatible) strings. You can build any hierarchy in this way, provided you stick with typeless meta values or can do with the default types. Those are set by key in `meta._typesForKeys`.

To build sequences, call `addMeta` repeatedly. To have a sequence of containers, call `addMeta` with `None` or an empty string as value, like this:

```
m.addMeta("p.q", "x") m.addMeta("p.r", "y") m.addMeta("p", None)
m.addMeta("p.q", "u") m.addMeta("p.r", "v")
```

More complex structures require direct construction of `MetaValues`. Use the `makeMetaValue` factory for this. This function takes a value (default empty), and possibly a key and/or type arguments. All additional arguments depend on the meta type desired. These are documented in the [reference manual](#).

The type argument selects an entry in the `meta._typesForKeys` table that specifies that, e.g., `_related` meta items always are links. You can also give the type directly (which overrides any specification through a key).

This can look like this:

```
m.addMeta("info", meta.makeMetaValue("content", type="info",
    infoName="someInfo", infoValue="GIVEN"))
```

## Managed Date-like Metadata

As almost everywhere, date-like metadata is a pain; it's not so much because of Babylonian formats (whenever you give a civil date in DaCHS, it should understand plain, basic DALI-flavoured ISO a.k.a. YYYY-MM-DDThh:mm:ss) but because there's so many dates around a resource and a resource descriptor, for instance:

- Date of RD creation
- Date of first publication (should that be the "creation date"?)
- Date of most recent `dachs pub`
- The `mtime` on the RD file
- Date of last change to underlying data

- Date of most recent import

and much more.

Dates like these you're communicating to the registry. This has

- `Resource/@created` -- in DaCHS, that's the manually managed `creationDate` meta.
- `Resource/@updated` -- in DaCHS, that's `datetimeUpdated`; see below
- `Resource/date` -- `_news` meta are turned into `role="updated"` dates. Plus, the `datetimeUpdated` meta is made into a date, too. Finally, you can manually create date meta items (with role children) that are just copied into `VOResource date`.

DaCHS keeps the the following date-like (i.e., values are ISO strings) metadata on RDs (warning: could still be wrong; this is a plan as of now).

- `creationDate` -- manually defined in RDs
- `_dataUpdated` -- the date the last time any `dachs imp` was run on this RD
- `_metadataUpdated` -- on the RD, this is the `mtime` of the RD source file (if it exists; otherwise that meta is missing). On published items, it's the time of the last `dachs pub`. This latter rule is so the `dataUpdated` on the registry record remains meaningful.

## Memoization

The `base.caches` module should be the central point for all kinds of memoization/caching tasks; in particular, if you use `base.caches`, your caches will automatically be cleared on `dachs serve reload`. To keep dependencies and risks of recursive imports low, it is the providing modules' responsibility to register caching functions. The idea is that, e.g., `rsdesc` wants a cache of resource descriptors. Therefore, it says:

```
base.caches.makeCache("getRD", getRD)
```

Clients then say:

```
base.caches.getRD(id).
```

This mechanism for now is restricted to items that come with a unique `id` (the argument). It would be easy to extend this to multiple-argument functions, but I don't think that's a good idea -- the "identities" of the cached objects should be kept simple.

No provision is made to prevent accidental overwriting of function names.

And, of course, individual functions can do `functools.lru_cache`-ing to their heart's delight but should keep in mind that `dachs serve reload` will not clear this.

## Profiling

If you want to profile server actions, try a script like this:

```
"""
Make a profile of server responses.

Call as

trial --profile createProfile.py
"""

import sys

from gavo import api
from gavo.web import dispatcher

sys.path.append("/home/msdemlei/gavo/trunk/tests")

import trialhelpers

class ProfileThis(trialhelpers.RenderTest):
    renderer = dispatcher.ArchiveService()

    def testOneService(self):
        self.assertGETHasStrings("/ppmx/res/ppmx/scs/form",
            {"hscs_pos": "12 2", "hscs_sr": "20.0"},
            ["PPMX"])
```

After running, you can use `pstats` on the file `profile.data`.

To profile actually running DaCHS operations, use the `--profile-to <profile file>` option of the `dachs` program. For the server, you must make sure it cleanly exists in order to have meaningful stats. Do this by accessing `/test/exit` on a debug server.

## Debugging

Just insert lines like:

```
import pdb;pdb.Pdb(nosigint=True).set_trace()
```

wherever required to have python dump you into the debugger and let you look around, single-step, etc.

When you want to inspect what's going on within the server, in particular when something only manifests itself after a long time, you may want to have a look at twisted's manhole; quite a bit easier, however, is to use the `debug/q rd` that you can get from <http://svn.ari.uni-heidelberg.de/svn/gavo/hdinputs/debug> and adapt it to your needs. The idea here is that within `q.rd##1` you create customDFs or customRFs exposing what you're interested in. You can then use those in `res/page1.html`. You can edit both files "live", they will both be reloaded as necessary.

## Debugging memory leaks

Sometimes one is careless and leaves a reference somewhere, perhaps in an RD. Since this really only matters in the server, such situations are particularly insidious to debug. To help there, there's some scaffolding in `web.root`.

To activate things, you set `MEM_DEBUG` to `True`. Down in `locateChild` of `ArchiveService`, there's code like:

```
if MEM_DEBUG:
    from gavo.utils import codetricks
    import gc
    gr = gc.get_referrers
    if hasattr(base, "getNewStructs"):
        ns = base.getNewStructs()
        print ">>>>> new structs:", len(ns)
```

What this lets you do is see when new structs are left somewhere in DaCHS' guts. What you do when such a thing happens is higher magic. I've found it helps to put something like a mini-memory debugger right into that handler. There's a rough one in `testtricks`, so you could put in something like:

```
if len(ns)==147:
    from gavo.helpers import testtricks
    ob = ns[0]
    del ns
    testtricks.debugReferenceChain(ob)
```

after the print (of course, this only makes sense if you're running `dachs serve debug`, as the actual server detaches from its tty). This lets you go through the objects referring to the first struct left over by hitting Return.

Enter anything to follow the (inverse) reference, except that a `d` will drop you in the debugger and `x` will continue normal execution. Do this until you see where the reference comes from. Just be aware that many references are harmless -- in particular, this function will hold a reference to the object in question, so you'll need some experience to figure out where to look.

## Core dumps

If you're desperate and need to get core dumps out of a crashing operational server (core dumps from `dachs serve debug` should just work as normal), you need to install the `python3-prctl` package. The core dumps will be in `stateDir`.

## Delimited SQL identifiers

Although it may look like it, we do not really support delimited identifiers (DIs) as column names (and not at all as table names). I happen to regard them as an SQL misfeature and really only want to keep them out of my software.

However, TAP forces me to deal with them at least superficially. That means that using them elsewhere will lead to lots of mysterious error messages from inside of DaCHS's bowels. There still should not be any remote exploits possible when using them.

Here's the deal on them:

They are represented as `utils.miscutils.QuotedName` objects. These `QuotedNames` have some methods to control the impact the partial support for delimited identifiers has on the rest of the software. In particular, when you stringify them, they result in string ready for inclusion into SQL (i.e., hopefully properly escaped). The hash to the name, i.e., there are no implied quotes, and, unfortunately, `hash(di)!=hash(str(di))`.

The DC software right now assumes DIs are ASCII only (what do the standards people say?).

The one real painful thing is the representation of result rows with DIs -- I did not want to have lots of these ugly `QuotedNames` in the result rows, so they end up as SQL-escaped strings when used as keys. This is extra sad since in this way for a DI column `foo`, `rec[QName("foo")]` raises a `KeyError`. To work around this, fields have a key attribute, and `rec[f.key]` should never bomb.

## Grammars

Grammars are DaCHS' means of turning some external data to rowdicts, i.e., dictionaries that map grammar keys to values that are usually strings. They are fed to rowmakers to come up with rows suitable for ingestion (or formatting).

A grammar consists of a `Grammar` object, which is a structure inheriting from `grammars.Grammar`. It contains all the "configuration" (e.g., rules). Grammars have a `parse` method receiving some kind of source token (typically, a file name). You will normally not need to override it.

The real action happens in the row iterator, which is declared in the `rowliterator` class attribute of the grammar. Row iterators should inherit from `grammars.Rowliterator`.

TODO: `yieldsTyped`, `rowfilters`, `sourceFields`, `targetData`

Do not import modules from the grammars subpackage directly. Instead, use `rsedef.getGrammar` with the name of the grammar you want. If you define a new grammar, add a line in `rsedef.builtingrammars.grammarRegistry`. To inspect what grammars are available, consult the keys from `rsedef.grammarRegistry`.

## Procedures

To embed actual (python) code into RDs, you should use the infrastructure given in `rsedef.procdef`. It basically leads up to `ProcApp`, which is what's usually embedded in RDs.

`ProcApp` inherits from `ProcDef`, a procedure definition. Such a definition gives some (python) code that is executed when the procedure is applied. To set up the execution environment of this code, there's the definition's `setup` child.

The `setup` contains code and parameters. The code is executed to set up the namespace that the procedure will run in; it is thus executed once -- at construction -- per procedure. The parameters allow configuration of the procedure. This is the place to do relatively expensive operations like I/O or imports.

For example, `//procs#resolveObject` creates the resolver in its `setup` code; this happens only once per creation of the embedding RD:

```
<procDef type="apply" id="resolveObject">
  <setup>
    <par key="ignoreUnknowns">True</par>
```



```

<par key="identifier" late="True"/>
<code>
    from gavo.protocols import simbadinterface
    resolver = simbadinterface.Sesame(saveNew=True)
</code>
</setup>
<doc>...</doc>
<code>
ra, dec = None, None
try:
    ra, dec = resolver.getPositionFor(identifier)
except KeyError:
    if not ignoreUnknowns:
        raise base.Error("resolveObject could not resolve object"
            " %s."%identifier)
vars["simbadAlpha"] = ra
vars["simbadDelta"] = dec
</code>
</procDef>

```

The setup definition introduced two parameters. One is `ignoreUnknowns`, which is "immediate" and just lets the code see a name `ignoreUnknowns`. As with all `par` elements, the content of the element is a python expression providing a default.

The other parameter, `identifier`, is a "late" identifier. This means that it is evaluated on each application of the procedure, much like a function argument. These are just translated into assignments at the top of the function body, which means that everything available in the procedure code is available; e.g., for rowmaker procedures (i.e., `type="apply"`), you can access `vars` here.

Taken together, `late` and `immediate` `par` allow for all kinds of configuration of procedures. This is particularly convenient together with macros.

To actually execute the code, you need some kind of procedure application. These always inherit from `procdef.ProcApp` and add bindings. The `bind` element lets you give python expressions for all names defined using `par` in the `setup` child of the `ProcDef` given in the `procDef` attribute. You can also define just a procedure application without a `procDef` by giving `setup` and `code`.

Procedure application have "types" -- these give where they can be used. In particular, the type determines the signature of the python callable that the procedure application is compiled into. `procdef.ProcApp` has no type, and thus is "abstract"; it should never be a child factory of any `StructAttribute`.

Instead, inherit from it and give

- `name_` -- the element name, as always in structures. This is "apply" for rowmaker applies, "rowfilter" for grammar rowfilters, etc
- `formalArgs` -- a python argument list that gives the arguments of the callable a `ProcApp` of this type is compiled into. Thus, this defines the signature.
- `requiredType` -- a type name that specifies what kind of `ProcDef` the application will accept. This will in general be the same as `name_`. None would mean accept all, which probably is useless.

So, all you need to do to define a new sort of `ProcApp` is write something like:

```
class EmbeddedIterator(rscdef.ProcApp):
    name_ = "iterator"
    formalArgs = "self"
```

(of course, here, documentation as to what the code is supposed to do is particularly important, so don't leave out the docstring when actually doing anything.

Then, you could have:

```
_iterator = base.StructAttribute("iterator", default=base.Undefined,
    childFactory=EmbeddedIterator,
    description="Code yielding row dictionaries", copyable=True)
```

in some structure. To produce something you can execute, then say:

```
theIterator = self.iterator.compile()
for row in theIterator(self):
    print row
```

or somesuch.

## ADQL User Defined Functions

ADQL user defined functions currently all live in `adql.ufunctions`, and their tests are centralised in `ufunctest`. We should probably have a canonical place where individual operators can add them from reliably.

To write a UDF, write a function matching the signature explained in `adql.ufunctions.userFunction` and apply that decorator. For the names, you *must* use something starting with either `gavo_` or `ivo_` as per ADQL 2.1 (where you can only use `ivo_` if you've got someone else also implementing it).

If you can, produce nodes and raise a `ReplaceNode` (as in, e.g., `gavo_transform`). Most existing UDFs admittedly return strings, which leads to lousy tree annotation and makes it impossible to later morph the result – but I'll give you it's much simpler to write functions returning fixed strings. If you do that, be sure to never include unparsed literals; remember: `args` is quite strongly under user control. Hence, make it a habit of always writing `notes.flatten(args[n])` whenever you use `args`. This also has the advantage that expressions in the arguments of your UDFs will be flattened, too.

UDFs will typically start their existence as `gavo_whatever`. If, as is rather common, this later becomes an interoperable UDF (i.e., listed in the UDF catalogue), this should then become `ivo_whatever`. However, existing queries using the `gavo`-prefixed forms should keep working. To make that happen, do:

- pass the `ivo_` version as `userFunction`'s first argument
- add the `gavo_` version in the `additionalNames` keyword argument.

Users will no longer see the `gavo_` version in the capabilities, and hence TOPCAT will mark a syntax error if folks type the old name. I'd say that's ok. People should change their queries after all.

There's an example for that migration technique in `ufunctions._ivo_histogram`. In that case, there is the additional complication that that's just handing through to a SQL function created in `//adql`; this is still called `gavo_histogram`, and to unify the names this is using a custom node (which is convenient here because we want to fiddle with the node's annotation anyway).

## Schema updates

If you need to change the on-disk schema, you must provide an updater in `gavo.user.upgrade`. See the docstring on `Upgrade` on what you can and should do in there, and read on.

The basic idea is that each upgrade step is written as a class inheriting `Upgrader`. Its `version` attribute must be the value of `upgrader.CURRENT_SCHEMAVERSION` (defined near the top) when you start working. After you have defined your upgrader, increase `CURRENT_SCHEMAVERSION` by one.

The upgrader has attributes and class methods with magic names; if these are string-values, they are directly executed, if they are methods, they are called with a connection argument. Do not use any other connections in upgraders or you'll break the atomicity of the upgrades.

The magic names can either be `u_<nn>_<name>` or `s_<nn>_<name>`. use `nn` to determine the action sequence. The difference between `u` and `s` is that when upgrading over multiple versions, all `s` methods are being executed before the first `u` method. The idea is that schema-changing changes should be in `s` methods, content updates and similar should be in `u` methods.

When defining upgrades, it pays to make sure upgraders don't break if what they're doing has been done already; that reduces the requirements on upgrade atomicity and prevents upgrade crashes (which are always ugly) when people do odd things. For this, use the `relationExists(tableName, connection)` and `getColumnNamesFor(rdTableId, connection)` methods to figure things out. There's also `_updateTAP_SCHEMA(connection)`, which you should use on anything influencing `TAP_SCHEMA` (which might include, say, changes in column serialisation).

At Heidelberg, once an upgrade is defined, test the upgrader using:

```
testgavo upgrade
```

The effects should be visible in the `dachstest` database. To be able to roll back changes effected by the upgrade, you may want to backup the cluster first. Markus has a script `backup_postgres` for that.

If you follow the rules, upgrade should be atomic, i.e., either the upgrade succeeds or the database is untouched, letting operators downgrade and continue operations until a problem is figured out. To selectively re-run upgraders (and they should be idempotent), use `dachs upgrade's --force-dbversion` option.

## XML Schema updates

According to the new schema versioning policy of the VO, for minor updates, the target namespace of XSD does not change any more. Still, the file names are versioned upstream.

DaCHS, in general, only knows one version of each schema. Therefore, we remove the version names, and when there's a new schema, you just overwrite the corresponding schema.

There are a few exceptions; in particular, because several minor versions of `VOTable` have been out there and in common use, we keep the schemas for `VOTable 1.1` and `1.2` around with their own custom prefixes.

In case you actually need a new schema file, this is what you need to do:

- install the new file in `gavo/resources/schemata` and `svn add` it
- in `gavo/helpers/testtricks.py`, locate `VO_SCHEMATA` and add your new file name.
- add a `registerPrefix` call for your new schema, allocating a new prefix for it while doing so. See `registry/model.py` for a large selection of such `registerPrefix` calls; but do the registration wherever the schema actually belongs.
- write a unit test exercising your schema.

## Schema Evolution

DaCHS sometimes prototypes new schema elements years before there's any chance to get them into official VO schemas. Many validators ignore `schemaLocation`, and so it's quite likely that DaCHS services would count as invalid for years.

Where schemas have built-in extensibility (e.g., Registry's capabilities), there's the DaCHS schema (mapped in `registry.model.DaFut`) where you can keep mirrors of types and elements. The idea is that you manually copy your new XSD into `resources/schemata/DaCHS.xsd` and the corresponding element declarations to the `DaFut` class. Before the upstream schema is updated as part of PR, you take your elements from `DaFut` in your code, after that, from whatever namespace object the things end up in.

I'd say things should disappear from `DaCHS.xsd` perhaps four years after they've gone official; people not updating their software for four years in a row deserve to have them go invalid.

## Javascript

While it's our goal to let people operate the web-based part of DaCHS without javascript enabled, it's ok if fancier functionality depends on javascript.

After some hesitation, we decided to use the jquery javascript library (we used to have MochiKit but left that when we wanted nice in-browser plotting; so, if you still see MochiKit somewhere, please disregard). We also include some of jquery-ui.

We keep all javascript in "full" source form (in `resources/web/js`). DaCHS performs on-the-fly minimisation (unless `[web]jsSource` is `False`).

For development of that, it's much more convenient if the stuff that gets served out is in source. To enable that, set `[web]jsSource` to `true`. This needs actual code support; right now this only works for files served out in `commonhead`. You need to restart the server for the setting to take effect.

### **gavo.js**

The `commonhead` renderer that's applied to almost all pages pulls in the javascript from `resources/web/js/gavo.js`. This includes some utility functions in the global namespace (and some that should be moved elsewhere). In particular, it contains quite a bit of ugly mess for managing the output formats.

Here's a discussion of some features that may be interesting to template authors.

## Built-in templating

There's a very plain templating engine in javascript included, using an idea due to John Resig, <http://ejohn.org/>. According to this, you define a template in your HTML as a script of type text/html:

```
<script type="text/html" id="tpl_authorHeader">
  <li>
    <a class="arrow-e"
      onclick="toggleAuthorResources(this)" name="$author"/>
      $author ($nummatch)
    </li>
</script>
```

The \$varName parts can then be filled – properly HTML-escaped – by calling:

```
renderTemplate("tpl_authorHeader", {
  author: 'Thor, A. U',
  nummatch: 8})
```

Currently, filling variables is the only thing the engine knows how to do.

## Fairly Simple Tabs

There's built-in javascript and CSS for switching tabs. The tabs require Javascript, so you'll usually want to hide them from non-JS-browsers. Thus, to define the tabs, do something along the lines of:

```
<script type="text/html" id="tabbar_store">
<ul id="tabset_tabs">
  <li class="selected"><a name="by-title">By Title</a></li>
  <li><a name="by-subject">By Subject</a></li>
  <li><a name="by-author">By Author</a></li>
</ul>
</script>

<p id="tab_placeholder" style="border:2pt dashed #bb9999;padding: 0.5ex">
  Enable Javascript for more choices.</p>
```

Note how the tab headings are within a elements that have a name – it's this name that lends identity to them. You could have hrefs for better non-javascript fallback if you have the tabs without javascript; remove the href attributes when you have javascript active, though.

Then, in your javascript, say:

```
$(document).ready(function() {
  $("#tab_placeholder").replaceWith(
    $(document.getElementById("tabbar_store").innerHTML));
  $("#tabset_tabs li").bind("click", makeTabCallback({
    'by-subject': func1,
    'by-author': func2,
    'by-title': func3,
  }));
});
```

(or do something equivalent, if you don't like the innerHTML here). The functions in the dictionary passed to makeTabCallback must then work on the container below the tabs. Here's CSS you could base the container css on:

```
position: relative;
background-color: #EAEBEE;
margin-top: 0px;
min-height: 70ex;
```

The CSS that styles the tabs is in `resources/web/css/gavo_dc.css`, the images necessary in `resources/web/img`.

### **samp.js**

This is Mark Taylor's `samp.js`, checked out from <https://github.com/astrojs/sampjs.git>.

### **jquery and flot**

We're distributing both `jquery` and `flot` in our tarballs because they're rather painful to fiddle together on non-Debian platforms. However, the Debian package doesn't carry them because re-distributing packaged stuff is being frowned upon (and it stinks).

To keep things in sync as good as we can, we need to update the built-in javascript files as we go to a new Debian. To do that, go to `gavo/resources/web/js` in a checkout and run:

```
sudo apt install libjs-jquery libjs-jquery-flot
python3 ../../../../web/ifpages.py
```

This will re-write the two files `jquery-gavo.js` and `jquery.flot.js` based on what Debian currently distributes.

### **Stuff in gavo.imp**

`gavo.imp` has some external dependencies of DaCHS. Shortly after release 1.0, many were dropped in favour of their packaged/native counterparts (`argparse`, `pyarsing`...). What's currently left is:

- `rjsmin` -- Debian packaged, and the Debian version will be picked up automatically if installed (so, this should not go into the Debian package)

### **Different Database Backends**

A request we get fairly regularly is to make DaCHS work with database engines other than Postgres, with MySQL and Oracle being the most popular alternatives for external requests and SQLite something we personally would like to see for ease of deployment. The short answer to all this: It's tricky. You might get away with using [foreign data wrappers](#) in some cases; a group at Paris Observatory reports fairly good results with them.

Here's the longer answer: DaCHS does a lot of inspection of the database, while at the same time worrying about different access levels, reconnection on database restarts, and similar; it also creates extension types. We are not aware of any abstraction layer that would let us keep all this code generic, and that's why we let DaCHS slide into a fairly deep entanglement with `psycpg2` and Postgres.

Seeing such an entanglement reduces the scope of DaCHS, we'd certainly help pulling it out of the entanglement. We probably won't do it ourselves. Here's a list of things

that would need to be done for un-entanglement, that's probably somewhat incomplete and also contains some project mines (innocuous-looking things that blow up into a lot of refactoring once you step on them):

- (a) separate what's specific to postgresql+psycopg2 from sqlsupport, put that into a module (backend\_postgres, say), devise some sort of dispatcher to backends, and have, to work out things, a second backend, that would then contain different implementations for tableExists, indexExists, and so on. Actually, throwing out some cruft from sqlsupport that should have gone ages ago would be a good thing, too.
- (b) figure out what other hidden dependencies exist; the most worrisome part probably is the extension types DaCHS uses and registers as well as the pgSphere interface; this is built into typesystems and used left and right. If there's no way to hide DB-specific differences, there'll have to be some major redesign. Also, DaCHS implicitly assumes TEXT in the database is cheap. If that's not true of a DB (and I think in Oracle TEXT can't be properly indexed) and you'll want much more VARCHARs and similar, minor adjustments might be in order.
- (c) The ADQL translator would need to get another "morpher" (the thing that turns ADQL parse trees into the language of the backend database) That's already foreseen, but figuring out how to enable maximum reuse of code between the different morphers might take some thought. Also, again the question of spherical geometry in the backend will have to be looked at.
- (d) Some mixins directly depend on postgres features (//scs#q3cindex is an obvious example). I believe it'd be ok to say "well, don't use these on non-Postgres", and we'd provide similar things for the other DBs. But that would make RDs non-portable, which I don't like too much either.
- (e) The C boosters generate material for Postgres binary copy. Obviously, one would need to figure out the analogon on other databases (which may not be well-documented; I had to check the Postgres source for some details, too) and then split up boosterskel.c into generic and postgres-specific parts. Or there'd be no support for C boosters on different databases, which might not be unreasonable, either.

## Writing Documentation

Documentation on DaCHS is maintained in ReStructuredText format with some minor extensions (see below). While there's documentation in the tarball and the main SVN, in order to encourage external contributions (including, but not restricted to, typo fixes and the like) the main copy now is at <https://github.com/gavodachs/dachs-doc.git>.

When authoring, you can use some extra RST features (the price: stock rst2pdf and friends don't work properly; use `dachs gendoc latex` or `dachs gendoc html`, or a special sphinx configuration). These include:

- `dachsref`: Give a reference documentation heading (as in [The //ob-score#publishSIAP Mixin](#)), and you'll get a link there.

- dachsdoc: Works like normal links, including explicit targets, but it prepends the root URL of the DaCHS documentation
- bibcode: Adds an ADS link to a bibcode

## Random Stuff

### Tracing imports

Sometimes it's nice to see what gets imported when. Futzing with PEP 302-style import hooks is a pain, and indeed a simple shell line produces more useful output than naive hooks:

```
strace dachs imp -h 2>&1 | grep 'open' | grep -v ENOENT | grep -v "pyc" | sed -e 's/.*"\(.*\)".*/\1/'
```

### matplotlib

To use matplotlib and pyplot within renderers or some other server context, use the following import pattern:

```
import matplotlib
matplotlib.use("Agg")
from matplotlib import pyplot
```

It is crucial that the use("Agg") happens before the import of pyplot. If you fail to do this properly, your code will fail complaining about missing DISPLAYs.

I *guess* we'll soon properly depend on matplotlib and to that initialization in a good place in utils, but don't hold your breath.

### Making a New Version of VOTable the Default

The default VOTable is currently encoded in too many places. Until we clean that up, here's what you need to do when making new VOTable version default (assuming the namespace stays constant, as it should).

- Pull the new schema into resources/schemata. As long as the namespace is constant, you can drop the pervious version.
- In formats/votablewrite.py: Change the default in VOTableContext's constructor. Check the predefined formats at the foot if anything should be updated there; as a rule, I'd suggest there's no reason to define a format for the old version; 1.1 and 1.2 are special cases because we got some things pretty wrong for them.
- In formats/votablewrite.py's makeVOTable: See that the new version is mapped to V.VOTABLE; you probably want to reject attempts to generate the previous version.
- In votable/model.py, add the new version to the NAMESPACES mapping.
- In votable/model.py, change the schemaURL in the registerPrefix line (just overwrite the previous value; you've dropped the old schema above).
- In votable/model.py, in the VOTABLE element definition, change the version attribute (it's still overwritten for the legacy versions).



- Then run tests; a few actually test for the the VOTable version spit out; fix these. All others shouldn't be affected.
- You may also want to change the declarations of the FORMAT parameters in //pql and //soda, and the corresponding key in GETDATA\_FORMATS in sdm.py; but that should only be necessary if there's experimental formats around.

### **Parsing Text Files**

See `utils.iterSimpleText`.

### **Licensing**

<https://matija.suklje.name/how-and-why-to-properly-write-copyright-statements-in-your-code> sounds rather knowledgeable and sensible. Let's put it in next year.