

---

# **Pimlico Documentation**

***Release 0.7***

**Mark Granroth-Wilding**

**Dec 01, 2017**



---

## Contents

---

<b>1</b>	<b>Contents</b>	<b>3</b>
	<b>Python Module Index</b>	<b>163</b>



The **Pimlico Processing Toolkit** is a toolkit for building pipelines of tasks for **processing large datasets** (corpora). It is especially focussed on processing linguistic corpora and provides wrappers around many existing, widely used **NLP** (Natural Language Processing) tools.

It makes it easy to write large, potentially complex pipelines with the following key goals:

- to provide **clear documentation** of what has been done;
- to make it easy to **incorporate standard NLP tasks**,
- and to extend the code with **non-standard tasks, specific to a pipeline**;
- to support simple **distribution of code** for reproduction, for example, on other datasets.

The toolkit takes care of managing data between the steps of a pipeline and checking that everything's executed in the right order.

The core toolkit is written in Python. Pimlico is open source, released under the GPLv3 license. It is available from [its Github repository](#). To get started with a Pimlico project, follow the [getting-started guide](#).

Pimlico is short for *P*ipelined *M*odular *L*inguistic *C*Orpus *p*rocessing.

More NLP tools will gradually be added. See [my wishlist](#) for current plans.



## 1.1 Pimlico guides

Step-by-step guides through common tasks while using Pimlico.

### 1.1.1 Super-quick Pimlico setup

This is a very quick walk-through of the process of starting a new project using Pimlico. For more details, explanations, etc see *the longer getting-started guide*.

First, make sure Python is installed.

#### System-wide configuration

Choose a location on your file system where Pimlico will store all the output from pipeline modules. For example, `/home/me/.pimlico_store/`.

Create a file in your home directory called `.pimlico` that looks like this:

```
long_term_store=/home/me/.pimlico_store
short_term_store=/home/me/.pimlico_store
```

This is not specific to a pipeline: separate pipelines use separate subdirectories.

#### Set up new project

Create a new, empty directory to put your project in. E.g.:

```
cd ~
mkdir myproject
```

Download `newproject.py` into this directory and run it:

```
wget https://raw.githubusercontent.com/markgw/pimlico/master/admin/newproject.py
python newproject.py myproject
```

This fetches the latest Pimlico codebase (in `pimlico/`) and creates a template pipeline (`myproject.conf`).

### Customizing the pipeline

You've got a basic pipeline config file now (`myproject.conf`).

Add sections to it to configure modules that make up your pipeline.

For guides to doing that, see the [the longer setup guide](#) and individual module documentation.

### Running Pimlico

Check the pipeline can be loaded and take a look at the list of modules you've configured:

```
./pimlico.sh myproject.conf status
```

Tell the modules to fetch all the dependencies you need:

```
./pimlico.sh myproject.conf install all
```

If there's anything that can't be installed automatically, this should output instructions for manual installation.

Check the pipeline's ready to run a module that you want to run:

```
./pimlico.sh myproject.conf run MODULE --dry-run
```

To run the next unexecuted module in the list, use:

```
./pimlico.sh myproject.conf run
```

## 1.1.2 Setting up a new project using Pimlico

You've decided to use Pimlico to implement a data processing pipeline. So, where do you start?

This guide steps through the basic setup of your project. You don't have to do everything exactly as suggested here, but it's a good starting point and follows Pimlico's recommended procedures. It steps through the setup for a very basic pipeline.

### System-wide configuration

Pimlico needs you to specify certain parameters regarding your local system. In the simplest case, this is just a file in your home directory called `.pimlico`. See [Local configuration](#) for more details.

It needs to know where to put output files as it executes. Settings are given in a config file in your home directory and apply to all Pimlico pipelines you run. Note that Pimlico will make sure that different pipelines don't interfere with each other's output (provided you give them different names).

There are two locations you need to specify: **short-term** and **long-term storage**. For more details, see [Long-term and short-term stores](#).



For a simple setup, these could be just two subdirectories of the same directory. However, it can be useful to distinguish them.

Create a file `~/ .pimlico` that looks like this:

```
long_term_store=/path/to/long-term/store
short_term_store=/path/to/short-term/store
```

Remember, these paths are not specific to a pipeline: all pipelines will use different subdirectories of these ones.

## Getting started with Pimlico

The procedure for starting a new Pimlico project, using the latest release, is very simple.

Create a new, empty directory to put your project in. Download `newproject.py` into the project directory.

Choose a name for your project (e.g. `myproject`) and run:

```
python newproject.py myproject
```

This fetches the latest version of Pimlico (now in the `pimlico/` directory) and creates a basic config file template, which will define your pipeline.

It also retrieves some libraries that Pimlico needs to run. Other libraries required by specific pipeline modules will be installed as necessary when you use the modules.

## Building the pipeline

You've now got a config file in `myproject.conf`. This already includes a `pipeline` section, which gives the basic pipeline setup. It will look something like this:

```
[pipeline]
name=myproject
release=<release number>
python_path=%(project_root)s/src/python
```

The `name` needs to be distinct from any other pipelines that you run &ndash; it's what distinguishes the storage locations.

`release` is the release of Pimlico that you're using: it's automatically set to the latest one, which has been downloaded.

If you later try running the same pipeline with an updated version of Pimlico, it will work fine as long as it's the same major version (the first digit). Otherwise, there may be backwards incompatible changes, so you'd need to update your config file, ensuring it plays nicely with the later Pimlico version.

## Getting input

Now we add our first module to the pipeline. This reads input from XML files and iterates of `<doc>` tags to get documents. This is how the Gigaword corpus is stored, so if you have Gigaword, just set the path to point to it.

---

### Todo

Use a dataset that everyone can get to in the example

---

```
[input-text]
type=pimlico.datatypes.XmlDocumentIterator
path=/path/to/data/dir
```

Perhaps your corpus is very large and you'd rather try out your pipeline on a small subset. In that case, add the following option:

```
truncate=1000
```

---

**Note:** For a neat way to define a small test version of your pipeline and keep its output separate from the main pipeline, see *Pipeline variants*.

---

## Grouping files

The standard approach to storing data between modules in Pimlico is to group them together into batches of documents, storing each batch in a tar archive, containing a file for every document. This works nicely with large corpora, where having every document as a separate file would cause filesystem difficulties and having all documents in the same file would result in a frustratingly large file.

We can do the grouping on the fly as we read data from the input corpus. The `tar_filter` module groups documents together and subsequent modules will all use the same grouping to store their output, making it easy to align the datasets they produce.

```
[tar-grouper]
type=pimlico.modules.corpora.tar_filter
input=input-text
```

## Doing something: tokenization

Now, some actual linguistic processing, albeit somewhat uninteresting. Many NLP tools assume that their input has been divided into sentences and tokenized. The OpenNLP-based tokenization module does both of these things at once, calling OpenNLP tools.

Notice that the output from the previous module feeds into the input for this one, which we specify simply by naming the module.

```
[tokenize]
type=pimlico.modules.opennlp.tokenize
input=tar-grouper
```

## Doing something more interesting: POS tagging

Many NLP tools rely on part-of-speech (POS) tagging. Again, we use OpenNLP, and a standard Pimlico module wraps the OpenNLP tool.

```
[pos-tag]
type=pimlico.modules.opennlp.pos
input=tokenize
```

## Running Pimlico

Now we've got our basic config file ready to go. It's a simple linear pipeline that goes like this:

read input docs -> group into batches -> tokenize -> POS tag

Before we can run it, there's one thing missing: three of these modules have their own dependencies, so we need to get hold of the libraries they use. The input reader uses the Beautiful Soup python library and the tokenization and POS tagging modules use OpenNLP.

## Checking everything's dandy

Now you can run the `status` command to check that the pipeline can be loaded and see the list of modules.

```
./pimlico.sh myproject.conf status
```

To check that specific modules are ready to run, with all software dependencies installed, use the `run` command with `--dry-run` (or `--dry`) switch:

```
./pimlico.sh myproject.conf run tokenize --dry
```

With any luck, all the checks will be successful. There might be some missing software dependencies.

## Fetching dependencies

All the standard modules provide easy ways to get hold of their dependencies automatically, or as close as possible. Most of the time, all you need to do is tell Pimlico to install them.

Use the `run` command, with a module name and `--dry-run`, to check whether a module is ready to run.

```
./pimlico.sh myproject.conf run tokenize --dry
```

In this case, it will tell you that some libraries are missing, but they can be installed automatically. Simply issue the `install` command for the module.

```
./pimlico.sh myproject.conf install tokenize
```

Simple as that.

There's one more thing to do: the tools we're using require statistical models. We can simply download the pre-trained English models from the OpenNLP website.

At present, Pimlico doesn't yet provide a built-in way for the modules to do this, as it does with software libraries, but it does include a GNU Makefile to make it easy to do:

```
cd ~/myproject/pimlico/models
make opennlp
```

Note that the modules we're using default to these standard, pre-trained models, which you're now in a position to use. However, if you want to use different models, e.g. for other languages or domains, you can specify them using extra options in the module definition in your config file.

If there are any other library problems shown up by the dry run, you'll need to address them before going any further.

## Running the pipeline

### What modules to run?

Pimlico suggests an order in which to run your modules. In our case, this is pretty obvious, seeing as our pipeline is entirely linear – it’s clear which ones need to be run before others.

```
./pimlico.sh myproject.conf status
```

The output also tells you the current status of each module. At the moment, all the modules are `UNEXECUTED`.

You’ll notice that the `tar-grouper` module doesn’t feature in the list. This is because it’s a filter – it’s run on the fly while reading output from the previous module (i.e. the input), so doesn’t have anything to run itself.

You might be surprised to see that `input-text` *does* feature in the list. This is because, although it just reads the data out of a corpus on disk, there’s not quite enough information in the corpus, so we need to run the module to collect a little bit of metadata from an initial pass over the corpus. Some input types need this, others not. In this case, all we’re lacking is a count of the total number of documents in the corpus.

---

**Note:** To make running your pipeline even simpler, you can abbreviate the command by using a **shebang** in the config file. Add a line at the top of `myproject.conf` like this:

```
#!/pimlico.sh
```

Then make the conf file executable by running (on Linux):

```
chmod ug+x myproject.conf
```

Now you can run Pimlico for your pipeline by using the config file as an executable command:

```
./myproject.conf status
```

---

## Running the modules

The modules can be run using the `run` command and specifying the module by name. We do this manually for each module.

```
./pimlico.sh myproject.conf run input-text
./pimlico.sh myproject.conf run tokenize
./pimlico.sh myproject.conf run pos-tag
```

## Adding custom modules

Most likely, for your project you need to do some processing not covered by the built-in Pimlico modules. At this point, you can start implementing your own modules, which you can distribute along with the config file so that people can replicate what you did.

The `newproject.py` script has already created a directory where our custom source code will live: `src/python`, with some subdirectories according to the standard code layout, with module types and datatypes in separate packages.

The template pipeline also already has an option `python_path` pointing to this directory, so that Pimlico knows where to find your code. Note that the code’s in a subdirectory of that containing the pipeline config and we specify the custom code path relative to the config file, so it’s easy to distribute the two together.

Now you can create Python modules or packages in `src/python`, following the same conventions as the built-in modules and overriding the standard base classes, as they do. The following articles tell you more about how to do this:

- [Writing Pimlico modules](#)
- [Writing document map modules](#)
- [Pimlico module structure](#)

Your custom modules and datatypes can then simply be used in the config file as module types.

### 1.1.3 Writing Pimlico modules

Pimlico comes with a fairly large number of *module types* that you can use to run many standard NLP, data processing and ML tools over your datasets.

For some projects, this is all you need to do. However, often you'll want to mix standard tools with your own code, for example, using the output from the tools. And, of course, there are many more tools you might want to run that aren't built into Pimlico: you can still benefit from Pimlico's framework for data handling, config files and so on.

For a detailed description of the structure of a Pimlico module, see [Pimlico module structure](#). This guide takes you through building a simple module.

---

**Note:** In any case where a module will process a corpus one document at a time, you should write a *document map module*, which takes care of a lot of things for you, so you only need to say what to do with each document.

---

#### Code layout

If you've followed the [basic project setup guide](#), you'll have a project with a directory structure like this:

```
myproject/
  pipeline.conf
  pimlico/
    bin/
    lib/
    src/
    ...
  src/
    python/
```

If you've not already created the `src/python` directory, do that now.

This is where your custom Python code will live. You can put all of your custom module types and datatypes in there and use them in the same way as you use the Pimlico core modules and datatypes.

Add this option to the `[pipeline]` section of your config file, so Pimlico knows where to find your code:

```
python_path=src/python
```

To follow the conventions used in Pimlico's codebase, we'll create the following package structure in `src/python`:

```
src/python/myproject/
  __init__.py
  modules/
    __init__.py
```

```
datatypes/  
    __init__.py
```

## Write a module

A Pimlico module consists of a Python package with a special layout. Every module has a file `info.py`. This contains the definition of the module's metadata: its inputs, outputs, options, etc.

Most modules also have a file `execute.py`, which defines the routine that's called when it's run. You should take care when writing `info.py` not to import any non-standard Python libraries or have any time-consuming operations that get run when it gets imported.

`execute.py`, on the other hand, will only get imported when the module is to be run, after dependency checks.

For the example below, let's assume we're writing a module called `nmf` and create the following directory structure for it:

```
src/python/myproject/modules/  
    __init__.py  
    nmf/  
        __init__.py  
        info.py  
        execute.py
```

## Easy start

To help you get started, Pimlico provides a wizard in the `newmodule` command.

This will ask you a series of questions, guiding you through the most common tasks in creating a new module. At the end, it will generate a template to get you started with your module's code. You then just need to fill in the gaps and write the code for what the module actually does.

Read on to learn more about the structure of modules, including things not covered by the wizard.

## Metadata

Module metadata (everything apart from what happens when it's actually run) is defined in `info.py` as a class called `ModuleInfo`.

Here's a sample basic `ModuleInfo`, which we'll step through. (It's based on the Scikit-learn `matrix_factorization` module.)

```
from pimlico.core.dependencies.python import PythonPackageOnPip  
from pimlico.core.modules.base import BaseModuleInfo  
from pimlico.datatypes.arrays import ScipySparseMatrix, NumpyArray  
  
class ModuleInfo(BaseModuleInfo):  
    module_type_name = "nmf"  
    module_readable_name = "Sklearn non-negative matrix factorization"  
    module_inputs = [("matrix", ScipySparseMatrix)]  
    module_outputs = [("w", NumpyArray), ("h", NumpyArray)]  
    module_options = {  
        "components": {  
            "help": "Number of components to use for hidden representation",
```

```

        "type": int,
        "default": 200,
    },
}

def get_software_dependencies(self):
    return super(ModuleInfo, self).get_software_dependencies() + \
        [PythonPackageOnPip("sklearn", "Scikit-learn")]

```

The `ModuleInfo` should always be a subclass of `BaseModuleInfo`. There are some subclasses that you might want to use instead (e.g., see [Writing document map modules](#)), but here we just use the basic one.

Certain class-level attributes should pretty much always be overridden:

- `module_type_name`: A name used to identify the module internally
- `module_readable_name`: A human-readable short description of the module
- `module_inputs`: Most modules need to take input from another module (though not all)
- `module_outputs`: Describes the outputs that the module will produce, which may then be used as inputs to another module

**Inputs** are given as pairs `(name, type)`, where `name` is a short name to identify the input and `type` is the datatype that the input is expected to have. Here, and most commonly, this is a subclass of `PimlicoDatatype` and Pimlico will check that a dataset supplied for this input is either of this type, or has a type that is a subclass of this.

Here we take just a single input: a sparse matrix.

**Outputs** are given in a similar way. It is up to the module's executor (see below) to ensure that these outputs get written, but here we describe the datatypes that will be produced, so that we can use them as input to other modules.

Here we produce two Numpy arrays, the factorization of the input matrix.

**Dependencies:** Since we require Scikit-learn to execute this module, we override `get_software_dependencies()` to specify this. As Scikit-learn is available through Pip, this is very easy: all we need to do is specify the Pip package name. Pimlico will check that Scikit-learn is installed before executing the module and, if not, allow it to be installed automatically.

Finally, we also define some **options**. The values for these can be specified in the pipeline config file. When the `ModuleInfo` is instantiated, the processed options will be available in its `options` attribute. So, for example, we can get the number of components (specified in the config file, or the default of 200) using `info.options["components"]`.

## Executor

Here is a sample executor for the module info given above, placed in the file `execute.py`.

```

from pimlico.core.modules.base import BaseModuleExecutor
from pimlico.datatypes.arrays import NumpyArrayWriter
from sklearn.decomposition import NMF

class ModuleExecutor(BaseModuleExecutor):
    def execute(self):
        input_matrix = self.info.get_input("matrix").array
        self.log.info("Loaded input matrix: %s" % str(input_matrix.shape))

        # Convert input matrix to CSR
        input_matrix = input_matrix.tocsr()

```

```
# Initialize the transformation
components = self.info.options["components"]
self.log.info("Initializing NMF with %d components" % components)
nmf = NMF(components)

# Apply transformation to the matrix
self.log.info("Fitting NMF transformation on input matrix" % transform_type)
transformed_matrix = transformer.fit_transform(input_matrix)

self.log.info("Fitting complete: storing H and W matrices")
# Use built-in Numpy array writers to output results in an appropriate format
with NumpyArrayWriter(self.info.get_absolute_output_dir("w")) as w_writer:
    w_writer.set_array(transformed_matrix)
with NumpyArrayWriter(self.info.get_absolute_output_dir("h")) as h_writer:
    h_writer.set_array(transformer.components_)
```

The executor is always defined as a class in `execute.py` called `ModuleExecutor`. It should always be a subclass of `BaseModuleExecutor` (though, again, note that there are more specific subclasses and class factories that we might want to use in other circumstances).

The `execute()` method defines what happens when the module is executed.

The instance of the module's `ModuleInfo`, complete with **options** from the pipeline config, is available as `self.info`. A standard Python **logger** is also available, as `self.log`, and should be used to keep the user updated on what's going on.

Getting hold of the **input data** is done through the module info's `get_input()` method. In the case of a Scipy matrix, here, it just provides us with the matrix as an attribute.

Then we do whatever our module is designed to do. At the end, we write the output data to the appropriate output directory. This should always be obtained using the `get_absolute_output_dir()` method of the module info, since Pimlico takes care of the exact location for you.

Most Pimlico datatypes provide a corresponding **writer**, ensuring that the output is written in the correct format for it to be read by the datatype's reader. When we leave the `with` block, in which we give the writer the data it needs, this output is written to disk.

## Pipeline config

Our module is now ready to use and we can refer to it in a pipeline config file. We'll assume we've prepared a suitable Scipy sparse matrix earlier in the pipeline, available as the default output of a module called `matrix`. Then we can add section like this to use our new module:

```
[matrix]
...(Produces sparse matrix output)...

[factorize]
type=myproject.modules.nmf
components=300
input=matrix
```

Note that, since there's only one input, we don't need to give its name. If we had defined multiple inputs, we'd need to specify this one as `input_matrix=matrix`.

You can now run the module as part of your pipeline in the usual ways.



## Skeleton new module

To make developing a new module a little quicker, here's a skeleton module info and executor.

```
from pimlico.core.modules.base import BaseModuleInfo

class ModuleInfo(BaseModuleInfo):
    module_type_name = "NAME"
    module_readable_name = "READABLE NAME"
    module_inputs = [("NAME", REQUIRED_TYPE)]
    module_outputs = [("NAME", PRODUCED_TYPE)]
    # Delete module_options if you don't need any
    module_options = {
        "OPTION_NAME": {
            "help": "DESCRIPTION",
            "type": TYPE,
            "default": VALUE,
        },
    }

    def get_software_dependencies(self):
        return super(ModuleInfo, self).get_software_dependencies() + [
            # Add your own dependencies to this list
            # Remove this method if you don't need to add any
        ]
```

```
from pimlico.core.modules.base import BaseModuleExecutor

class ModuleExecutor(BaseModuleExecutor):
    def execute(self):
        input_data = self.info.get_input("NAME")
        self.log.info("MESSAGES")

        # DO STUFF

        with SOME_WRITER(self.info.get_absolute_output_dir("NAME")) as writer:
            # Do what the writer requires
```

## 1.1.4 Writing document map modules

### Todo

Write a guide to building document map modules.

For now, the skeletons below are a useful starting point, but there should be a more fulsome explanation here of what document map modules are all about and how to use them.

### Skeleton new module

To make developing a new module a little quicker, here's a skeleton module info and executor for a document map module. It follows the most common method for defining the executor, which is to use the multiprocessing-based executor factory.

```
from pimlico.core.modules.map import DocumentMapModuleInfo
from pimlico.datatypes.tar import TarredCorpusType

class ModuleInfo(DocumentMapModuleInfo):
    module_type_name = "NAME"
    module_readable_name = "READABLE NAME"
    module_inputs = [("NAME", TarredCorpusType(DOCUMENT_TYPE))]
    module_outputs = [("NAME", PRODUCED_TYPE)]
    module_options = {
        "OPTION_NAME": {
            "help": "DESCRIPTION",
            "type": TYPE,
            "default": VALUE,
        },
    }

    def get_software_dependencies(self):
        return super(ModuleInfo, self).get_software_dependencies() + [
            # Add your own dependencies to this list
        ]

    def get_writer(self, output_name, output_dir, append=False):
        if output_name == "NAME":
            # Instantiate a writer for this output, using the given output dir
            # and passing append in as a kwarg
            return WRITER_CLASS(output_dir, append=append)
```

A bare-bones executor:

```
from pimlico.core.modules.map.multiproc import multiprocessing_executor_factory

def process_document(worker, archive_name, doc_name, *data):
    # Do something to process the document...

    # Return an object to send to the writer
    return output

ModuleExecutor = multiprocessing_executor_factory(process_document)
```

Or getting slightly more sophisticated:

```
from pimlico.core.modules.map.multiproc import multiprocessing_executor_factory

def process_document(worker, archive_name, doc_name, *data):
    # Do something to process the document

    # Return a tuple of objects to send to each writer
    # If you only defined a single output, you can just return a single object
    return output1, output2, ...

# You don't have to, but you can also define pre- and postprocessing
# both at the executor level and worker level

def preprocess(executor):
```

```

    pass

def postprocess(executor, error=None):
    pass

def set_up_worker(worker):
    pass

def tear_down_worker(worker, error=None):
    pass

ModuleExecutor = multiprocessing_executor_factory(
    process_document,
    preprocess_fn=preprocess, postprocess_fn=postprocess,
    worker_set_up_fn=set_up_worker, worker_tear_down_fn=tear_down_worker,
)

```

### 1.1.5 Filter modules

Filter modules appear in pipeline config, but never get executed directly, instead producing their output on the fly when it is needed.

There are two types of filter modules in Pimlico:

- All *document map modules* can be used as filters.
- Other modules may be defined in such a way that they always function as filters.

#### Using document map modules as filters

See [this guide](#) for how to create document map modules, which process each document in an input iterable corpus, producing one document in the output corpus for each. Many of the core Pimlico modules are document map modules.

Any document map module can be used as a filter simply by specifying `filter=True` in its options. It will then not appear in the module execution schedule (output by the `status` command), but will get executed on the fly by any module that uses its output. It will be initialized when the downstream module starts accessing the output, and then the single-document processing routine will be run on each document to produce the corresponding output document as the downstream module iterates over the corpus.

It is possible to chain together filter modules in sequence.

#### Other filter modules

A module can be defined so that it always functions as a filter by setting `module_executable=False` on its module-info class. Pimlico will assume that its outputs are ready as soon as its inputs are ready and will not try to execute it. The module developer must ensure that the outputs get produced when necessary.

This form of filter is typically appropriate for very simple transformations of data. For example, it might perform a simple conversion of one datatype into another to allow the output of a module to be used as if it had a different datatype. However, it is possible to do more sophisticated processing in a filter module, though the implementation is a little more tricky (*tar\_filter* is an example of this).

## Defining

Define a filter module something like this:

```
class ModuleInfo(BaseModuleInfo):
    module_type_name = "my_module_name"
    module_executable = False # This is the crucial instruction to treat this as a
    ↪filter
    module_inputs = []        # Define inputs
    module_outputs = []       # Define at least one output, which we'll produce as
    ↪needed
    module_options = {}       # Any options you need

    def instantiate_output_datatype(self, output_name, output_datatype, **kwargs):
        # Here we produce the desired output datatype,
        # using the inputs acquired from self.get_input(name)
        return MyOutputDatatype()
```

You don't need to create an `execute.py`, since it's not executable, so Pimlico will not try to load a module executor. Any processing you need to do should be put inside the datatype, so that it's performed when the datatype is used (e.g. when iterating over it), but not when `instantiate_output_datatype()` is called or when the datatype is instantiated, as these happen every time the pipeline is loaded.

A trick that can be useful to wrap up functionality in a filter datatype is to define a new datatype that does the necessary processing on the fly and to set its class attribute `emulated_datatype` to point to a datatype class that should be used instead for the purposes of type checking. The built-in `tar_filter` module uses this trick.

Either way, you should **take care with imports**. Remember that the `execute.py` of executable modules is only imported when a module is to be run, meaning that we can load the pipeline config without importing any dependencies needed to run the module. If you put processing in a specially defined datatype class that has dependencies, make sure that they're not imported at the top of `info.py`, but only when the datatype is used.

### 1.1.6 Multistage modules

Multistage modules are used to encapsulate a module than is executed in several consecutive runs. You can think of each stage as being its own module, but where the whole sequence of modules is always executed together. The multistage module simply chains together these individual modules so that you only include a single module instance in your pipeline definition.

One common example of a use case for multistage modules is where some fairly time-consuming preprocessing needs to be done on an input dataset. If you put all of the processing into a single module, you can end up in an irritating situation where the lengthy data preprocessing succeeds, but something goes wrong in the main execution code. You then fix the problem and have to run all the preprocessing again.

Most obvious solution to this is to separate the preprocessing and main execution into two separate modules. But then, if you want to reuse you module sometime in the future, you have to remember to always put the preprocessing module before the main one in your pipeline (or infer this from the datatypes!). And if you have more than these two modules (say, a sequence of several, or preprocessing of several inputs) this starts to make pipeline development frustrating.

A multistage module groups these internal modules into one logical unit, allowing them to be used together by including a single module instance and also to share parameters.

## Defining a multistage module

### Component stages

The first step in defining a multistage module is to define its individual stages. These are actually defined in exactly the same way as normal modules. (This means that they can also be used separately.)

If you're writing these modules specifically to provide the stages of your multistage module (rather than tying together already existing modules for convenience), you probably want to put them all in subpackages.

For an ordinary module, *we used the directory structure*:

```
src/python/myproject/modules/
    __init__.py
    mymodule/
        __init__.py
        info.py
        execute.py
```

Now, we'll use something like this:

```
src/python/myproject/modules/
    __init__.py
    my_ms_module/
        __init__.py
        info.py
        module1/
            __init__.py
            info.py
            execute.py
        module2/
            __init__.py
            info.py
            execute.py
```

Note that `module1` and `module2` both have the typical structure of a module definition: an `info.py` to define the module-info, and an `execute.py` to define the executor. At the top level, we've just got an `info.py`. It's in here that we'll define the multistage module. We don't need an `execute.py` for that, since it just ties together the other modules, using their executors at execution time.

### Multistage module-info

With our component modules that constitute the stages defined, we now just need to tie them together. We do this by defining a module-info for the multistage module in its `info.py`. Instead of subclassing `BaseModuleInfo`, as usual, we create the `ModuleInfo` class using the factory function `multistage_module()`.

```
ModuleInfo = multistage_module("module_name",
    [
        # Stages to be defined here...
    ]
)
```

In other respects, this module-info works in the same way as usual: it's a class (return by the factory) called `ModuleInfo` in the `info.py`.

`multistage_module()` takes two arguments: a module name (equivalent to the `module_name` attribute of a normal module-info) and a list of instances of `ModuleStage`.

## Connecting inputs and outputs

Connections between the outputs and inputs of the stages work in a very similar way to connections between module instances in a pipeline. The same type checking system is employed and data is passed between the stages (i.e. between consecutive executions) as if the stages were separate modules.

Each stage is defined as an instance of *ModuleStage*:

```
[
    ModuleStage("stage_name", TheModuleInfoClass, connections=[...], output_
↪connections=[...])
]
```

The parameter `connections` defines how the stage's inputs are connected up to either the outputs of previous stages or inputs to the multistage module. Just like in pipeline config files, if no explicit input connections are given, the default input to a stage is connected to the default output from the previous one in the list.

There are two classes you can use to define input connections.

***InternalModuleConnection*** This makes an explicit connection to the output of another stage.

You must specify the name of the input (to this stage) that you're connecting. You may specify the name of the output to connect it to (defaults to the default output). You may also give the name of the stage that the output comes from (defaults to the previous one).

```
[
    ModuleStage("stage1", FirstInfo),
    # FirstInfo has an output called "corpus", which we connect explicitly to the_
↪next stage
    # We could leave out the "corpus" here, if it's the default output from_
↪FirstInfo
    ModuleStage("stage2", SecondInfo, connections=[InternalModuleConnection("data"
↪", "corpus")]),
    # We connect the same output from stage1 to stage3
    ModuleStage("stage3", ThirdInfo, connections=[InternalModuleConnection("data",
↪ "corpus", "stage1")]),
]
```

***ModuleInputConnection***: This makes a connection to an input to the whole multistage module.

Note that you don't have to explicitly define the multistage module's inputs anywhere: you just mark certain inputs to certain stages as coming from outside the multistage module, using this class.

```
[
    ModuleStage("stage1", FirstInfo, [ModuleInputConnection("raw_data")]),
    ModuleStage("stage2", SecondInfo, [InternalModuleConnection("data", "corpus"
↪")]),
    ModuleStage("stage3", ThirdInfo, [InternalModuleConnection("data", "corpus",
↪"stage1")]),
]
```

Here, the module type `FirstInfo` has an input called `raw_data`. We've specified that this needs to come in directly as an input to the multistage module – when we use the multistage module in a pipeline, it must be connected up with some earlier module.

The multistage module's input created by doing this will also have the name `raw_data` (specified using a parameter `input_raw_data` in the config file). You can override this, if you want to use a different name:

```
[
    ModuleStage("stage1", FirstInfo, [ModuleInputConnection("raw_data", "data
↪")]),
    ModuleStage("stage2", SecondInfo, [InternalModuleConnection("data", "corpus
↪")]),
    ModuleStage("stage3", ThirdInfo, [InternalModuleConnection("data", "corpus",
↪"stage1")]),
]
```

This would be necessary if two stages both had inputs called `raw_data`, which you want to come from different data sources. You would then simply connect them to different inputs to the multistage module:

```
[
    ModuleStage("stage1", FirstInfo, [ModuleInputConnection("raw_data", "first_
↪data")]),
    ModuleStage("stage2", SecondInfo, [ModuleInputConnection("raw_data", "second_
↪data")]),
    ModuleStage("stage3", ThirdInfo, [InternalModuleConnection("data", "corpus",
↪"stage1")]),
]
```

Conversely, you might deliberately connect the inputs from two stages to the same input to the multistage module, by using the same multistage input name twice. (Of course, the two stages are not required to have overlapping input names for this to work.) This will result in the multistage just requiring one input, which get used by both stages.

```
[
    ModuleStage("stage1", FirstInfo,
        [ModuleInputConnection("raw_data", "first_data"), ↪
↪ModuleInputConnection("dict", "vocab")]),
    ModuleStage("stage2", SecondInfo,
        [ModuleInputConnection("raw_data", "second_data"), ↪
↪ModuleInputConnection("vocabulary", "vocab")]),
    ModuleStage("stage3", ThirdInfo, [InternalModuleConnection("data", "corpus",
↪"stage1")]),
]
```

By default, the multistage module has just a single output: the default output of the last stage in the list. You can specify any of the outputs of any of the stages to be provided as an output to the multistage module. Use the `output_connections` parameter when defining the stage.

This parameter should be a list of instances of *ModuleOutputConnection*. Just like with input connections, if you don't specify otherwise, the multistage module's output will have the same name as the output from the stage module. But you can override this when giving the output connection.

```
[
    ModuleStage("stage1", FirstInfo, [ModuleInputConnection("raw_data", "first_data
↪")]),
    ModuleStage("stage2", SecondInfo, [ModuleInputConnection("raw_data", "second_data
↪")]),
        output_connections=[ModuleOutputConnection("model")]), # This ↪
↪output will just be called "model"
    ModuleStage("stage3", ThirdInfo, [InternalModuleConnection("data", "corpus",
↪"stage1"),
        output_connections=[ModuleOutputConnection("model", "stage3_model")]),
]
```

## Module options

The parameters of the multistage module that can be specified when it is used in a pipeline config (those usually defined in the `module_options` attribute) include all of the options to all of the stages. The option names are simply `<stage_name>_<option_name>`.

So, in the above example, if `FirstInfo` has an option called `threshold`, the multistage module will have an option `stage1_threshold`, which gets passed through to `stage1` when it is run.

Often you might wish to specify one parameter to the multistage module that gets used by several stages. Say `stage2` had a `cutoff` parameter and we always wanted to use the same value as the `threshold` for `stage1`. Instead of having to specify `stage1_threshold` and `stage2_cutoff` every time in your config file, you can assign a single name to an option (say `threshold`) for the multistage module, whose value gets passed through to the appropriate options of the stages.

Do this by specifying a dictionary as the `option_connections` parameter to `ModuleStage`, whose keys are names of the stage module type's options and whose values are the new option names for the multistage module that you want to map to those stage options. You can use the same multistage module option name multiple times, which will cause only a single option to be added to the multistage module (using the definition from the first stage), which gets mapped to multiple stage options.

To implement that above example, you would give:

```
[
    ModuleStage("stage1", FirstInfo, [ModuleInputConnection("raw_data", "first_data
↪"),
                                option_connections={"threshold": "threshold"}),
    ModuleStage("stage2", SecondInfo, [ModuleInputConnection("raw_data", "second_data
↪"),
                                [ModuleOutputConnection("model")],
                                option_connections={"cutoff": "threshold"}),
    ModuleStage("stage3", ThirdInfo, [InternalModuleConnection("data", "corpus",
↪"stage1"),
                                [ModuleOutputConnection("model", "stage3_model")]]),
]
```

If you know that the different stages have distinct option name, or that they should always tie their values together where their option names overlap, you can set `use_stage_option_names=True` on the stages. This will cause the stage-name prefix not to be added to the option name when connecting it to the multistage module's option.

You can also force this behaviour for all stages by setting `use_stage_option_names=True` when you call `multistage_module()`. Any explicit option name mappings you provide via `option_connections` will override this.

## Running

To run a multistage module once you've used it in your pipeline config, you run one stage at a time, as if they were separate module instances.

Say we've used the above multistage module in a pipeline like so:

```
[model_train]
type=myproject.modules.my_ms_module
stage1_threshold=10
stage2_cutoff=10
```

The normal way to run this module would be to use the `run` command with the module name:



```
./pimlico.sh mypipeline.conf run model_train
```

If we do this, Pimlico will choose the next unexecuted stage that's ready to run (presumably `stage1` at this point). Once that's done, you can run the same command again to execute `stage2`.

You can also select a specific stage to execute by using the module name `<ms_module_name>:<stage_name>`, e.g. `model_train:stage2`. (Note that `stage2` doesn't actually depend on `stage1`, so it's perfectly plausible that we might want to execute them in a different order.)

If you want to execute multiple stages at once, just use this scheme to specify each of them as a module name for the run command. Remember, Pimlico can take any number of modules and execute them in sequence:

```
./pimlico.sh mypipeline.conf run model_train:stage1 model_train:stage2
```

Or, if you want to execute all of them, you can use the stage name `*` or `all` as a shorthand:

```
./pimlico.sh mypipeline.conf run model_train:all
```

Finally, if you're not sure what stages a multistage module has, use the module name `<ms_module_name>:?.` The run command will then just output a list of stages and exit.

## 1.1.7 Running one pipeline on multiple computers

### Multiple servers

In most of the examples, we've been setting up a pipeline, with a config file, some source code and some data, all on one machine. Then we run each module in turn, checking that it has all the software and data that it needs to run.

But it's not unusual to find yourself needing to process a dataset across different computers. For example, you have access to a server with lots of CPUs and one module in your pipeline would benefit greatly from parallelizing lots of little tasks over them. However, you don't have permission to install software on that server that you need for another module.

This is not a problem: you can simply put your config file and code on both machines. After running one module on one machine, you copy over its output to the place on the other machine where Pimlico expects to find it. Then you're ready to run the next module on the second machine.

Pimlico is designed to handle this situation nicely.

- **It doesn't expect software requirements for all modules to be satisfied before you can run any of them.** Software dependencies are checked only for modules about to be run and the code used to execute a module is not even loaded until you actually run the module.
- **It doesn't require you to execute your pipeline in order.** If the output from a module is available where it's expected to be, you can happily run any modules that take that data as input, even if the pipeline up to that point doesn't appear to have been executed (e.g. if it's been run on another machine).
- **It provides you with tools to make it easier to copy data between machines.** You can easily copy the output data from one module to the appropriate location on another server, so it's ready to be used as input to another module there.

### Copying data between computers

Let's assume you've got your pipeline set up, with identical config files, on two computers: `server_a` and `server_b`. You've run the first module in your pipeline, `module1`, on `server_a` and want to run the next, `module2`, which takes input from `module1`, on `server_b`.

The procedure is as follows:

- **Dump** the data from the pipeline on `server_a`. This packages up the output data for a module in a single file.
- **Copy** the dumped file from `server_a` to `server_b`, in whatever way is most convenient, e.g., using `scp`.
- **Load** the dumped file into the pipeline on `server_b`. This unpacks the data directory for the file and puts it in Pimlico's data directory for the module.

For example, on `server_a`:

```
$ ./pimlico.sh pipeline.conf dump module1
$ scp ~/module1.tar.gz server_b:~/
```

Note that the `dump` command created a `.tar.gz` file in your home directory. If you want to put it somewhere else, use the `--output` option to specify a directory. The file is named after the module that you're dumping.

Now, log into `server_b` and load the data.

```
$ ./pimlico.sh pipeline.conf load ~/module1.tar.gz
```

Now `module1`'s output data is in the right place and ready for use by `module2`.

The `dump` and `load` commands can also process data for multiple modules at once. For example:

```
$ mkdir ~/modules
$ ./pimlico.sh pipeline.conf dump module1 ... module10 --output ~/modules
$ scp -r ~/modules server_b:~/
```

Then on `server_b`:

```
$ ./pimlico.sh pipeline.conf load ~/modules/*
```

## Other issues

Aside from getting data between the servers, there are certain issues that often arise when running a pipeline across multiple servers.

- **Shared Pimlico codebase.** If you share the directory that contains Pimlico's code across servers (e.g. NFS or `rsync`), you can have problems resulting from sharing the libraries it installs. See [instructions for using multiple virtualenvs](#) for the solution.
- **Shared home directory.** If you share your home directory across servers, using the same `.pimlico` local config file might be a problem. See [Local configuration](#) for various possible solutions.

## 1.1.8 Documenting your own code

Pimlico's documentation is produced using [Sphinx](#). The Pimlico codebase includes a tool for generating documentation of Pimlico's built-in modules, including things like a table of the module's available config options and its input and outputs.

You can also use this tool yourself to generate documentation of your own code that uses Pimlico. Typically, you will use in your own project some of Pimlico's built-in modules and some of your own.

Refer to Sphinx's documentation for how to build normal Sphinx documentation – writing your own ReST documents and using the `apidoc` tool to generate API docs. Here we describe how to create a basic Sphinx setup that will generate a reference for your custom Pimlico modules.

It is assumed that you've got a working Pimlico setup and have already successfully written some modules.

## Basic doc setup

Create a `docs` directory in your project root (the directory in which you have `pimlico/` and your own `src/`, etc).

Put a Sphinx `conf.py` in there. You can start from the very basic skeleton [here](#).

You'll also want a `Makefile` to build your docs with. You can use the basic Sphinx one as a starting point. Here's a version of that that already includes an extra target for building your module docs.

Finally, create a root document for your documentation, `index.rst`. This should include a table of contents which includes the generated module docs. You can use [this one](#) as a template.

## Building the module docs

Take a look in the `Makefile` (if you've used our one as a starting point) and set the variables at the top to point to the Python package that contains the Pimlico modules you want to document.

The make target there runs the tool `modulegen` in the Pimlico codebase. Just run, in the `docs/`:

```
make modules
```

You can also do this manually:

```
python -m pimlico.utils.docs.modulegen --path python.path.to.modules modules/
```

(The Pimlico codebase must, of course, be importable. The simplest way to ensure this is to use Pimlico's `python` alias in its `bin/` directory.)

There is now a set of `.rst` files in the `modules/` output directory, which can be built using Sphinx by running `make html`.

Your beautiful docs are now in the `_build/` directory!

## 1.2 Core docs

A set of articles on the core aspects and features of Pimlico.

### 1.2.1 Downloading Pimlico

To start a new project using Pimlico, download the [newproject.py](#) script. It will create a template pipeline config file to get you started and download the latest version of Pimlico to accompany it.

See *Setting up a new project using Pimlico* for more detail.

Pimlico's source code is available on [Github](#).

## Manual setup

If for some reason you don't want to use the `newproject.py` script, you can set up a project yourself. Download Pimlico [from Github](#).

Simply download the whole source code as a `.zip` or `.tar.gz` file and uncompress it. This will produce a directory called `pimlico`, followed by a long incomprehensible string, which you can rename simply `pimlico`.

Pimlico has a few basic dependencies, but these will be automatically downloaded the first time you load it.

### 1.2.2 Pipeline config

A Pimlico pipeline, as read from a config file (`pimlico.core.config.PipelineConfig`) contains all the information about the pipeline being processed and provides access to specific modules in it. A config file looks much like a standard `.ini` file, with sections headed by `[section_name]` headings, containing key-value parameters of the form `key=value`.

Each section, except for `vars` and `pipeline`, defines a module instance in the pipeline. Some of these can be executed, others act as filters on the outputs of other modules, or input readers.

Each section that defines a module has a `type` parameter. Usually, this is a fully-qualified Python package name that leads to the module type's Python code (that package containing the `info` Python module). A special type is `alias`. This simply defines a module alias – an alternative name for an already defined module. It should have exactly one other parameter, `input`, specifying the name of the module we're aliasing.

#### Special sections

- **vars:** May contain any variable definitions, to be used later on in the pipeline. Further down, expressions like `%(varname)s` will be expanded into the value assigned to `varname` in the `vars` section.
- **pipeline:** Main pipeline-wide configuration. The following options are required for every pipeline:
  - `name`: a single-word name for the pipeline, used to determine where files are stored
  - `release`: the release of Pimlico for which the config file was written. It is considered compatible with later minor versions of the same major release, but not with later major releases. Typically, a user receiving the pipeline config will get hold of an appropriate version of the Pimlico codebase to run it with.

Other optional settings:

- `python_path`: a path or paths, relative to the directory containing the config file, in which Python modules/packages used by the pipeline can be found. Typically, a config file is distributed with a directory of Python code providing extra modules, datatypes, etc. Multiple paths are separated by colons (`:`).

#### Special variable substitutions

Certain variable substitutions are always available, in addition to those defined in `vars` sections. Use them anywhere in your config file with an expression like `%(varname)s` (note the `s` at the end).

- **pimlico\_root:** Root directory of Pimlico, usually the directory `pimlico/` within the project directory.
- **project\_root:** Root directory of the whole project. Current assumed to always be the parent directory of `pimlico_root`.
- **output\_dir:** Path to output dir (usually `output` in Pimlico root).
- **long\_term\_store:** Long-term store base directory being used under the current config. Can be used to link to data from other pipelines run on the same system. This is the value specified in the *local config file*.
- **short\_term\_store:** Short-term store base directory being used under the current config. Can be used to link to data from other pipelines run on the same system. This is the value specified in the *local config file*.
- **home:** Running user's home directory (on Unix and Windows, see Python's `os.path.expanduser()`).
- **test\_data\_dir:** Directory in Pimlico distribution where test data is stored (`test/data` in Pimlico root). Used in test pipelines, which take all their input data from this directory.

For example, to point a parameter to a file located within the project root:

```
param=%(project_root)s/data/myfile.txt
```

## Directives

Certain special directives are processed when reading config files. They are lines that begin with `%%`, followed by the directive name and any arguments.

- **variant:**

Allows a line to be included only when loading a particular variant of a pipeline. For more detail on pipeline variants, see [Pipeline variants](#).

The variant name is specified as part of the directive in the form: `variant:variant_name`. You may include the line in more than one variant by specifying multiple names, separated by commas (and no spaces). You can use the default variant “main”, so that the line will be left out of other variants. The rest of the line, after the directive and variant name(s) is the content that will be included in those variants.

```
[my_module]
type=path.to.module
%%variant:main size=52
%%variant:smaller size=7
```

An alternative notation for the variant directive is provided to make config files more readable. Instead of `variant:variant_name`, you can write `(variant_name)`. So the above example becomes:

```
[my_module]
type=path.to.module
%%(main) size=52
%%(smaller) size=7
```

- **novariant:** A line to be included only when not loading a variant of the pipeline. Equivalent to `variant:main`.

```
[my_module]
type=path.to.module
%%novariant size=52
%%variant:smaller size=7
```

- **include:** Include the entire contents of another file. The filename, specified relative to the config file in which the directive is found, is given after a space.
- **abstract:** Marks a config file as being abstract. This means that Pimlico will not allow it to be loaded as a top-level config file, but only allow it to be included in another config file.
- **copy:** Copies all config settings from another module, whose name is given as the sole argument. May be used multiple times in the same module and later copies will override earlier. Settings given explicitly in the module’s config override any copied settings.

All parameters are copied, including things like `type`. Any parameter can be overridden in the copying module instance.

The directive even allows you to copy parameters from multiple modules by using the directive multiple times, though this is not very often useful. In this case, the values are copied (and overridden) in the order of the directives.

For example, to reuse all the parameters from `module1` in `module2`, only specifying them once:

```
[module1]
type=some.module.type
input=moduleA
param1=56
param2=never
param3=0.75

[module2]
# Copy all params from module1
%%copy module1
# Override the input module
input=moduleB
```

## Multiple parameter values

Sometimes you want to write a whole load of modules that are almost identical, varying in just one or two parameters. You can give a parameter multiple values by writing them separated by vertical bars (`|`). The module definition will be expanded to produce a separate module for each value, with all the other parameters being identical.

For example, this will produce three module instances, all having the same `num_lines` parameter, but each with a different `num_chars`:

```
[my_module]
type=module.type.path
num_lines=10
num_chars=3|10|20
```

You can even do this with multiple parameters of the same module and the expanded modules will cover all combinations of the parameter assignments.

For example:

```
[my_module]
type=module.type.path
num_lines=10|50|100
num_chars=3|10|20
```

## Tying alternatives

You can change the behaviour of alternative values using the `tie_alts` option. `tie_alts=T` will cause parameters within the same module that have multiple alternatives to be expanded in parallel, rather than taking the product of the alternative sets. So, if `option_a` has 5 values and `option_b` has 5 values, instead of producing 25 pipeline modules, we'll only produce 5, matching up each pair of values in their alternatives.

```
[my_module]
type=module.type.path
tie_alts=T
option_a=1|2|3|4|5
option_b=one|two|three|four|five
```

If you want to tie together the alternative values on some parameters, but not others, you can specify groups of parameter names to tie using the `tie_alts` option. Each group is separated by spaces and the names of parameters to tie within a group are separated by `|` s. Any parameters that have alternative values but are not specified in one of the groups are not tied to anything else.

For example, the following module config will tie together `option_a`'s alternatives with `option_b`'s, but produce all combinations of them with `option_c`'s alternatives, resulting in  $3 \times 2 = 6$  versions of the module (`my_module[option_a=1~option_b=one~option_c=x]`, `my_module[option_a=1~option_b=one~option_c=y]`, `my_module[option_a=2~option_b=two~option_c=x]` etc).

```
[my_module]
type=module.type.path
tie_alts=option_a|option_b
option_a=1|2|3
option_b=one|two|three
option_c=x|y
```

Using this method, you must give the parameter names in `tie_alts` exactly as you specify them in the config. For example, although for a particular module you might be able to specify a certain input (the default) using the name `input` or a specific name like `input_data`, these will not be recognised as being the same parameter in the process of expanding out the combinations of alternatives.

## Naming alternatives

Each module will be given a distinct name, based on the varied parameters. If just one is varied, the names will be of the form `module_name[param_value]`. If multiple parameters are varied at once, the names will be `module_name[param_name0=param_value0~param_name1=param_value1~...]`. So, the first example above will produce: `my_module[3]`, `my_module[10]` and `my_module[20]`. And the second will produce: `my_module[num_lines=10~num_chars=3]`, `my_module[num_lines=10~num_chars=10]`, etc.

You can also specify your own identifier for the alternative parameter values, instead of using the values themselves (say, for example, if it's a long file path). Specify it surrounded by curly braces at the start of the value in the alternatives list. For example:

```
[my_module]
type=module.type.path
file_path={small}/home/me/data/corpus/small_version|{big}/home/me/data/corpus/big_
↪version
```

This will result in the modules `my_module[small]` and `my_module[big]`, instead of using the whole file path to distinguish them.

An alternative approach to naming the expanded alternatives can be selected using the `alt_naming` parameter. The default behaviour described above corresponds to `alt_naming=full`. If you choose `alt_naming=pos`, the alternative parameter settings (using names where available, as above) will be distinguished like positional arguments, without making explicit what parameter each value corresponds to. This can make for nice concise names in cases where it's clear what parameters the values refer to.

If you specify `alt_naming=full` explicitly, you can also give a further option `alt_naming=full(inputnames)`. This has the effect of removing the `input_` from the start of named inputs. This often makes for intuitive module names, but is not the default behaviour, since there's no guarantee that the input name (without the initial `input_`) does not clash with an option name.

Another possibility, which is occasionally appropriate, is `alt_naming=option(<name>)`, where `<name>` is the name of an option that has alternatives. In this case, the names of the alternatives for the whole module will be taken directly from the alternative names on that option only. (E.g. specified by `{name}` or inherited from a previous module, see below). You may specify multiple option names, separated by commas, and the corresponding alt names will be separated by `~`. If there's only one option with alternatives, this is equivalent to `alt_naming=pos`. If there are multiple, it might often lead to name clashes. The circumstance in which this is most commonly appropriate is

where you use `tie_alts=T`, so it's sufficient to distinguish the alternatives by the name associated with just one option.

## Expanding alternatives down the pipeline

If a module takes input from a module that has been expanded into multiple versions for alternative parameter values, it too will automatically get expanded, as if all the multiple versions of the previous module had been given as alternative values for the input parameter. For example, the following will result in 3 versions of `my_module` (`my_module[1]`, etc) and 3 corresponding versions of `my_next_module` (`my_next_module[1]`, etc):

```
[my_module]
type=module.type.path
option_a=1|2|3

[my_next_module]
type=another.module.type.path
input=my_module
```

Where possible, names given to the alternative parameter values in the first module will be carried through to the next.

## Module variables: passing information through the pipeline

When a pipeline is read in, each module instance has a set of *module variables* associated with it. In your config file, you may specify assignments to the variables for a particular module. Each module inherits all of the variable assignments from modules that it receives its inputs from.

The main reason for having module variables is to be able to do things in later modules that depend on what path through the pipeline an input came from. Once you have defined the sequence of processing steps that pass module variables through the pipeline, apply mappings to them, etc, you can use them in the parameters passed into modules.

### Basic assignment

Module variables are set by including parameters in a module's config of the form `modvar_<name> = <value>`. This will assign `value` to the variable `name` for this module. The simplest form of assignment is just a string literal, enclosed in double quotes:


```
[my_module]
type=module.type.path
modvar_myvar = "Value of my variable"
```

## Names of alternatives

Say we have a simple pipeline that has a single source of data, with different versions of the dataset for different languages (English and Swedish). A series of modules process each language in an identical way and, at the end, outputs from all languages are collected by a single `summary` module. This final module may need to know what language each of its incoming datasets represents, so that it can output something that we can understand.

The two languages are given as alternative values for a parameter `path`, and the whole pipeline gets automatically expanded into two paths for the two alternatives:





The `summary` module gets its two inputs for the two different languages as a multiple-input: this means we could expand this pipeline to as many languages as we want, just by adding to the `input_src` module's `path` parameter.

However, as far as `summary` is concerned, this is just a list of datasets – it doesn't know that one of them is English and one is Swedish. But let's say we want it to output a table of results. We're going to need some labels to identify the languages.

The solution is to add a module variable to the first module that takes different values when it gets expanded into two modules. For this, we can use the `altname` function in a `modvar` assignment: this assigns the name of the expanded module's alternative for a given parameter that has alternatives in the config.

```
[input_src]
path={en}/to/english | {sv}/to/swedish
modvar_lang=altname(path)
```

Now the expanded module `input_src[en]` will have the module variable `lang="en"` and the Swedish version `lang="sv"`. This value gets passed from module to module down the two paths in the pipeline.

### Other assignment syntax

A further function `map` allows you to apply a mapping to a value, rather like a Python dictionary lookup. Its first argument is the value to be mapped (or anything that expands to a value, using `modvar` assignment syntax). The second is the mapping. This is simply a space-separated list of source-target mappings of the form `source -> target`. Typically both the sources and targets will be string literals.

Now we can give our languages legible names. (Here we're splitting the definition over multiple indented lines, as permitted by config file syntax, which makes the mapping easier to read.)

```
[input_src]
path={en}/to/english | {sv}/to/swedish
modvar_lang=map(
    altname(path),
    "en" -> "English"
    "sv" -> "Svenska")
```

The assignments may also reference variable names, including those previously assigned to in the same module and those received from the input modules.

```
[input_src]
path={en}/to/english | {sv}/to/swedish
modvar_lang=altname(path)
modvar_lang_name=map(
    lang,
    "en" -> "English"
    "sv" -> "Svenska")
```

If a module gets two values for the same variable from multiple inputs, the first value will simply be overridden by the second. Sometimes it's useful to map module variables from specific inputs to different modvar names. For example, if we're combining two different languages, we might need to keep track of what the two languages we combined were. We can do this using the notation `input_name.var_name`, which refers to the value of module variable `var_name` that was received from input `input_name`.

```
[input_src]
path={en}/to/english | {sv}/to/swedish
modvar_lang=altname(path)

[combiner]
type=my.language.combiner
input_lang_a=lang_data
input_lang_b=lang_data
modvar_first_lang=lang_a.lang
modvar_second_lang=lang_b.lang
```

## Use in module parameters

To make something in a module's execution dependent on its module variables, you can insert them into module parameters.

For example, say we want one of the module's parameters to make use of the `lang` variable we defined above:

```
[input_src]
path={en}/to/english | {sv}/to/swedish
modvar_lang=altname(path)
some_param=$(lang)
```

Note the difference to other variable substitutions, which use the `%(varname)s` notation. For modvars, we use the notation `$(varname)`.

We can also put the value in the middle of other text:

```
[input_src]
path={en}/to/english | {sv}/to/swedish
modvar_lang=altname(path)
some_param=myval-$(lang)-continues
```

The modvar processing to compute a particular module's set of variable assignments is performed before the substitution. This means that you can do any modvar processing specific to the module instance, in the various ways defined above, and use the resulting value in other parameters. For example:

```
[input_src]
path={en}/to/english | {sv}/to/swedish
modvar_lang=altname(path)
```

```
modvar_mapped_lang=map(lang,
    "en" -> "eng"
    "sv" -> "swe"
)
some_param=$(mapped_lang)
```

You can also place in the `$(...)` construct any of the variable processing operations shown above for assignments to module variables. This is a little more concise than first assigning values to modvars, if you don't need to use the variables again anywhere else. For example:

```
[input_src]
path={en}/to/english | {sv}/to/swedish
some_param=$(map(altname(path)),
    "en" -> "eng"
    "sv" -> "swe"
))
```

## Usage in module code

A module's executor can also retrieve the values assigned to module variables from the `module_variables` attribute of the module-info associated with the input dataset. Sometimes this can be useful when you are writing your own module code, though the above usage to pass values from (or dependent on) module variables into module parameters is more flexible, so should generally be preferred.

```
# Code in executor
# This is a MultipleInput-type input, so we get a list of datasets
datasets = self.info.get_input()
for d in datasets:
    language = d.module.module_variables["lang"]
```

### 1.2.3 Pipeline variants

---

#### Todo

Document variants

---

### 1.2.4 Pimlico module structure

This document describes the code structure for Pimlico module types in full.

For a basic guide to writing your own modules, see [Writing Pimlico modules](#).

---

#### Todo

Write documentation for this

---

## 1.2.5 Module dependencies

In a Pimlico pipeline, you typically use lots of different external software packages. Some are Python packages, others system tools, Java libraries, whatever. Even the core modules that core with Pimlico between them depend on a huge amount of software.

Naturally, we don't want to have to install *all* of this software before you can run even a simple Pimlico pipeline that doesn't use all (or any) of it. So, we keep the core dependencies of Pimlico to an absolute minimum, and then check whether the necessary software dependencies are installed each time a pipeline module is going to be run.

### Core dependencies

Certain dependencies are required for Pimlico to run at all, or needed so often that you wouldn't get far without installing them. These are defined in `pimlico.core.dependencies.core`, and when you run the Pimlico command-line interface, it checks they're available and tries to install them if they're not.

### Module dependencies

Each module type defines its own set of software dependencies, if it has any. When you try to run the module, Pimlico runs some checks to try to make sure that all of these are available.

If some of them are not, it may be possible to install them automatically, straight from Pimlico. In particular, many Python packages can be very easily installed using `Pip`. If this is the case for one of the missing dependencies, Pimlico will tell you in the error output, and you can install them using the `install` command (with the module name/number as an argument).

### Virtualenv

In order to simplify automatic installation, Pimlico is always run within a virtual environment, using `Virtualenv`. This means that any Python packages installed by `Pip` will live in a local directory within the Pimlico codebase that you're running and won't interfere with anything else on your system.

When you run Pimlico for the first time, it will create a new virtualenv for this purpose. Every time you run it after that, it will use this same environment, so anything you install will continue to be available.

### Custom virtualenv

Most of the time, you don't even need to be aware of the virtualenv that Python's running in<sup>1</sup>. Under certain circumstances, you might need to use a custom virtualenv.

For example, say you're running your pipeline over different servers, but have the pipeline and Pimlico codebase on a shared network drive. Then you can find that the software installed in the virtualenv on one machine is incompatible with the system-wide software on the other.

You can specify a name for a custom virtualenv using the environment variable `PIMENV`. The first time you run Pimlico with this set, it will automatically create the new virtualenv.

```
$ PIMENV=myenv ./pimlico.sh mypipeline.conf status
```

Replace `myenv` with a name that better reflects its use (e.g. name of the server).

Every time you run Pimlico on that server, set the `PIMENV` environment variable in the same way.

---

<sup>1</sup> If you're interested, it lives in `pimlico/lib/virtualenv/default`

In case you want to get to the virtualenv itself, you can find it in `pimlico/lib/virtualenv/myenv`.

---

**Note:** Pimlico previously used another environment variable `VIRTUALENV`, which gave a path to the virtualenv. You can still use this, but, unless you have a good reason to, it's easier to use `PIMENV`.

---

## Defining module dependencies

---

### Todo

Describe how module dependencies are defined for different types of deps

---

### Some examples

---

### Todo

Include some examples from the core modules of how deps are defined and some special cases of software fetching

---

## 1.2.6 Local configuration

As well as knowing about the pipeline you're running, Pimlico also needs to know some things about the setup of the system on which you're running it. This is completely independent of the pipeline config: the same pipeline can be run on different systems with different local setups.

A couple of settings must always be provided for Pimlico: the **long-term** and **short-term stores** (see [Long-term and short-term stores](#) below). Other system settings may be specified as necessary. (At the time of writing, there aren't any, but they will be documented here as they arise.) See [Other Pimlico settings](#) below.

Specific modules may also have system-level settings. For example, a module that calls an external tool may need to know the location of that tool, or how much memory it can use on this system. Any that apply to the built-in Pimlico modules are listed below in [Settings for built-in modules](#).

### Local config file location

Pimlico looks in various places to find the local config settings. Settings are loaded in a particular order, overriding earlier versions of the same setting as we go (see `pimlico.core.config.PipelineConfig.load_local_config()`).

Settings are specified with the following order of precedence (those later override the earlier):

```
local config file < host-specific config file < cmd-line overrides
```

Most often, you'll just specify all settings in the main local config file. This is a file in your home directory named `.pimlico`. This must exist for Pimlico to be able to run at all.

## Host-specific config

If you share your home directory between different computers (e.g. a networked filesystem), the above setup could cause a problem, as you may need a different local config on the different computers. Pimlico allows you to have special config files that only get read on machines with a particular hostname.

For example, say I have two computers, `localbox` and `remotebox`, which share a home directory. I've created my `.pimlico` local config file on `localbox`, but need to specify a different storage location on `remotebox`. I simply create another config file called `.pimlico_remotebox`[#hostname]_`. Pimlico will load first the basic local config in ``.pimlico` and then override those settings with what it reads from the host-specific config file.

You can also specify a hostname prefix to match. Say I've got a whole load of computers I want to be able to run on, with hostnames `remotebox1`, `remotebox2`, etc. If I create a config file called `.pimlico_remotebox-`, it will be used on all of these hosts.

## Command-line overrides

Occasionally, you might want to specify a local config setting just for one run of Pimlico. Use the `--override-local-config` (or `-l`) to specify a value for an individual setting in the form `setting=value`. For example:

```
./pimlico.sh mypipeline.conf -l somesetting=5 run mymodule
```

If you want to override multiple settings, simply use the option multiple times.

## Custom location

If the above solutions don't work for you, you can also explicitly specify on the command line an alternative location from which to load the local config file that Pimlico typically expects to find in `~/pimlico`.

Use the `--local-config` parameter to give a filename to use instead of the `~/pimlico`.

For example, if your home directory is shared across servers and the above hostname-specific config solution doesn't work in your case, you can fall back to pointing Pimlico at your own host-specific config file.

## Long-term and short-term stores

Pimlico needs to know where to put and find output files as it executes. Settings are given in the local config, since they apply to all Pimlico pipelines you run and may vary from system to system. Note that Pimlico will make sure that different pipelines don't interfere with each other's output (provided you give them different names): all pipelines store their output and look for their input within these same base locations.

The **short-term store** should be on a disk that's as fast as possible to write to. For example, avoid using an NFS disk. It needs to be large enough to store output between pipeline stages, though you can easily move output from earlier stages into the long-term store as you go along.

The **long-term store** is where things are typically put at the end of a pipeline. It therefore doesn't need to be super-fast to access, but you may want it to be in a location that gets backed up, so you don't lose your valuable output.

For a simple setup, these could be just two subdirectories of the same directory, or actually the same directory. However, it can be useful to distinguish them.

Specify the locations in the local config like this:

```
long_term_store=/path/to/long-term/store
short_term_store=/path/to/short-term/store
```

Remember, these paths are not specific to a pipeline: all pipelines will use different subdirectories of these ones.

## Other Pimlico settings

In future, there will no doubt be more settings that you can specify at the system level for Pimlico. These will be documented here as they arise.

## Settings for built-in modules

Specific modules may consult the local config to allow you to specify settings for them. We cannot document them here for all modules, as we don't know what modules are being developed outside the core codebase. However, we can provide a list here of the settings consulted by built-in Pimlico modules.

There aren't any yet, but they will be listed here as they arise.

## Footnotes:

### 1.2.7 Python scripts

All the heavy work of your data-processing is implemented in Pimlico modules, either by loading core Pimlico modules from your pipeline config file or by writing your own modules. Sometimes, however, it can be handy to write a quick Python script to get hold of the output of one of your pipeline's modules and inspect it or do something with it.

This can be easily done writing a Python script and using the *python* shell command to run it. This command loads your pipeline config (just like all others) and then either runs a script you've specified on the command line, or enters an interactive Python shell. The advantages of this over just running the normal *python* command on the command line are that the script is run in the same execution context used for your pipeline (e.g. using the Pimlico instance's virtualenv) and that the loaded pipeline is available to you, so you can easily can hold of its data locations, datatypes, etc.

## Accessing the pipeline

At the top of your Python script, you can get hold of the loaded pipeline config instance like this:

```
from pimlico.cli.pyshell import get_pipeline

pipeline = get_pipeline()
```

Now you can use this to get to, among other things, the pipeline's modules and their input and output datasets. A module called `module1` can be accessed by treating the pipeline like a dict:

```
module = pipeline["module1"]
```

This gives you the `ModuleInfo` instance for that module, giving access to its inputs, outputs, options, etc:

```
data = module.get_output("output_name")
```

## Writing and running scripts

All of the above code to access a pipeline can be put in a Python script somewhere in your codebase and run from the command line. Let's say I create a script `src/python/scripts/myscript.py` containing:

```
from pimlico.cli.pyshell import get_pipeline

pipeline = get_pipeline()
module = pipeline["module1"]
data = module.get_output("output_name")
# Here we can start probing the data using whatever interface the datatype provides
print data
```

Now I can run this from the root directory of my project as follows:

```
./pimlico.sh mypipeline.conf python src/python/scripts/myscript.py
```

## 1.3 Core Pimlico modules

Pimlico comes with a substantial collection of module types that provide wrappers around existing NLP and machine learning tools, as well as a number of general tools for processing datasets that are useful for many applications.

### 1.3.1 CAEVO event extractor

Path	pimlico.modules.caevo
Executable	yes

**CAEVO** is Nate Chambers' CAscading EVent Ordering system, a tool for extracting events of many types from text and ordering them.

**CAEVO** is open source, implemented in Java, so is easily integrated into Pimlico using Py4J.

#### Inputs

Name	Type(s)
documents	TarredCorpus<RawTextDocumentType>

#### Outputs

Name	Type(s)
events	<i>CaevoCorpus</i>

#### Options

Name	Description	Type
sieves	Filename of sieve list file, or path to the file. If just a filename, assumed to be in Caevo model dir (models/caevo). Default: default.sieves (supplied with Caevo)	string



## Submodules

### CAEVO output convertor

Path	pimlico.modules.caevo.output
Executable	yes

Tool to split up the output from Caevo and convert it to other datatypes.

Caevo output includes the output from a load of NLP tools that it runs as prerequisites to event extraction, etc. The individual parts of the output can easily be retrieved from the output corpus via the output datatype. In order to be able to use them as input to other modules, they need to be converted to compatible standard datatypes.

For example, tokenization output is stored in Caevo's XML output using a special format. Instead of writing other modules in such a way as to be able to pull this information out of the :class:~pimlico.datatypes.CaevoCorpus, you can filter the output using this module to provide a :class:~pimlico.datatypes.TokenizedCorpus, which is a standard format for input to other module types.

As with other document map modules, you can use this as a filter (*filter=T*), so you can actually need to commit the converted data to disk.

---

### Todo

Add more output convertors: currently only provides tokenization

---

## Inputs

Name	Type(s)
documents	<i>CaevoCorpus</i>

## Outputs

No non-optional outputs

## Optional

Name	Type(s)
tokenized	<i>TokenizedCorpus</i>
parse	<i>ConstituencyParseTreeCorpus</i>
pos	WordAnnotationCorpusWithWordAndPos

## Options

Name	Description	Type
gzip	If True, each output, except annotations, for each document is gzipped. This can help reduce the storage occupied by e.g. parser or coref output. Default: False	bool

### 1.3.2 C&C parser

Path	pimlico.modules.candc
Executable	yes

Wrapper around the original [C&C parser](#).

Takes tokenized input and parses it with C&C. The output is written exactly as it comes out from C&C. It contains both GRs and supertags, plus POS-tags, etc.

The wrapper uses C&C's SOAP server. It sets the SOAP server running in the background and then calls C&C's SOAP client for each document. If parallelizing, multiple SOAP servers are set going and each one is kept constantly fed with documents.

#### Inputs

Name	Type(s)
documents	TarredCorpus<TokenizedDocumentType>

#### Outputs

Name	Type(s)
parsed	<i>CandcOutputCorpus</i>

#### Options

Name	Description	Type
model	Absolute path to models directory or name of model set. If not an absolute path, assumed to be a subdirectory of the candc models dir (see instructions in models/candc/README on how to fetch pre-trained models)	string

### 1.3.3 Stanford CoreNLP

Path	pimlico.modules.corenlp
Executable	yes

Process documents one at a time with the [Stanford CoreNLP toolkit](#). CoreNLP provides a large number of NLP tools, including a POS-tagger, various parsers, named-entity recognition and coreference resolution. Most of these tools can be run using this module.

The module uses the CoreNLP server to accept many inputs without the overhead of loading models. If parallelizing, only a single CoreNLP server is run, since this is designed to set multiple Java threads running if it receives multiple queries at the same time. Multiple Python processes send queries to the server and process the output.

The module has no non-optional outputs, since what sort of output is available depends on the options you pass in: that is, on which tools are run. Use the annotations option to choose which word annotations are added. Otherwise, simply select the outputs that you want and the necessary tools will be run in the CoreNLP pipeline to produce those outputs.

Currently, the module only accepts tokenized input. If pre-POS-tagged input is given, for example, the POS tags won't be handed into CoreNLP. In the future, this will be implemented.

We also don't currently provide a way of choosing models other than the standard, pre-trained English models. This is a small addition that will be implemented in the future.

## Inputs

Name	Type(s)
documents	TarredCorpus<WordAnnotationsDocumentType/TokenizedDocumentType/RawTextDocumentType>

## Outputs

No non-optional outputs

## Optional

Name	Type(s)
annotations	<i>AnnotationFieldsFromOptions</i>
tokenized	<i>TokenizedCorpus</i>
parse	<i>ConstituencyParseTreeCorpus</i>
parse-deps	<i>StanfordDependencyParseCorpus</i>
dep-parse	<i>StanfordDependencyParseCorpus</i>
raw	<i>JsonDocumentCorpus</i>
coref	<i>CorefCorpus</i>

## Options

Name	Description	Type
gzip	If True, each output, except annotations, for each document is gzipped. This can help reduce the storage occupied by e.g. parser or coref output. Default: False	bool
time-out	Timeout for the CoreNLP server, which is applied to every job (document). Number of seconds. By default, we use the server's default timeout (15 secs), but you may want to increase this for more intensive tasks, like coref	float
readable	If True, JSON outputs are formatted in a readable fashion, pretty printed. Otherwise, they're as compact as possible. Default: False	bool
an-no-ta-tors	Comma-separated list of word annotations to add, from CoreNLP's annotators. Choose from: word, pos, lemma, ner	string
dep_type	Type of dependency parse to output, when outputting dependency parses, either from a constituency parse or direct dependency parse. Choose from the three types allowed by CoreNLP: 'basic', 'collapsed' or 'collapsed-ccprocessed'	'basic', 'collapsed' or 'collapsed-ccprocessed'

### 1.3.4 Corpus manipulation

Core modules for generic manipulation of mainly iterable corpora.

#### Corpus concatenation

Path	pimlico.modules.corpora.concat
Executable	no

Concatenate two corpora to produce a bigger corpus.

They must have the same data point type, or one must be a subtype of the other.

In theory, we could find the most specific common ancestor and use that as the output type, but this is not currently implemented and may not be worth the trouble. Perhaps we will add this in future.

This is a filter module. It is not executable, so won't appear in a pipeline's list of modules that can be run. It produces its output for the next module on the fly when the next module needs it.

### Inputs

Name	Type(s)
corpora	<i>list</i> of <i>IterableCorpus</i>

### Outputs

Name	Type(s)
corpus	corpus with data-point from input

### Corpus statistics

Path	pimlico.modules.corpora.corpus_stats
Executable	yes

Some basic statistics about tokenized corpora

Counts the number of tokens, sentences and distinct tokens in a corpus.

### Inputs

Name	Type(s)
corpus	TarredCorpus<TokenizedDocumentType>

### Outputs

Name	Type(s)
stats	<i>NamedFile</i> ()

### Human-readable formatting

Path	pimlico.modules.corpora.format
Executable	yes

Corpus formatter

Pimlico provides a data browser to make it easy to view documents in a tarred document corpus. Some datatypes provide a way to format the data for display in the browser, whilst others provide multiple formatters that display the data in different ways.

This module allows you to use this formatting functionality to output the formatted data as a corpus. Since the formatting operations are designed for display, this is generally only useful to output the data for human consumption.

## Inputs

Name	Type(s)
corpus	<i>TarredCorpus</i>

## Outputs

Name	Type(s)
formatted	<i>TarredCorpus</i>

## Options

Name	Description	Type
for-mat-ter	Fully qualified class name of a formatter to use to format the data. If not specified, the default formatter is used, which uses the datatype's <code>browser_display</code> attribute if available, or falls back to just converting documents to unicode	string

## Corpus document list filter

Path	<code>pimlico.modules.corpora.list_filter</code>
Executable	yes

Similar to `:mod:pimlico.modules.corpora.split`, but instead of taking a random split of the dataset, splits it according to a given list of documents, putting those in the list in one set and the rest in another.

## Inputs

Name	Type(s)
corpus	<i>TarredCorpus</i>
list	<i>StringList</i>

## Outputs

Name	Type(s)
set1	same as input corpus
set2	same as input corpus

## Corpus split

Path	<code>pimlico.modules.corpora.split</code>
Executable	yes

Split a tarred corpus into two subsets. Useful for dividing a dataset into training and test subsets. The output datasets have the same type as the input. The documents to put in each set are selected randomly. Running the module multiple times will give different splits.

Note that you can use this multiple times successively to split more than two ways. For example, say you wanted a training set with 80% of your data, a dev set with 10% and a test set with 10%, split it first into training and non-training 80-20, then split the non-training 50-50 into dev and test.

The module also outputs a list of the document names that were included in the first set. Optionally, it outputs the same thing for the second input too. Note that you might prefer to only store this list for the smaller set: e.g. in a training-test split, store only the test document list, as the training list will be much larger. In such a case, just put the smaller set first and don't request the optional output *doc\_list2*.

## Inputs

Name	Type(s)
corpus	<i>TarredCorpus</i>

## Outputs

Name	Type(s)
set1	same as input corpus
set2	same as input corpus
doc_list1	<i>StringList</i>

## Optional

Name	Type(s)
doc_list2	<i>StringList</i>

## Options

Name	Description	Type
set1_size	Proportion of the corpus to put in the first set, float between 0.0 and 1.0. If an integer >1 is given, this is treated as the absolute number of documents to put in the first set, rather than a proportion. Default: 0.2 (20%)	float

## Corpus subset

Path	pimlico.modules.corpora.subset
Executable	no

Simple filter to truncate a dataset after a given number of documents, potentially offsetting by a number of documents. Mainly useful for creating small subsets of a corpus for testing a pipeline before running on the full corpus.

Can be run on an iterable corpus or a tarred corpus. If the input is a tarred corpus, the filter will emulate a tarred corpus with the appropriate datatype, passing through the archive names from the input.

This is a filter module. It is not executable, so won't appear in a pipeline's list of modules that can be run. It produces its output for the next module on the fly when the next module needs it.

## Inputs

Name	Type(s)
documents	<i>IterableCorpus</i>

## Outputs

Name	Type(s)
documents	corpus with data-point from input

## Options

Name	Description	Type
offset	Number of documents to skip at the beginning of the corpus (default: 0, start at beginning)	int
skip_invalid	Skip over any invalid documents so that the output subset contains the chosen number of (valid) documents (or as many as possible) and no invalid ones. By default, invalid documents are passed through and counted towards the subset size	bool
size	(required) Number of documents to include	int

## Tar archive grouper

Path	pimlico.modules.corpora.tar
Executable	yes

Group the files of a multi-file iterable corpus into tar archives. This is a standard thing to do at the start of the pipeline, since it's a handy way to store many (potentially small) files without running into filesystem problems.

The files are simply grouped linearly into a series of tar archives such that each (apart from the last) contains the given number.

After grouping documents in this way, document map modules can be called on the corpus and the grouping will be preserved as the corpus passes through the pipeline.

---

**Note:** There is a fundamental problem with this module. It stores the raw data that it gets as input, and reports the output type as the same as the input type. However, it doesn't correctly write that type. A lot of the time, this isn't a problem, but it means that it doesn't write corpus metadata that may be needed by the datatype to read the documents correctly.

The new datatypes system will provide a solution to this problem, but until then the safest approach is not to use this module, but always use `tar_filter` instead, which doesn't have this problem.

---

## Inputs

Name	Type(s)
documents	<i>IterableCorpus</i>

## Outputs

Name	Type(s)
documents	<i>tarred corpus with input doc type</i>

## Options

Name	Description	Type
archive_size	Number of documents to include in each archive (default: 1k)	string
archive_basename	Base name to use for archive tar files. The archive number is appended to this. (Default: 'archive')	string

## Tar archive grouper (filter)

Path	pimlico.modules.corpora.tar_filter
Executable	no

Like *tar*, but doesn't write the archives to disk. Instead simulates the behaviour of tar but as a filter, grouping files on the fly and passing them through with an archive name

This is a filter module. It is not executable, so won't appear in a pipeline's list of modules that can be run. It produces its output for the next module on the fly when the next module needs it.

## Inputs

Name	Type(s)
documents	<i>IterableCorpus</i>

## Outputs

Name	Type(s)
documents	<i>tarred corpus with input doc type</i>

## Options

Name	Description	Type
archive_size	Number of documents to include in each archive (default: 1k)	string
archive_basename	Base name to use for archive tar files. The archive number is appended to this. (Default: 'archive')	string

## Corpus vocab builder

Path	pimlico.modules.corpora.vocab_builder
Executable	yes

Builds a dictionary (or vocabulary) for a tokenized corpus. This is a data structure that assigns an integer ID to every distinct word seen in the corpus, optionally applying thresholds so that some words are left out.



Similar to `pimlico.modules.features.vocab_builder`, which builds two vocabs, one for terms and one for features.

### Inputs

Name	Type(s)
text	TarredCorpus<TokenizedDocumentType>

### Outputs

Name	Type(s)
vocab	<i>Dictionary</i>

### Options

Name	Description	Type
thresh-old	Minimum number of occurrences required of a term to be included	int
max_pr	Include terms that occur in max this proportion of documents	float
in-clude	Ensure that certain words are always included in the vocabulary, even if they don't make it past the various filters, or are never seen in the corpus. Give as a comma-separated list	comma-separated list of strings
limit	Limit vocab size to this number of most common entries (after other filters)	int

### Tokenized corpus to ID mapper

Path	pimlico.modules.corpora.vocab_mapper
Executable	yes

### Inputs

Name	Type(s)
text	TarredCorpus<TokenizedDocumentType>
vocab	<i>Dictionary</i>

### Outputs

Name	Type(s)
ids	<i>IntegerListsDocumentCorpus</i>

## 1.3.5 Embedding feature extractors and trainers

Modules for extracting features from which to learn word embeddings from corpora, and for training embeddings.

Some of these don't actually learn the embeddings, they just produce features which can then be fed into an embedding learning module, such as a form of matrix factorization. Note that you can train embeddings not only using the

trainers here, but also using generic matrix manipulation techniques, for example the factorization methods provided by sklearn.

### Dependency feature extractor for embeddings

Path	pimlico.modules.embeddings.dependencies
Executable	yes

---

#### Todo

Document this module

---

### Inputs

Name	Type(s)
dependencies	TarredCorpus<CoNLLDependencyParseDocumentType>

### Outputs

Name	Type(s)
term_features	<i>TermFeatureListCorpus</i>

### Options

Name	Description	Type
lemma	Use lemmas as terms instead of the word form. Note that if you didn't run a lemmatizer before dependency parsing the lemmas are probably actually just copies of the word forms	bool
con-dense_prep	Where a word is modified ...TODO	string
term_pos	Only extract features for terms whose POSs are in this comma-separated list. Put a * at the end to denote POS prefixes	comma-separated list of strings
skip_types	Dependency relations to skip, separated by commas	comma-separated list of strings

### Word2vec embedding trainer

Path	pimlico.modules.embeddings.word2vec
Executable	yes

Word2vec embedding learning algorithm, using [Gensim](#)'s implementation.

Find out more about [word2vec](#).

This module is simply a wrapper to call [Gensim](#)'s Python (+C) implementation of word2vec on a Pimlico corpus.

## Inputs

Name	Type(s)
text	TarredCorpus<TokenizedDocumentType>

## Outputs

Name	Type(s)
model	<i>Embeddings</i>

## Options

Name	Description	Type
iters	number of iterations over the data to perform. Default: 5	int
min_count	word2vec's min_count option: prunes the dictionary of words that appear fewer than this number of times in the corpus. Default: 5	int
negative_samples	number of negative samples to include per positive. Default: 5	int
size	number of dimensions in learned vectors. Default: 200	int

### 1.3.6 Feature set processing

Various tools for generic processing of extracted sets of features: building vocabularies, mapping to integer indices, etc.

#### Key-value to term-feature converter

Path	pimlico.modules.features.term_feature_compiler
Executable	yes

---

#### Todo

Document this module

---

## Inputs

Name	Type(s)
key_values	TarredCorpus<KeyValueListDocumentType>

## Outputs

Name	Type(s)
term_features	<i>TermFeatureListCorpus</i>

## Options

Name	Description	Type
term_keys	Name of keys (feature names in the input) which denote terms. The first one found in the keys of a particular data point will be used as the term for that data point. Any other matches will be removed before using the remaining keys as the data point's features. Default: just 'term'	comma-separated list of strings
include_feature_keys	If True, include the key together with the value from the input key-value pairs as feature names in the output. Otherwise, just use the value. E.g. for input [prop=wordy, poss=my], if True we get features [prop_wordy, poss_my] (both with count 1); if False we get just [wordy, my]. Default: False	bool

## Term-feature matrix builder

Path	pimlico.modules.features.term_feature_matrix_builder
Executable	yes

### Todo

Document this module

---

## Inputs

Name	Type(s)
data	<i>IndexedTermFeatureListCorpus</i>

## Outputs

Name	Type(s)
matrix	<i>ScipySparseMatrix</i>

## Term-feature corpus vocab builder

Path	pimlico.modules.features.vocab_builder
Executable	yes

### Todo

Document this module

---

## Inputs

Name	Type(s)
term_features	TarredCorpus<TermFeatureListDocumentType>

## Outputs

Name	Type(s)
term_vocab	<i>Dictionary</i>
feature_vocab	<i>Dictionary</i>

## Options

Name	Description	Type
feature_limit	Limit vocab size to this number of most common entries (after other filters)	int
feature_max_prop	Include features that occur in max this proportion of documents	float
term_max_prop	Include terms that occur in max this proportion of documents	float
term_threshold	Minimum number of occurrences required of a term to be included	int
feature_threshold	Minimum number of occurrences required of a feature to be included	int
term_limit	Limit vocab size to this number of most common entries (after other filters)	int

## Term-feature corpus vocab mapper

Path	pimlico.modules.features.vocab_mapper
Executable	yes

## Todo

Document this module

## Inputs

Name	Type(s)
data	<code>TarredCorpus&lt;TermFeatureListDocumentType&gt;</code>
term_vocab	<i>Dictionary</i>
feature_vocab	<i>Dictionary</i>

## Outputs

Name	Type(s)
data	<i>IndexedTermFeatureListCorpus</i>

### 1.3.7 Input readers

Various input readers for various datatypes. These are used to read in data from some external source, such as a corpus in its distributed format (e.g. XML files or a collection of text files), and present it to the Pimlico pipeline as a Pimlico dataset, which can be used as input to other modules.

They do not typically store the data as a Pimlico dataset, but produce it on the fly, although sometimes it could be appropriate to do otherwise.

Note that there can be multiple input readers for a single datatype. For example, there are many ways to read in a corpus of raw text documents, depending on the format they're stored in. They might be in one big XML file, text

files collected into compressed archives, a big text file with document separators, etc. These all require their own input reader and all of them produce the same output corpus type.

---

**Note:** These input readers are ultimately intended to replace reading input data using a datatype's `input_module_options`. That functionality will be removed altogether as part of the development of the new datatype system, so it should be phased out now and replaced by input reader modules for each datatype.

---

### Package `pimlico.modules.input.embeddings`

#### FastText embedding reader

Path	<code>pimlico.modules.input.embeddings.fasttext</code>
Executable	yes

Reads in embeddings from the [FastText](#) format, storing them in the format used internally in Pimlico for embeddings.

Can be used, for example, to read the [pre-trained embeddings](#) offered by Facebook AI.

Currently only reads the text format (`.vec`), not the binary format (`.bin`).

#### See also:

[`pimlico.modules.input.embeddings.fasttext\_gensim`](#): An alternative reader that uses Gensim's FastText format reading code and permits reading from the binary format, which contains more information.

#### Inputs

No inputs

#### Outputs

Name	Type(s)
embeddings	<a href="#">Embeddings</a>

#### Options

Name	Description	Type
path	(required) Path to the FastText embedding file	string
limit	Limit to the first N words. Since the files are typically ordered from most to least frequent, this limits to the N most common words	int

#### FastText embedding reader using Gensim

Path	<code>pimlico.modules.input.embeddings.fasttext_gensim</code>
Executable	yes

Reads in embeddings from the [FastText](#) format, storing them in the format used internally in Pimlico for embeddings. This version uses Gensim's implementation of the format reader, so depends on Gensim.

Can be used, for example, to read the [pre-trained embeddings](#) offered by Facebook AI.

Reads only the binary format (`.bin`), not the text format (`.vec`).

**See also:**

**`pimlico.modules.input.embeddings.fasttext`:** An alternative reader that does not use Gensim. It permits (only) reading the text format.

## Inputs

No inputs

## Outputs

Name	Type(s)
embeddings	<i>Embeddings</i>

## Options

Name	Description	Type
path	(required) Path to the FastText embedding file ( <code>.bin</code> )	string

## Text corpora

### Raw text files

Path	<code>pimlico.modules.input.text.raw_text_files</code>
Executable	no

Input reader for raw text file collections. Reads in files from arbitrary locations specified by a list of globs.

The input paths must be absolute paths (or globs), but remember that you can make use of various *special substitutions in the config file* to give paths relative to your project root, or other locations.

The file paths may use *globs* to match multiple files. By default, it is assumed that every filename should exist and every glob should match at least one file. If this does not hold, the dataset is assumed to be not ready. You can override this by placing a `?` at the start of a filename/glob, indicating that it will be included if it exists, but is not depended on for considering the data ready to use.

**See also:**

**Datatype `pimlico.datatypes.files.UnnamedFileCollection`** The datatype previously used for reading in file collections, now being phased out to be replaced by this input reader.

This is an input module. It takes no pipeline inputs and is used to read in data

## Inputs

No inputs

## Outputs

Name	Type(s)
corpus	OutputType

## Options

Name	Description	Type
files	(required) Comma-separated list of absolute paths to files to include in the collection. Paths may include globs. Place a '?' at the start of a filename to indicate that it's optional. You can specify a line range for the file by adding ':X-Y' to the end of the path, where X is the first line and Y the last to be included. Either X or Y may be left empty. (Line numbers are 1-indexed.)	comma-separated list of (line range-limited) file paths
ex-clude	A list of files to exclude. Specified in the same way as <i>files</i> (except without line ranges). This allows you to specify a glob in <i>files</i> and then exclude individual files from it (you can use globs here too)	comma-separated list of strings
encod-ing_errors	What to do in the case of invalid characters in the input while decoding (e.g. illegal utf-8 chars). Select 'strict' (default), 'ignore', 'replace'. See Python's <code>str.decode()</code> for details	string
encod-ing	Encoding to assume for input files. Default: utf8	string

### 1.3.8 Malt dependency parser

Wrapper around the [Malt dependency parser](#) and data format converters to support connections to other modules.

#### Annotated text to CoNLL dep parse input converter

Path	pimlico.modules.malt.conll_parser_input
Executable	yes

Converts word-annotations to CoNLL format, ready for input into the Malt parser. Annotations must contain words and POS tags. If they contain lemmas, all the better; otherwise the word will be repeated as the lemma.

## Inputs

Name	Type(s)
annotations	<i>WordAnnotationCorpus</i> with 'word' and 'pos' fields

## Outputs

Name	Type(s)
conll_data	<i>CoNLLEntityParseInputCorpus</i>

#### Malt dependency parser

Path	pimlico.modules.malt.parse
Executable	yes



---

**Todo**Document this module

---

---

**Todo**Replace `check_runtime_dependencies()` with `get_software_dependencies()`

---

**Inputs**

Name	Type(s)
documents	TarredCorpus<CoNLLDependencyParseDocumentType>

**Outputs**

Name	Type(s)
parsed	<i>CoNLLDependencyParseCorpus</i>

**Options**

Name	Description	Type
model	Filename of parsing model, or path to the file. If just a filename, assumed to be Malt models dir (models/malt). Default: engmalt.linear-1.7.mco, which can be acquired by ‘make malt’ in the models dir	string
no_gzip	By default, we gzip each document in the output data. If you don’t do this, the output can get very large, since it’s quite a verbose output format	bool

### 1.3.9 OpenNLP modules

A collection of module types to wrap individual OpenNLP tools.

**OpenNLP coreference resolution**

Path	pimlico.modules.opennlp.coreference
Executable	yes

---

**Todo**Document this module

---

---

**Todo**Replace `check_runtime_dependencies()` with `get_software_dependencies()`

---

Use local config setting `opennlp_memory` to set the limit on Java heap memory for the OpenNLP processes. If parallelizing, this limit is shared between the processes. That is, each OpenNLP worker will have a memory limit of `opennlp_memory / processes`. That setting can use *g*, *G*, *m*, *M*, *k* and *K*, as in the Java setting.

## Inputs

Name	Type(s)
parses	TarredCorpus<TreeStringsDocumentType>

## Outputs

Name	Type(s)
coref	<i>CorefCorpus</i>

## Options

Name	Description	Type
gzip	If True, each output, except annotations, for each document is gzipped. This can help reduce the storage occupied by e.g. parser or coref output. Default: False	bool
model	Coreference resolution model, full path or directory name. If a filename is given, it is expected to be in the OpenNLP model directory (models/opennlp/). Default: '' (standard English opennlp model in models/opennlp/)	string
readable	If True, pretty-print the JSON output, so it's human-readable. Default: False	bool
time-out	Timeout in seconds for each individual coref resolution task. If this is exceeded, an InvalidDocument is returned for that document	int

## OpenNLP coreference resolution

Path	pimlico.modules.opennlp.coreference_pipeline
Executable	yes

Runs the full coreference resolution pipeline using OpenNLP. This includes sentence splitting, tokenization, pos tagging, parsing and coreference resolution. The results of all the stages are available in the output.

Use local config setting `opennlp_memory` to set the limit on Java heap memory for the OpenNLP processes. If parallelizing, this limit is shared between the processes. That is, each OpenNLP worker will have a memory limit of `opennlp_memory / processes`. That setting can use *g*, *G*, *m*, *M*, *k* and *K*, as in the Java setting.

## Inputs

Name	Type(s)
text	TarredCorpus<RawTextDocumentType>

## Outputs

Name	Type(s)
coref	<i>CorefCorpus</i>

## Optional

Name	Type(s)
tokenized	<i>TokenizedCorpus</i>
pos	<i>WordAnnotationCorpusWithPos</i>
parse	<i>ConstituencyParseTreeCorpus</i>

## Options

Name	Description	Type
gzip	If True, each output, except annotations, for each document is gzipped. This can help reduce the storage occupied by e.g. parser or coref output. Default: False	bool
to-ken_model	Tokenization model. Specify a full path, or just a filename. If a filename is given it is expected to be in the opennlp model directory (models/opennlp/)	string
parse_model	Parser model, full path or directory name. If a filename is given, it is expected to be in the OpenNLP model directory (models/opennlp/)	string
timeout	Timeout in seconds for each individual coref resolution task. If this is exceeded, an InvalidDocument is returned for that document	int
coref_model	Coreference resolution model, full path or directory name. If a filename is given, it is expected to be in the OpenNLP model directory (models/opennlp/). Default: “ (standard English opennlp model in models/opennlp/)	string
readable	If True, pretty-print the JSON output, so it’s human-readable. Default: False	bool
pos_model	POS tagger model, full path or filename. If a filename is given, it is expected to be in the opennlp model directory (models/opennlp/)	string
sen- tence_model	Sentence segmentation model. Specify a full path, or just a filename. If a filename is given it is expected to be in the opennlp model directory (models/opennlp/)	string

## OpenNLP NER

Path	pimlico.modules.opennlp.ner
Executable	yes

Named-entity recognition using OpenNLP’s tools.

By default, uses the pre-trained English model distributed with OpenNLP. If you want to use other models (e.g. for other languages), download them from the OpenNLP website to the models dir (*models/opennlp*) and specify the model name as an option.

Note that the default model is for identifying person names only. You can identify other name types by loading other pre-trained OpenNLP NER models. Identification of multiple name types at the same time is not (yet) implemented.

## Inputs

Name	Type(s)
text	<i>TarredCorpus</i> < <i>TokenizedDocumentType</i>   <i>WordAnnotationsDocumentType</i> >

## Outputs

Name	Type(s)
documents	<i>SentenceSpansCorpus</i>

## Options

Name	Description	Type
model	NER model, full path or filename. If a filename is given, it is expected to be in the opennlp model directory (models/opennlp/)	string

## OpenNLP constituency parser

Path	pimlico.modules.opennlp.parse
Executable	yes

## Todo

Document this module

---

## Inputs

Name	Type(s)
documents	TarredCorpus<TokenizedDocumentType> or <i>WordAnnotationCorpus</i> with ‘word’ field

## Outputs

Name	Type(s)
parser	<i>ConstituencyParseTreeCorpus</i>

## Options

Name	Description	Type
model	Parser model, full path or directory name. If a filename is given, it is expected to be in the OpenNLP model directory (models/opennlp/)	string

## OpenNLP POS-tagger

Path	pimlico.modules.opennlp.pos
Executable	yes

Part-of-speech tagging using OpenNLP’s tools.

By default, uses the pre-trained English model distributed with OpenNLP. If you want to use other models (e.g. for other languages), download them from the OpenNLP website to the models dir (*models/opennlp*) and specify the model name as an option.

## Inputs

Name	Type(s)
text	TarredCorpus<TokenizedDocumentType WordAnnotationsDocumentType>

## Outputs

Name	Type(s)
documents	<i>AddAnnotationField</i>

## Options

Name	Description	Type
model	POS tagger model, full path or filename. If a filename is given, it is expected to be in the opennlp model directory (models/opennlp/)	string

## OpenNLP tokenizer

Path	pimlico.modules.opennlp.tokenize
Executable	yes

Sentence splitting and tokenization using OpenNLP's tools.

## Inputs

Name	Type(s)
text	TarredCorpus<RawTextDocumentType>

## Outputs

Name	Type(s)
documents	<i>TokenizedCorpus</i>

## Options

Name	Description	Type
to-ken_model	Tokenization model. Specify a full path, or just a filename. If a filename is given it is expected to be in the opennlp model directory (models/opennlp/)	string
tokenize_only	By default, sentence splitting is performed prior to tokenization. If tokenize_only is set, only the tokenization step is executed	bool
sen-tence_model	Sentence segmentation model. Specify a full path, or just a filename. If a filename is given it is expected to be in the opennlp model directory (models/opennlp/)	string

## 1.3.10 R interfaces

Modules for interfacing with the [statistical programming language R](#). Currently, we provide just a simple way to pass data from the output of another module into an R script and run it. In the future, it may be appropriate to add more sophisticated interfaces, or expose R's functionality in a more specialised way, integrating more closely with Pimlico's datatype system.

## R script executor

Path	pimlico.modules.r.script
Executable	yes

Simple interface to R that just involves running a given R script, first substituting in some paths from the pipeline, making it easy to pass in data from the output of other modules.

## Inputs

Name	Type(s)
sources	<i>list of PimlicoDatatype</i>

## Outputs

Name	Type(s)
output	<i>NamedFile()</i>

## Options

Name	Description	Type
script	(required) Path to the script to be run. The script itself may include substitutions of the form ‘{{inputX}}’, which will be replaced with the absolute path to the data dir of the Xth input, and ‘{{output}}’, which will be replaced with the absolute path to the output dir. The latter allows the script to output things other than the output file, which always exists and contains the full script’s output	string

## 1.3.11 Regular expressions

### Regex annotated text matcher

Path	pimlico.modules.regex.annotated_text
Executable	yes

## Todo

Document this module

---

## Inputs

Name	Type(s)
documents	TarredCorpus<WordAnnotationsDocumentType>

## Outputs

Name	Type(s)
documents	<i>KeyValueListCorpus</i>

## Options

Name	Description	Type
expr	(required) An expression to determine what to search for in sentences. Consists of a sequence of tokens, each matching one field in the corresponding token's annotations in the data. These are specified in the form <code>field[x]</code> , where <code>field</code> is the name of a field supplied by the input data and <code>x</code> is the value required of that field. If <code>x</code> ends in a <code>,</code> it will match prefixes: e.g. <code>pos[NN]</code> . If no field name is given, the default 'word' is used. A token of the form ' <code>x=y</code> ' matches the expression <code>y</code> as above and assigns the matching word to the extracted variable <code>x</code> (to be output). You may also extract a different annotation field by specifying <code>x=f:y</code> , where <code>f</code> is the field name to be extracted. E.g. 'what a=lemma:pos[NN*] lemma[come] with b=pos[NN*]' matches phrases like 'what meals come with fries', producing 'a=meal' and 'b=fries'. Both <code>pos</code> and <code>lemma</code> need to be fields in the dataset'. If you give multiple whole expressions separated by <code> </code> s, matches will be collected from all of them	string

### 1.3.12 Scikit-learn tools

Scikit-learn ('sklearn') provides easy-to-use implementations of a large number of machine-learning methods, based on Numpy/Scipy.

You can build Numpy arrays from your corpus using the *feature processing tools* and then use them as input to Scikit-learn's tools using the modules in this package.

#### Sklearn matrix factorization

Path	pimlico.modules.sklearn.matrix_factorization
Executable	yes

Provides a simple interface to Scikit-Learn's various matrix factorization models.

Since they provide a consistent training interface, you can simply choose the class name of the method you want to use and specify options relevant to that method in the `options` option. For available options, take a look at the table of parameters in the [Scikit-Learn documentation](#) for each class.

## Inputs

Name	Type(s)
matrix	<i>ScipySparseMatrix</i>

## Outputs

Name	Type(s)
w	<i>NumpyArray</i>
h	<i>NumpyArray</i>

## Options

Name	Description	Type
class	(required) Scikit-learn class to use to fit the matrix factorization. Should be the name of a class in the package <code>sklearn.decomposition</code> that has a <code>fit_transform()</code> method and a <b>components_</b> attribute. Supported classes: NMF, SparsePCA, ProjectedGradientNMF, FastICA, FactorAnalysis, PCA, RandomizedPCA, LatentDirichletAllocation, TruncatedSVD	'NMF', 'SparsePCA', 'ProjectedGradientNMF', 'FastICA', 'FactorAnalysis', 'PCA', 'RandomizedPCA', 'LatentDirichletAllocation' or 'TruncatedSVD'
options	Options to pass into the constructor of the sklearn class, formatted as a JSON dictionary (potentially without the <code>{}</code> s). E.g.: <code>'n_components=200, solver="cd", tol=0.0001, max_iter=200'</code>	string

### 1.3.13 General utilities

General utilities for things like filesystem manipulation.

#### Module output alias

Path	<code>pimlico.modules.utility.alias</code>
Executable	no

Alias a datatype coming from the output of another module.

Used to assign a handy identifier to the output of a module, so that we can just refer to this alias module later in the pipeline and use its default output. This can help make for a more readable pipeline config.

For example, say we use `split` to split a dataset into two random subsets. The two splits can be accessed by referring to the two outputs of that module: `split_module.set1` and `split_module.set2`. However, it's easy to lose track of what these splits are supposed to be used for, so we might want to give them names:

```
[split_module]
type=pimlico.modules.corpora.split
set1_size=0.2

[test_set]
type=pimlico.modules.utility.alias
input=split_module.set1

[training_set]
type=pimlico.modules.utility.alias
input=split_module.set2

[training_routine]
type=...
input_corpus=training_set
```

Note the difference between using this module and using the special *alias* module type. The *alias* type creates an alias for a whole module, allowing you to refer to all of its outputs, inherit its settings, and anything else you could do with the original module name. This module, however, provides an alias for exactly one output of a module and generates a module instance of its own in the pipeline (albeit a filter module).



This is a filter module. It is not executable, so won't appear in a pipeline's list of modules that can be run. It produces its output for the next module on the fly when the next module needs it.

### Inputs

Name	Type(s)
input	<i>PimlicoDatatype</i>

### Outputs

Name	Type(s)
output	same as input corpus

### Copy file

Path	pimlico.modules.utility.copy_file
Executable	yes

Copy a file

Simple utility for copying a file (which presumably comes from the output of another module) into a particular location. Useful for collecting together final output at the end of a pipeline.

### Inputs

Name	Type(s)
source	<i>list of File</i>

### Outputs

Name	Type(s)
documents	<i>TarredCorpus</i>

### Options

Name	Description	Type
tar-get_name	Name to rename the target file to. If not given, it will have the same name as the source file. Ignored if there's more than one input file	string
tar-get_dir	(required) Path to directory into which the file should be copied. Will be created if it doesn't exist	string

## 1.3.14 Visualization tools

Modules for plotting and suchlike

## Bar chart plotter

Path	pimlico.modules.visualization.bar_chart
Executable	yes

### Inputs

Name	Type(s)
values	<i>list</i> of <i>NumericResult</i>

### Outputs

Name	Type(s)
plot	<i>PlotOutput</i>

## Embedding space plotter

Path	pimlico.modules.visualization.embeddings_plot
Executable	yes

Plot vectors from embeddings, trained by some other module, in a 2D space using a MDS reduction and Matplotlib.

They might, for example, come from *pimlico.modules.embeddings.word2vec*. The embeddings are read in using Pimlico's generic word embedding storage type.

Uses scikit-learn to perform the MDS/TSNE reduction.

### Inputs

Name	Type(s)
vectors	<i>list</i> of <i>Embeddings</i>

### Outputs

Name	Type(s)
plot	<i>PlotOutput</i>

## Options

Name	Description	Type
skip	Number of most frequent words to skip, taking the next most frequent after these. Default: 0	int
metric	Distance metric to use. Choose from 'cosine', 'euclidean', 'manhattan'. Default: 'cosine'	'cosine', 'euclidean' or 'manhattan'
reduction	Dimensionality reduction technique to use to project to 2D. Available: mds (Multi-dimensional Scaling), tsne (t-distributed Stochastic Neighbor Embedding). Default: mds	'mds' or 'tsne'
colours	List of colours to use for different embedding sets. Should be a list of matplotlib colour strings, one for each embedding set given in input_vectors	comma-separated list of strings
cmap	Mapping from word prefixes to matplotlib plotting colours. Every word beginning with the given prefix has the prefix removed and is plotted in the corresponding colour. Specify as a JSON dictionary mapping prefix strings to colour strings	JSON string
words	Number of most frequent words to plot. Default: 50	int

## 1.4 Command-line interface

The main Pimlico command-line interface (usually accessed via *pimlico.sh* in your project root) provides subcommands to perform different operations. Call it like so, using one of the subcommands documented below to access particular functionality:

```
./pimlico.sh <config-file> [general options...] <subcommand> [subcommand args/options]
```

The commands you are likely to use most often are: *status*, *run*, *reset* and maybe *browse*.

For a reference for each command's options, see the command-line documentation: `./pimlico.sh --help`, for a general reference and `./pimlico.sh <config_file> <command> --help` for a specific subcommand's reference.

Below is a more detailed guide for each subcommand, including all of the documentation available via the command line.

<i>browse</i>	View the data output by a module
<i>clean</i>	Remove all module output directories that do not correspond to a module in the pipeline
<i>deps</i>	List information about software dependencies: whether they're available, versions, etc
<i>dump</i>	Dump the entire available output data from a given pipeline module to a tarball
<i>email</i>	Test email settings and try sending an email using them
<i>inputs</i>	Show the (expected) locations of the inputs of a given module
<i>install</i>	Install missing module library dependencies
<i>load</i>	Load a module's output data from a tarball previously created by the dump command
<i>longstore</i>	Move a particular module's output from the short-term store to the long-term store
<i>newmodule</i>	Create a new module type
<i>output</i>	Show the location where the given module's output data will be (or has been) stored
<i>python</i>	Load the pipeline config and enter a Python interpreter with access to it in the environment
<i>reset</i>	Delete any output from the given module and restore it to unexecuted state
<i>run</i>	Execute an individual pipeline module, or a sequence
<i>shell</i>	Open a shell to give access to the data output by a module
<i>status</i>	Output a module execution schedule for the pipeline and execution status for every module
<i>unlock</i>	Forcibly remove an execution lock from a module
<i>variants</i>	List the available variants of a pipeline config
<i>visualize</i>	Coming soon...visualize the pipeline in a pretty way

### 1.4.1 status

*Command-line tool subcommand*

Output a module execution schedule for the pipeline and execution status for every module.

Usage:

```
pimlico.sh [...] status [module_name] [-h] [--all] [--short] [--history] [--deps-of_
↪DEPS_OF] [--no-color]
```

#### Positional arguments

Arg	Description
[module_name]	Optionally specify a module name (or number). More detailed status information will be output for this module. Alternatively, use this arg to limit the modules whose status will be output to a range by specifying 'A...B', where A and B are module names or numbers

#### Options

Option	Description
--all, -a	Show all modules defined in the pipeline, not just those that can be executed
--short, -s	Use a brief format when showing the full pipeline's status. Only applies when module names are not specified. This is useful with very large pipelines, where you just want a compact overview of the status
--history, -i	When a module name is given, even more detailed output is given, including the full execution history of the module
--deps-of, -d	Restrict to showing only the named/numbered module and any that are (transitive) dependencies of it. That is, show the whole tree of modules that lead through the pipeline to the given module
--no-color, --nc	Don't include terminal color characters, even if the terminal appears to support them. This can be useful if the automatic detection of color terminals doesn't work and the status command displays lots of horrible escape characters

### 1.4.2 variants

*Command-line tool subcommand*

List the available variants of a pipeline config.

Usage:

```
pimlico.sh [...] variants [-h]
```

### 1.4.3 run

*Command-line tool subcommand*

Main command for executing Pimlico modules from the command line *run* command.

Usage:

```
pimlico.sh [...] run [modules [modules ...]] [-h] [--force-rerun] [--all-deps] [--  
↪all] [--dry-run] [--step] [--preliminary] [--exit-on-error] [--email {modend,end}]
```

### Positional arguments

Arg	Description
[modules [modules ...]]	The name (or number) of the module to run. To run a stage from a multi-stage module, use ‘module:stage’. Use ‘status’ command to see available modules. Use ‘module:?’ or ‘module:help’ to list available stages. If not given, defaults to next incomplete module that has all its inputs ready. You may give multiple modules, in which case they will be executed in the order specified

## Options

Option	Description
<code>--force-rerun</code> , <code>-f</code>	Force running the module(s), even if it's already been run to completion
<code>--all-deps</code> , <code>-a</code>	If the given module(s) has dependent modules that have not been completed, executed them first. This allows you to specify a module late in the pipeline and execute the full pipeline leading to that point
<code>--all</code>	Run all currently unexecuted modules that have their inputs ready, or will have by the time previous modules are run. (List of modules will be ignored)
<code>--dry-run</code> , <code>--dry</code> , <code>--check</code>	Perform all pre-execution checks, but don't actually run the module(s)
<code>--step</code>	Enabled super-verbose debugging mode, which steps through a module's processing outputting a lot of information and allowing you to control the output as it goes. Useful for working out what's going on inside a module if it's mysteriously not producing the output you expected
<code>--preliminary</code> , <code>--pre</code>	Perform a preliminary run of any modules that take multiple datasets into one of their inputs. This means that we will run the module even if not all the datasets are yet available (but at least one is) and mark it as preliminarily completed
<code>--exit-on-error</code> , <code>-e</code>	If an error is encountered while executing a module that causes the whole module execution to fail, output the error and exit. By default, Pimlico will send error output to a file (or print it in debug mode) and continue to execute the next module that can be executed, if any
<code>--email</code>	Send email notifications when processing is complete, including information about the outcome. Choose from: 'modend' (send notification after module execution if it fails and a summary at the end of everything), 'end' (send only the final summary). Email sending must be configured: see 'email' command to test

### 1.4.4 browse

*Command-line tool subcommand*

View the data output by a module.

Usage:

```
pimlico.sh [...] browse module_name [output_name] [-h] [--raw] [--skip-invalid] [--  
↪formatter FORMATTER]
```

#### Positional arguments

Arg	Description
<code>module_name</code>	The name (or number) of the module whose output to look at. Use 'module:stage' for multi-stage modules
<code>[output_name]</code>	The name of the output from the module to browse. If blank, load the default output

## Options

Option	Description
<code>--raw, -r</code>	Don't parse the data using the output datatype (i.e. just read the raw text). If not set, we output the result of applying <code>unicode()</code> to the parsed data structure, or a custom formatting if the datatype loaded defines one
<code>--skip-invalid</code>	Skip over invalid documents, instead of showing the error that caused them to be invalid
<code>--formatter</code> <code>-f</code>	Fully qualified class name of a subclass of <code>DocumentBrowserFormatter</code> to use to determine what to output for each document. If specified, <code>--raw</code> is ignored. You may also choose from the named standard formatters for the datatype in question. Use <code>'-f help'</code> to see a list of available formatters

### 1.4.5 shell

*Command-line tool subcommand*

Open a shell to give access to the data output by a module.

Usage:

```
pimlico.sh [...] shell module_name [output_name] [-h]
```

#### Positional arguments

Arg	Description
<code>module_name</code>	The name (or number) of the module whose output to look at
<code>[output_name]</code>	The name of the output from the module to browse. If blank, load the default output

### 1.4.6 python

*Command-line tool subcommand*

Load the pipeline config and enter a Python interpreter with access to it in the environment.

Usage:

```
pimlico.sh [...] python [script] [-h] [-i]
```

#### Positional arguments

Arg	Description
<code>[script]</code>	Script file to execute. Omit to enter interpreter

## Options

Option	Description
<code>-i</code>	Enter interactive shell after running script

### 1.4.7 reset

*Command-line tool subcommand*

Delete any output from the given module and restore it to unexecuted state.

Usage:

```
pimlico.sh [...] reset [modules [modules ...]] [-h] [-n]
```

#### Positional arguments

Arg	Description
[modules [modules ...]]	The names (or numbers) of the modules to reset, or ‘all’ to reset the whole pipeline

#### Options

Option	Description
-n, --no-deps	Only reset the state of this module, even if it has dependent modules in an executed state, which could be invalidated by resetting and re-running this one

### 1.4.8 clean

*Command-line tool subcommand*

Cleans up module output directories that have got left behind.

Often, when developing a pipeline incrementally, you try out some modules, but then remove them, or rename them to something else. The directory in the Pimlico output store that was created to contain their metadata, status and output data is then left behind and no longer associated with any module.

Run this command to check all storage locations for such directories. If it finds any, it prompts you to confirm before deleting them. (If there are things in the list that don’t look like they were left behind by the sort of things mentioned above, don’t delete them! I don’t want you to lose your precious output data if I’ve made a mistake in this command.)

Note that the operation of this command is specific to the loaded pipeline variant. If you have multiple variants, make sure to select the one you want to clean with the general *-variant* option.

Usage:

```
pimlico.sh [...] clean [-h]
```

### 1.4.9 longstore

*Command-line tool subcommand*

Move a particular module’s output from the short-term store to the long-term store. It will still be found here by input readers. You might want to do this if your long-term store is bigger, to keep down the short-term store size.

Usage:

```
pimlico.sh [...] longstore [modules [modules ...]] [-h]
```



## Positional arguments

Arg	Description
[modules [modules ...]]	The names (or numbers) of the module whose output to move

### 1.4.10 unlock

#### *Command-line tool subcommand*

Forcibly remove an execution lock from a module. If a lock has ended up getting left on when execution exited prematurely, use this to remove it.

When a module starts running, it is locked to avoid making a mess of your output data by running the same module from another terminal, or some other silly mistake (I know, for some of us this sort of behaviour is frustratingly common).

Usually shouldn't be necessary, even if there's an error during execution, since the module should be unlocked when Pimlico exits, but occasionally (e.g. if you have to forcibly kill Pimlico during execution) the lock gets left on.

Usage:

```
pimlico.sh [...] unlock module_name [-h]
```

## Positional arguments

Arg	Description
module_name	The name (or number) of the module to unlock

### 1.4.11 dump

#### *Command-line tool subcommand*

Dump the entire available output data from a given pipeline module to a tarball, so that it can easily be loaded into the same pipeline on another system. This is primarily to support spreading the execution of a pipeline between multiple machines, so that the output from a module can easily be transferred and loaded into a pipeline.

Dump to a tarball using this command, transfer the file between machines and then run the [load command](#) to import it there.

**See also:**

[Running one pipeline on multiple computers](#): for a more detailed guide to transferring data across servers.

Usage:

```
pimlico.sh [...] dump [modules [modules ...]] [-h] [--output OUTPUT]
```

## Positional arguments

Arg	Description
[modules [modules ...]]	Names or numbers of modules whose data to dump. If multiple are given, a separate file will be dumped for each

## Options

Option	Description
<code>--output, -o</code>	Path to directory to output to. Defaults to the current user's home directory

### 1.4.12 load

#### *Command-line tool subcommand*

Load the output data for a given pipeline module from a tarball previously created by the *dump* command (typically on another machine). This is primarily to support spreading the execution of a pipeline between multiple machines, so that the output from a module can easily be transferred and loaded into a pipeline.

Dump to a tarball using the *dump command*, transfer the file between machines and then run this command to import it there.

#### See also:

*Running one pipeline on multiple computers*: for a more detailed guide to transferring data across servers.

Usage:

```
pimlico.sh [...] load [paths [paths ...]] [-h] [--force-overwrite]
```

## Positional arguments

Arg	Description
<code>[paths [paths ...]]</code>	Paths to dump files (tarballs) to load into the pipeline

## Options

Option	Description
<code>--force-overwrite</code> <code>-f</code>	If data already exists for a module being imported, overwrite without asking. By default, the user will be prompted to check whether they want to overwrite

### 1.4.13 deps

#### *Command-line tool subcommand*

Output information about module dependencies.

Usage:

```
pimlico.sh [...] deps [modules [modules ...]] [-h]
```

## Positional arguments

Arg	Description
<code>[modules [modules ...]]</code>	Check dependencies for named modules and install any that are automatically installable. Use 'all' to install dependencies for all modules

### 1.4.14 install

*Command-line tool subcommand*

Install missing dependencies.

Usage:

```
pimlico.sh [...] install [modules [modules ...]] [-h] [--trust-downloaded]
```

#### Positional arguments

Arg	Description
[modules [modules ...]]	Check dependencies for named modules and install any that are automatically installable. Use ‘all’ to install dependencies for all modules

#### Options

Option	Description
--trust-downloaded -t	If an archive file to be downloaded is found to be in the lib dir already, trust that it is the file we’re after. By default, we only reuse archives we’ve just downloaded, so we know they came from the right URL, avoiding accidental name clashes

### 1.4.15 inputs

*Command-line tool subcommand*

Show the locations of the inputs of a given module. If the input datasets are available, their actual location is shown. Otherwise, all directories in which the data is being checked for are shown.

Usage:

```
pimlico.sh [...] inputs module_name [-h]
```

#### Positional arguments

Arg	Description
module_name	The name (or number) of the module to display input locations for

### 1.4.16 output

*Command-line tool subcommand*

Show the location where the given module’s output data will be (or has been) stored.

Usage:

```
pimlico.sh [...] output module_name [-h]
```

## Positional arguments

Arg	Description
module_name	The name (or number) of the module to display input locations for

### 1.4.17 newmodule

#### *Command-line tool subcommand*

Interactive tool to create a new module type, generating a skeleton for the module's code. Currently only works for certain module types. May be extended in future to help with creating a broader range of sorts of modules.

Usage:

```
pimlico.sh [...] newmodule [-h]
```

### 1.4.18 visualize

#### *Command-line tool subcommand*

(Not yet fully implemented!) Visualize the pipeline, with status information for modules.

Usage:

```
pimlico.sh [...] visualize [-h] [--all]
```

## Options

Option	Description
--all, -a	Show all modules defined in the pipeline, not just those that can be executed

### 1.4.19 email

#### *Command-line tool subcommand*

Test email settings and try sending an email using them.

Usage:

```
pimlico.sh [...] email [-h]
```

## 1.5 API Documentation

API documentation for the main Pimlico codebase, excluding the *built-in Pimlico module types*.

## 1.5.1 pimlico package

### Subpackages

#### pimlico.cli package

### Subpackages

#### pimlico.cli.browser package

### Submodules

#### pimlico.cli.browser.formatter module

The command-line iterable corpus browser displays one document at a time. It can display the raw data from the corpus files, which sometimes is sufficiently human-readable to not need any special formatting. It can also parse the data using its datatype and output text either from the datatype's standard unicode representation or, if the document datatype provides it, a special browser formatting of the data.

When viewing output data, particularly during debugging of modules, it can be useful to provide special formatting routines to the browser, rather than using or overriding the datatype's standard formatting methods. For example, you might want to pull out specific attributes for each document to get an overview of what's coming out.

The browser command accepts a command-line option that specifies a Python class to format the data. This class should be a subclass of :class:pimlico.cli.browser.formatter.DocumentBrowserFormatter that accepts a datatype compatible with the datatype being browsed and provides a method to format each document. You can write these in your custom code and refer to them by their fully qualified class name.

**class DocumentBrowserFormatter** (*corpus*)

Bases: object

Base class for formatters used to post-process documents for display in the iterable corpus browser.

#### DATATYPE

Whether the formatter expects to get the parsed document, according to the datatype's usual input reading routine, or just the raw text stored in the file. Usually it's the former (default), but subclasses may override this.

alias of `DataPointType`

#### RAW\_INPUT = False

#### format\_document (doc)

Format a single document and return the result as a string (or unicode, but it will be converted to ASCII for display).

Must be overridden by subclasses.

#### filter\_document (doc)

Each doc is passed through this function directly after being read from the corpus. If None is returned, the doc is skipped. Otherwise, the result is used instead of the doc data. The default implementation does nothing.

**class DefaultFormatter** (*corpus*, *raw\_data=False*)

Bases: `pimlico.cli.browser.formatter.DocumentBrowserFormatter`

Generic implementation of a browser formatter that's used if no other formatter is given.

**DATATYPE**alias of `IterableCorpus`**format\_document** (*doc*)**class InvalidDocumentFormatter** (*corpus*)Bases: `pimlico.cli.browser.formatter.DocumentBrowserFormatter`

Formatter that skips over all docs other than invalid results. Uses standard formatting for InvalidDocument information.

**DATATYPE**alias of `IterableCorpus`**format\_document** (*doc*)**filter\_document** (*doc*)**typecheck\_formatter** (*formatted\_doc\_type, formatter\_cls*)

Check that a document type is compatible with a particular formatter.

**load\_formatter** (*dataset, formatter\_name=None, parse=True*)

Load a formatter specified by its fully qualified Python class name. If None, loads the default formatter. You may also specify a formatter by name, choosing from one of the standard ones that the formatted datatype gives.

**Parameters**

- **formatter\_name** – class name, or class
- **dataset** – dataset that will be formatted
- **parse** – only used if the default formatter is loaded, determines *raw\_data* (= *not parse*)

**Returns** instantiated formatter**pimlico.cli.browser.tool module**

Tool for browsing datasets, reading from the data output by pipeline modules.

**browse\_cmd** (*pipeline, opts*)

Command for main Pimlico CLI

**browse\_data** (*data, formatter, parse=False, skip\_invalid=False*)**class CorpusState** (*corpus*)Bases: `object`

Keep track of which document we're on.

**next\_document** ()**skip** (*n*)**class InputDialog** (*text, input\_edit*)Bases: `urwid.widget.WidgetWrap`

A dialog that appears with an input

**signals** = ['close', 'cancel']**keypress** (*size, k*)**class MessageDialog** (*text, default=None*)Bases: `urwid.widget.WidgetWrap`

A dialog that appears with a message

```
class InputPopupLauncher (original_widget, text, input_edit, callback=None)
    Bases: urwid.wimp.PopUpLauncher

    create_pop_up ()

    get_pop_up_parameters ()

skip_popup_launcher (original_widget, text, default=None, callback=None)

save_popup_launcher (original_widget, text, default=None, callback=None)

class MessagePopupLauncher (original_widget, text)
    Bases: urwid.wimp.PopUpLauncher

    create_pop_up ()

    get_pop_up_parameters ()
```

## Module contents

### pimlico.cli.debug package

#### Submodules

#### pimlico.cli.debug.stepper module

##### **class Stepper**

Bases: object

Type that stores the state of the stepping process. This allows information and parameters to be passed around through the process and updated as we go. For example, if particular type of output is disabled by the user, a parameter can be updated here so we know not to output it later.

##### **enable\_step\_for\_pipeline** (*pipeline*)

Prepares a pipeline to run in step mode, modifying modules and wrapping methods to supply the extra functionality.

This approach means that we don't have to consume extra computation time checking whether step mode is enabled during normal runs.

**Parameters** *pipeline* – instance of PipelineConfig

**instantiate\_output\_datatype\_decorator** (*instantiate\_output\_datatype, module\_name, output\_names, stepper*)

**wrap\_tarred\_corpus** (*dtype, module\_name, output\_name, stepper*)

**archive\_iter\_decorator** (*archive\_iter, module\_name, output\_name, stepper*)

**get\_input\_decorator** (*get\_input, module\_name, stepper*)

Decorator to wrap a module info's `get_input()` method so when know where inputs are being used.

**option\_message** (*message\_lines, stepper, options=None, stack\_trace\_option=True, category=None*)

## Module contents

Extra-verbose debugging facility

Tools for very slowly and verbosely stepping through the processing that a given module does to debug it.

Enabled using the `-step` switch to the run command.

**fmt\_frame\_info** (*info*)

**output\_stack\_trace** (*frame=None*)

## pimlico.cli.shell package

### Submodules

#### pimlico.cli.shell.base module

**class ShellCommand**

Bases: `object`

Base class used to provide commands for exploring a particular datatype. A basic set of commands is provided for all datatypes, but specific datatype classes may provide their own, by overriding the `shell_commands` attribute.

**commands** = []

**help\_text** = None

**execute** (*shell, \*args, \*\*kwargs*)

**class DataShell** (*data, commands, \*args, \*\*kwargs*)

Bases: `cmd.Cmd`

Terminal shell for querying datatypes.

**prompt** = '>>> '

**get\_names** ()

**do\_EOF** (*line*)

Exits the shell

**preloop** ()

**postloop** ()

**emptyline** ()

Don't repeat the last command (default): ignore empty lines

**default** (*line*)

We use this to handle commands that can't be handled using the `do_` pattern. Also handles the default fallback, which is to execute Python.

**cmdloop** (*intro=None*)

**exception ShellError**

Bases: `exceptions.Exception`

#### pimlico.cli.shell.commands module

Basic set of shell commands that are always available.

**class MetadataCmd**

Bases: `pimlico.cli.shell.base.ShellCommand`



```
    commands = ['metadata']
    help_text = "Display the loaded dataset's metadata"
    execute(shell, *args, **kwargs)

class PythonCmd
    Bases: pimlico.cli.shell.base.ShellCommand
    commands = ['python', 'py']
    help_text = "Run a Python interpreter using the current environment, including import availability of all the project c
    execute(shell, *args, **kwargs)
```

## pimlico.cli.shell.runner module

```
class ShellCLICmd
    Bases: pimlico.cli.subcommands.PimlicoCLISubcommand
    command_name = 'shell'
    command_help = 'Open a shell to give access to the data output by a module'
    add_arguments(parser)
    run_command(pipeline, opts)

launch_shell(data)
    Starts a shell to view and query the given datatype instance.
```

## Module contents

### Submodules

## pimlico.cli.check module

```
class InstallCmd
    Bases: pimlico.cli.subcommands.PimlicoCLISubcommand
    Install missing dependencies.
    command_name = 'install'
    command_help = 'Install missing module library dependencies'
    add_arguments(parser)
    run_command(pipeline, opts)

class DepsCmd
    Bases: pimlico.cli.subcommands.PimlicoCLISubcommand
    Output information about module dependencies.
    command_name = 'deps'
    command_help = "List information about software dependencies: whether they're available, versions, etc"
    add_arguments(parser)
    run_command(pipeline, opts)
```

## pimlico.cli.clean module

### class CleanCmd

Bases: `pimlico.cli.subcommands.PimlicoCLISubcommand`

Cleans up module output directories that have got left behind.

Often, when developing a pipeline incrementally, you try out some modules, but then remove them, or rename them to something else. The directory in the Pimlico output store that was created to contain their metadata, status and output data is then left behind and no longer associated with any module.

Run this command to check all storage locations for such directories. If it finds any, it prompts you to confirm before deleting them. (If there are things in the list that don't look like they were left behind by the sort of things mentioned above, don't delete them! I don't want you to lose your precious output data if I've made a mistake in this command.)

Note that the operation of this command is specific to the loaded pipeline variant. If you have multiple variants, make sure to select the one you want to clean with the general `-variant` option.

**command\_name** = 'clean'

**command\_help** = 'Remove all module directories that do not correspond to a module in the pipeline in all storage locations'

**command\_desc** = 'Remove all module output directories that do not correspond to a module in the pipeline'

**run\_command** (*pipeline*, *opts*)

## pimlico.cli.loaddump module

### class DumpCmd

Bases: `pimlico.cli.subcommands.PimlicoCLISubcommand`

Dump the entire available output data from a given pipeline module to a tarball, so that it can easily be loaded into the same pipeline on another system. This is primarily to support spreading the execution of a pipeline between multiple machines, so that the output from a module can easily be transferred and loaded into a pipeline.

Dump to a tarball using this command, transfer the file between machines and then run the *load command* to import it there.

**See also:**

*Running one pipeline on multiple computers:* for a more detailed guide to transferring data across servers

**command\_name** = 'dump'

**command\_help** = 'Dump the entire available output data from a given pipeline module to a tarball, so that it can easily be loaded into a pipeline on another system'

**command\_desc** = 'Dump the entire available output data from a given pipeline module to a tarball'

**add\_arguments** (*parser*)

**run\_command** (*pipeline*, *opts*)

### class LoadCmd

Bases: `pimlico.cli.subcommands.PimlicoCLISubcommand`

Load the output data for a given pipeline module from a tarball previously created by the *dump* command (typically on another machine). This is primarily to support spreading the execution of a pipeline between multiple machines, so that the output from a module can easily be transferred and loaded into a pipeline.

Dump to a tarball using the *dump command*, transfer the file between machines and then run this command to import it there.

See also:

*Running one pipeline on multiple computers*: for a more detailed guide to transferring data across servers

`command_name = 'load'`

`command_help = "Load a module's output data from a tarball previously created by the dump command, usually on a d"`

`command_desc = "Load a module's output data from a tarball previously created by the dump command"`

`add_arguments (parser)`

`run_command (pipeline, opts)`

## pimlico.cli.locations module

`class InputsCmd`

Bases: `pimlico.cli.subcommands.PimlicoCLISubcommand`

`command_name = 'inputs'`

`command_help = 'Show the locations of the inputs of a given module. If the input datasets are available, their actual loca'`

`command_desc = 'Show the (expected) locations of the inputs of a given module'`

`add_arguments (parser)`

`run_command (pipeline, opts)`

`class OutputCmd`

Bases: `pimlico.cli.subcommands.PimlicoCLISubcommand`

`command_name = 'output'`

`command_help = "Show the location where the given module's output data will be (or has been) stored"`

`add_arguments (parser)`

`run_command (pipeline, opts)`

## pimlico.cli.main module

Main command-line script for running Pimlico, typically called from *pimlico.sh*.

Provides access to many subcommands, acting as the primary interface to Pimlico's functionality.

`class VariantsCmd`

Bases: `pimlico.cli.subcommands.PimlicoCLISubcommand`

`command_name = 'variants'`

`command_help = 'List the available variants of a pipeline config'`

`add_arguments (parser)`

`run_command (pipeline, opts)`

`class LongStoreCmd`

Bases: `pimlico.cli.subcommands.PimlicoCLISubcommand`

`command_name = 'longstore'`

`command_help = "Move a particular module's output from the short-term store to the long-term store. It will still be fo"`

`command_desc = "Move a particular module's output from the short-term store to the long-term store"`

```
add_arguments (parser)
```

```
run_command (pipeline, opts)
```

```
class UnlockCmd
```

```
Bases: pimlico.cli.subcommands.PimlicoCLISubcommand
```

Forcibly remove an execution lock from a module. If a lock has ended up getting left on when execution exited prematurely, use this to remove it.

When a module starts running, it is locked to avoid making a mess of your output data by running the same module from another terminal, or some other silly mistake (I know, for some of us this sort of behaviour is frustratingly common).

Usually shouldn't be necessary, even if there's an error during execution, since the module should be unlocked when Pimlico exits, but occasionally (e.g. if you have to forcibly kill Pimlico during execution) the lock gets left on.

```
command_name = 'unlock'
```

```
command_help = "Forcibly remove an execution lock from a module. If a lock has ended up getting left on when executi
```

```
command_desc = 'Forcibly remove an execution lock from a module'
```

```
add_arguments (parser)
```

```
run_command (pipeline, opts)
```

```
class BrowseCmd
```

```
Bases: pimlico.cli.subcommands.PimlicoCLISubcommand
```

```
command_name = 'browse'
```

```
command_help = 'View the data output by a module'
```

```
add_arguments (parser)
```

```
run_command (pipeline, opts)
```

```
class VisualizeCmd
```

```
Bases: pimlico.cli.subcommands.PimlicoCLISubcommand
```

```
command_name = 'visualize'
```

```
command_help = '(Not yet fully implemented!) Visualize the pipeline, with status information for modules'
```

```
command_desc = 'Comming soon...visualize the pipeline in a pretty way'
```

```
add_arguments (parser)
```

```
run_command (pipeline, opts)
```

## pimlico.cli.newmodule module

```
class NewModuleCmd
```

```
Bases: pimlico.cli.subcommands.PimlicoCLISubcommand
```

```
command_name = 'newmodule'
```

```
command_help = "Interactive tool to create a new module type, generating a skeleton for the module's code. Currently o
```

```
command_desc = 'Create a new module type'
```

```
run_command (pipeline, opts)
```

```
ask (prompt, strip_space=True)
```

## pimlico.cli.pyshell module

### class PimlicoPythonShellContext

Bases: `object`

A class used as a static global data structure to provide access to the loaded pipeline when running the Pimlico Python shell command.

This should never be used in any other context to pass around loaded pipelines or other global data. We don't do that sort of thing.

### class PythonShellCmd

Bases: `pimlico.cli.subcommands.PimlicoCLISubcommand`

**command\_name** = 'python'

**command\_help** = 'Load the pipeline config and enter a Python interpreter with access to it in the environment'

**add\_arguments** (*parser*)

**run\_command** (*pipeline, opts*)

### get\_pipeline ()

This function may be used in scripts that are expected to be run exclusively from the Pimlico Python shell command (`python`) to get hold of the pipeline that was specified on the command line and loaded when the shell was started.

### exception ShellContextError

Bases: `exceptions.Exception`

## pimlico.cli.reset module

### class ResetCmd

Bases: `pimlico.cli.subcommands.PimlicoCLISubcommand`

**command\_name** = 'reset'

**command\_help** = 'Delete any output from the given module and restore it to unexecuted state'

**add\_arguments** (*parser*)

**run\_command** (*pipeline, opts*)

## pimlico.cli.run module

### class RunCmd

Bases: `pimlico.cli.subcommands.PimlicoCLISubcommand`

Main command for executing Pimlico modules from the command line *run* command.

**command\_name** = 'run'

**command\_help** = 'Execute an individual pipeline module, or a sequence'

**add\_arguments** (*parser*)

**run\_command** (*pipeline, opts*)

### pimlico.cli.status module

**class** `StatusCmd`

Bases: `pimlico.cli.subcommands.PimlicoCLISubcommand`

**command\_name** = 'status'

**command\_help** = 'Output a module execution schedule for the pipeline and execution status for every module'

**add\_arguments** (*parser*)

**run\_command** (*pipeline, opts*)

**module\_status\_color** (*module*)

**status\_colored** (*module, text=None*)

Colour the text according to the status of the given module. If text is not given, the module's name is returned.

**module\_status** (*module*)

Detailed module status, shown when a specific module's status is requested.

### pimlico.cli.subcommands module

**class** `PimlicoCLISubcommand`

Bases: `object`

Base class for defining subcommands to the main command line tool.

This allows us to split up subcommands, together with all their arguments/options and their functionality, since there are quite a lot of them.

Documentation of subcommands should be supplied in the following ways:

- Include help texts for positional args and options in the `add_arguments()` method. They will all be included in the doc page for the command.
- Write a very short description of what the command is for (a few words) in `command_desc`. This will be used in the summary table / TOC in the docs.
- Write a short description of what the command does in `command_help`. This will be available in command-line help and used as a fallback if you don't do the next point.
- Write a good guide to using the command (or at least say what it does) in the class' docstring (i.e. overriding this). This will form the bulk of the command's doc page.

**command\_name** = `None`

**command\_help** = `None`

**command\_desc** = `None`

**add\_arguments** (*parser*)

**run\_command** (*pipeline, opts*)

### pimlico.cli.testemail module

**class** `EmailCmd`

Bases: `pimlico.cli.subcommands.PimlicoCLISubcommand`

**command\_name** = 'email'

```
command_help = 'Test email settings and try sending an email using them'
```

```
run_command(pipeline, opts)
```

## pimlico.cli.util module

```
module_number_to_name(pipeline, name)
```

```
module_numbers_to_names(pipeline, names)
```

Convert module numbers to names, also handling ranges of numbers (and names) specified with "...". Any "..." will be filled in by the sequence of intervening modules.

Also, if an unexpanded module name is specified for a module that's been expanded into multiple corresponding to alternative parameters, all of the expanded module names are inserted in place of the unexpanded name.

```
format_execution_error(error)
```

Produce a string with lots of error output to help debug a module execution error.

**Parameters** **error** – the exception raised (ModuleExecutionError or ModuleInfoLoadError)

**Returns** formatted output

```
print_execution_error(error)
```

## Module contents

### pimlico.core package

#### Subpackages

### pimlico.core.dependencies package

#### Submodules

### pimlico.core.dependencies.base module

Base classes for defining software dependencies for module types and routines for fetching them.

```
class SoftwareDependency(name, url=None, dependencies=None)
```

Bases: object

Base class for all Pimlico module software dependencies.

```
available(local_config)
```

Return True if the dependency is satisfied, meaning that the software/library is installed and ready to use.

```
problems(local_config)
```

Returns a list of problems standing in the way of the dependency being available. If the list is empty, the dependency is taken to be installed and ready to use.

Overriding methods should call super method.

```
installable()
```

Return True if it's possible to install this library automatically. If False, the user will have to install it themselves. Instructions for doing this may be provided by `installation_instructions()`, which will only generally be called if `installable()` returns False.

This might be the case, for example, if the software is not available to download freely, or if it requires a system-wide installation.

**installation\_instructions ()**

Where a dependency can't be installed programmatically, we typically want to be able to output instructions for the user to tell them how to go about doing it themselves. Any subclass that doesn't provide an automatic installation routine should override this to provide instructions.

You may also provide this even if the class does provide automatic installation. For example, you might want to provide instructions for other ways to install the software, like a system-wide install. This instructions will be shown together with missing dependency information.

**dependencies ()**

Returns a list of instances of `:class:SoftwareDependency` subclasses representing this library's own dependencies. If the library is already available, these will never be consulted, but if it is to be installed, we will check first that all of these are available (and try to install them if not).

**install (local\_config, trust\_downloaded\_archives=False)**

Should be overridden by any subclasses whose library is automatically installable. Carries out the actual installation.

You may assume that all dependencies returned by `:method:dependencies` have been satisfied prior to calling this.

**all\_dependencies ()**

Recursively fetch all dependencies of this dependency (not including itself).

**get\_installed\_version (local\_config)**

If `available()` returns True, this method should return a `SoftwareVersion` object (or subclass) representing the software's version.

The base implementation returns an object representing an unknown version number.

If `available()` returns False, the behaviour is undefined and may raise an error.

**class SystemCommandDependency (name, test\_command, \*\*kwargs)**

Bases: `pimlico.core.dependencies.base.SoftwareDependency`

Dependency that tests whether a command is available on the command line. Generally requires system-wide installation.

**installable ()**

Usually not automatically installable

**problems (local\_config)****exception InstallationError**

Bases: `exceptions.Exception`

**check\_and\_install (deps, local\_config, trust\_downloaded\_archives=False)**

Check whether dependencies are available and try to install those that aren't. Returns a list of dependencies that can't be installed.

**install (dep, local\_config, trust\_downloaded\_archives=False)****install\_dependencies (pipeline, modules=None, trust\_downloaded\_archives=True)**

Install dependencies for pipeline modules

**Parameters**

- **pipeline** –
- **modules** – list of module names, or None to install for all



**Returns****recursive\_deps** (*dep*)

Collect all recursive dependencies of this dependency. Does a depth-first search so that everything comes later in the list than things it depends on.

**pimlico.core.dependencies.core module**

Basic Pimlico core dependencies

**CORE\_PIMLICO\_DEPENDENCIES** = [**PythonPackageSystemwideInstall**<**Virtualenv**>, **PythonPackageOnPip**<**Progressbar** (pro

Core dependencies required by the basic Pimlico installation, regardless of what pipeline is being processed. These will be checked when Pimlico is run, using the same dependency-checking mechanism that Pimlico modules use, and installed automatically if they're not found.

**pimlico.core.dependencies.java module****class JavaDependency** (*name, classes=[], jars=[], class\_dirs=[], \*\*kwargs*)

Bases: *pimlico.core.dependencies.base.SoftwareDependency*

Base class for Java library dependencies.

In addition to the usual functionality provided by dependencies, subclasses of this provide contributions to the Java classpath in the form of directories of jar files.

The instance has a set of representative Java classes that the checker will try to load to check whether the library is available and functional. It will also check that all jar files exist.

Jar paths and class directory paths are assumed to be relative to the Java lib dir (lib/java), unless they are absolute paths.

Subclasses should provide `install()` and override `installable()` if it's possible to install them automatically.

**problems** (*local\_config*)

**installable** ()

**get\_classpath\_components** ()

**class JavaJarsDependency** (*name, jar\_urls, \*\*kwargs*)

Bases: *pimlico.core.dependencies.java.JavaDependency*

Simple way to define a Java dependency where the library is packaged up in a jar, or a series of jars. The jars should be given as a list of (name, url) pairs, where name is the filename the jar should have and url is a url from which it can be downloaded.

URLs may also be given in the form "url->member", where url is a URL to a tar.gz or zip archive and member is a member to extract from the archive. If the type of the file isn't clear from the URL (i.e. if it doesn't have ".zip" or ".tar.gz" in it), specify the intended extension in the form "[ext]url->member", where ext is "tar.gz" or "zip".

**installable** ()

**install** (*local\_config, trust\_downloaded\_archives=False*)

**class PimlicoJavaLibrary** (*name, classes=[], additional\_jars=[]*)

Bases: *pimlico.core.dependencies.java.JavaDependency*

Special type of Java dependency for the Java libraries provided with Pimlico. These are packages up in jars and stored in the build dir.

**check\_java\_dependency** (*class\_name*, *classpath=None*)

Utility to check that a java class is able to be loaded.

**check\_java** ()

Check that the JVM executable can be found. Raises a `DependencyError` if it can't be found or can't be run.

**get\_classpath** (*deps*, *as\_list=False*)

Given a list of `JavaDependency` subclass instances, returns all the components of the classpath that will make sure that the dependencies are available.

If *as\_list=True*, returned as a list. Get the full classpath by `"".join(x)` on the list. If *as\_list=False*, returns classpath string.

**get\_module\_classpath** (*module*)

Builds a classpath that includes all of the classpath elements specified by Java dependencies of the given module. These include the dependencies from `get_software_dependencies()` and also any dependencies of the datatype.

Used to ensure that Java modules that depend on particular jars or classes get all of those files included on their classpath when Java is run.

**class Py4JSoftwareDependency**

Bases: `pimlico.core.dependencies.java.JavaDependency`

Java component of Py4J. Use this one as the main dependency, as it depends on the Python component and will install that first if necessary.

**dependencies** ()

**jars**

**installable** ()

**install** (*local\_config*, *trust\_downloaded\_archives=False*)

## pimlico.core.dependencies.python module

Tools for Python library dependencies.

Provides superclasses for Python library dependencies and a selection of commonly used dependency instances.

**class PythonPackageDependency** (*package*, *name*, *\*\*kwargs*)

Bases: `pimlico.core.dependencies.base.SoftwareDependency`

Base class for Python dependencies. Provides import checks, but no installation routines. Subclasses should either provide `install()` or `installation_instructions()`.

The import checks do not (as of 0.6rc) actually import the package, as this may have side-effects that are difficult to account for, causing odd things to happen when you check multiple times, or try to import later. Instead, it just checks whether the package finder is about to locate the package. This doesn't guarantee that the import will succeed.

**problems** (*local\_config*)

**import\_package** ()

Try importing `package_name`. By default, just uses `__import__`. Allows subclasses to allow for special import behaviour.

Should raise an `ImportError` if import fails.

**get\_installed\_version** (*local\_config*)

Tries to import a `__version__` variable from the package, which is a standard way to define the package version.

```
class PythonPackageSystemwideInstall (package_name, name, pip_package=None,
                                     apt_package=None, yum_package=None, **kwargs)
    Bases: pimlico.core.dependencies.python.PythonPackageDependency
    Dependency on a Python package that needs to be installed system-wide.

    installable ()

    installation_instructions ()

class PythonPackageOnPip (package, name=None, pip_package=None, **kwargs)
    Bases: pimlico.core.dependencies.python.PythonPackageDependency
    Python package that can be installed via pip. Will be installed in the virtualenv if not available.

    installable ()

    install (local_config, trust_downloaded_archives=False)

    get_installed_version (local_config)
```

## pimlico.core.dependencies.versions module

```
class SoftwareVersion (string_id)
    Bases: object

    Base class for representing version numbers / IDs of software. Different software may use different conventions
    to represent its versions, so it may be necessary to subclass this class to provide the appropriate parsing and
    comparison of versions.

compare_dotted_versions (version0, version1)
    Comparison function for reasonably standard version numbers, with subversions to any level of nesting specified
    by dots.
```

## Module contents

### pimlico.core.external package

#### Submodules

### pimlico.core.external.java module

```
call_java (class_name, args=[], classpath=None)

start_java_process (class_name, args=[], java_args=[], wait=0.1, classpath=None)

class Py4JInterface (gateway_class, port=None, python_port=None, gateway_args=[], pipeline=None,
                    print_stdout=True, print_stderr=True, env={}, system_properties={}, java_opts=[],
                    timeout=10.0, prefix_classpath=None)
    Bases: object

    start (timeout=None, port_output_prefix=None)
        Start a Py4J gateway server in the background on the given port, which will then be used for communicat-
        ing with the Java app.

        If a port has been given, it is assumed that the gateway accepts a -port option. Likewise with python_port
        and a -python-port option.

        If timeout is given, it overrides any timeout given in the constructor or specified in local config.
```

**new\_client()**

**stop()**

**clear\_output\_queues()**

**no\_retry\_gateway** (*\*\*kwargs*)

A wrapper around the constructor of JavaGateway that produces a version of it that doesn't retry on errors. The default gateway keeps retrying and outputting millions of errors if the server goes down, which makes responding to interrupts horrible (as the server might die before the Python process gets the interrupt).

TODO This isn't working: it just gets worse when I use my version!

**gateway\_client\_to\_running\_server** (*port*)

**launch\_gateway** (*gateway\_class='py4j.GatewayServer', args=[], javaopts=[], redirect\_stdout=None, redirect\_stderr=None, daemonize\_redirect=True, env={}, port\_output\_prefix=None, startup\_timeout=10.0, prefix\_classpath=None*)

Our own more flexible version of Py4J's launch\_gateway.

**get\_redirect\_func** (*redirect*)

**class OutputConsumer** (*redirects, stream, \*args, \*\*kwargs*)

Bases: `threading.Thread`

Thread that consumes output Modification of Py4J's OutputConsumer to allow multiple redirects.

**remove\_temporary\_redirects()**

**run()**

**output\_p4j\_error\_info** (*command, returncode, stdout, stderr*)

**make\_py4j\_errors\_safe** (*fn*)

Decorator for functions/methods that call Py4J. Py4J's exceptions include information that gets retrieved from the Py4J server when they're displayed. This is a problem if the server is not longer running and raises another exception, making the whole situation very confusing.

If you wrap your function with this, Py4JJavaErrors will be replaced by our own exception type Py4JSafeJavaError, containing some of the information about the Java exception if possible.

**exception Py4JSafeJavaError** (*java\_exception=None, str=None*)

Bases: `exceptions.Exception`

**exception DependencyCheckerError**

Bases: `exceptions.Exception`

**exception JavaProcessError**

Bases: `exceptions.Exception`

## Module contents

Tools for calling external (non-Python) tools.

## pimlico.core.modules package

### Subpackages

## pimlico.core.modules.map package

### Submodules

## pimlico.core.modules.map.filter module

**class DocumentMapOutputTypeWrapper** (\*args, \*\*kwargs)

Bases: object

**non\_filter\_datatype** = None

**wrapped\_module\_info** = None

**output\_name** = None

**archive\_iter** (subsample=None, start\_after=None)

Provides an iterator just like TarredCorpus, but instead of iterating over data read from disk, gets it on the fly from the input datatype.

**data\_ready** ()

Ready to supply this data as soon as all the wrapper module's inputs are ready to produce their data.

**wrap\_module\_info\_as\_filter** (module\_info\_instance)

Create a filter module from a document map module so that it gets executed on the fly to provide its outputs as input to later modules. Can be applied to any document map module simply by adding *filter=T* to its config.

This function is called when *filter=T* is given.

**Parameters** *module\_info\_instance* – basic module info to wrap the outputs of

**Returns** a new non-executable ModuleInfo whose outputs are produced on the fly and will be identical to the outputs of the wrapper module.

## pimlico.core.modules.map.multiproc module

Document map modules can in general be easily parallelized using multiprocessing. This module provides implementations of a pool and base worker processes that use multiprocessing, making it dead easy to implement a parallelized module, simply by defining what should be done on each document.

In particular, use :fun:multiprocessing\_executor\_factory wherever possible.

**class MultiprocessingMapProcess** (input\_queue, output\_queue, exception\_queue, executor, docs\_per\_batch=1)

Bases: multiprocessing.process.Process, pimlico.core.modules.map.DocumentMapProcessMixin

A base implementation of document map parallelization using multiprocessing. Note that not all document map modules will want to use this: e.g. if you call a background service that provides parallelization itself (like the CoreNLP module) there's no need for multiprocessing in the Python code.

**notify\_no\_more\_inputs** ()

**run** ()

**class MultiprocessingMapPool** (*executor, processes*)

Bases: *pimlico.core.modules.map.DocumentProcessorPool*

A base implementation of document map parallelization using multiprocessing.

**PROCESS\_TYPE** = None

**SINGLE\_PROCESS\_TYPE** = None

**start\_worker** ()

**static create\_queue** (*maxsize=None*)

**shutdown** ()

**notify\_no\_more\_inputs** ()

**empty\_all\_queues** ()

**class MultiprocessingMapModuleExecutor** (*module\_instance\_info, \*\*kwargs*)

Bases: *pimlico.core.modules.map.DocumentMapModuleExecutor*

**POOL\_TYPE** = None

**create\_pool** (*processes*)

**postprocess** (*error=False*)

**multiprocessing\_executor\_factory** (*process\_document\_fn, preprocess\_fn=None, post-  
process\_fn=None, worker\_set\_up\_fn=None,  
worker\_tear\_down\_fn=None, batch\_docs=None, multi-  
processing\_single\_process=False*)

Factory function for creating an executor that uses the multiprocessing-based implementations of document-map pools and worker processes. This is an easy way to implement a parallelizable executor, which is suitable for a large number of module types.

*process\_document\_fn* should be a function that takes the following arguments (unless *batch\_docs* is given):

- the worker process instance (allowing access to things set during setup)
- archive name
- document name
- the rest of the args are the document itself, from each of the input corpora

If *preprocess\_fn* is given, it is called from the main process once before execution begins, with the executor as an argument.

If *postprocess\_fn* is given, it is called from the main process at the end of execution, including on the way out after an error, with the executor as an argument and a kwarg *error* which is True if execution failed.

If *worker\_set\_up\_fn* is given, it is called within each worker before execution begins, with the worker process instance as an argument. Likewise, *worker\_tear\_down\_fn* is called from within the worker process before it exits.

Alternatively, you can supply a worker type, a subclass of :class:MultiprocessingMapProcess, as the first argument. If you do this, *worker\_set\_up\_fn* and *worker\_tear\_down\_fn* will be ignored.

If *batch\_docs* is not None, *process\_document\_fn* is treated differently. Instead of supplying the *process\_document()* of the worker, it supplies a *process\_documents()*. The second argument is a list of tuples, each of which is assumed to be the args to *process\_document()* for a single document. In this case, *docs\_per\_batch* is set on the worker processes, so that the given number of docs are collected from the input and passed into *process\_documents()* at once.

By default, if only a single process is needed, we use the threaded implementation of a map process instead of multiprocessing. If this doesn't work out in your case, for some reason, specify *multiprocessing\_single\_process=True* and a multiprocessing process will be used even when only creating one.

### pimlico.core.modules.map.singleproc module

Sometimes the simple multiprocessing-based approach to map module parallelization just isn't suitable. This module provides an equivalent set of implementations and convenience functions that don't use multiprocessing, but conform to the pool-based execution pattern by creating a single-thread pool.

**class SingleThreadMapModuleExecutor** (*module\_instance\_info*, *\*\*kwargs*)

Bases: *pimlico.core.modules.map.threaded.ThreadingMapModuleExecutor*

**create\_pool** (*processes*)

**single\_process\_executor\_factory** (*process\_document\_fn*, *preprocess\_fn=None*, *post-*  
*process\_fn=None*, *worker\_set\_up\_fn=None*,  
*worker\_tear\_down\_fn=None*, *batch\_docs=None*)

Factory function for creating an executor that uses the single-process implementations of document-map pools and workers. This is an easy way to implement a non-parallelized executor

*process\_document\_fn* should be a function that takes the following arguments:

- the executor instance (allowing access to things set during setup)
- archive name
- document name
- the rest of the args are the document itself, from each of the input corpora

If *preprocess\_fn* is given, it is called once before execution begins, with the executor as an argument.

If *postprocess\_fn* is given, it is called at the end of execution, including on the way out after an error, with the executor as an argument and a kwarg *error* which is True if execution failed.

### pimlico.core.modules.map.threaded module

Just like multiprocessing, but using threading instead. If you're not sure which you should use, it's probably multiprocessing.

**class ThreadingMapThread** (*input\_queue*, *output\_queue*, *exception\_queue*, *executor*)

Bases: *threading.Thread*, *pimlico.core.modules.map.DocumentMapProcessMixin*

**notify\_no\_more\_inputs** ()

**run** ()

**shutdown** (*timeout=3.0*)

**class ThreadingMapPool** (*executor*, *processes*)

Bases: *pimlico.core.modules.map.DocumentProcessorPool*

**THREAD\_TYPE** = None

**start\_worker** ()

**static create\_queue** (*maxsize=None*)

**shutdown** ()

```
class ThreadingMapModuleExecutor (module_instance_info, **kwargs)
    Bases: pimlico.core.modules.map.DocumentMapModuleExecutor
```

```
    POOL_TYPE = None
```

```
    create_pool (processes)
```

```
    postprocess (error=False)
```

```
threading_executor_factory (process_document_fn, preprocess_fn=None, postprocess_fn=None,
                             worker_set_up_fn=None, worker_tear_down_fn=None)
```

Factory function for creating an executor that uses the threading-based implementations of document-map pools and worker processes.

process\_document\_fn should be a function that takes the following arguments:

- the worker process instance (allowing access to things set during setup)
- archive name
- document name
- the rest of the args are the document itself, from each of the input corpora

If preprocess\_fn is given, it is called from the main thread once before execution begins, with the executor as an argument.

If postprocess\_fn is given, it is called from the main thread at the end of execution, including on the way out after an error, with the executor as an argument and a kwarg *error* which is True if execution failed.

If worker\_set\_up\_fn is given, it is called within each worker before execution begins, with the worker thread instance as an argument. Likewise, worker\_tear\_down\_fn is called from within the worker thread before it exits.

Alternatively, you can supply a worker type, a subclass of :class:ThreadingMapThread, as the first argument. If you do this, worker\_set\_up\_fn and worker\_tear\_down\_fn will be ignored.

## Module contents

```
class DocumentMapModuleInfo (module_name, pipeline, **kwargs)
    Bases: pimlico.core.modules.base.BaseModuleInfo
```

Abstract module type that maps each document in turn in a corpus. It produces a single output document for every input.

Subclasses should specify the input types, which should all be subclasses of TarredCorpus, and output types, the first of which (i.e. default) should also be a subclass of TarredCorpus. The base class deals with iterating over the input(s) and writing the outputs to a new TarredCorpus. The subclass only needs to implement the mapping function applied to each document (in its executor).

```
    module_outputs = [('documents', <class 'pimlico.datatypes.tar.TarredCorpus'>)]
```

```
    input_corpora
```

```
    get_writer (output_name, output_dir, append=False)
```

Get the writer instance that will be given processed documents to write. Should return a subclass of TarredCorpusWriter. The default implementation instantiates a plain TarredCorpusWriter.

```
    get_writers (append=False)
```

```
    get_detailed_status ()
```

```
class DocumentMapModuleExecutor (module_instance_info, **kwargs)
    Bases: pimlico.core.modules.base.BaseModuleExecutor
```



Base class for executors for document map modules. Subclasses should provide the behaviour for each individual document by defining a pool (and worker processes) to handle the documents as they're fed into it.

Note that in most cases it won't be necessary to override the pool and worker base classes yourself. Unless you need special behaviour, use the standard implementations and factory functions.

Although the pattern of execution for all document map modules is based on parallel processing (creating a pool, spawning worker processes, etc), this doesn't mean that all such modules have to be parallelizable. If you have no reason not to parallelize, it's recommended that you do (with single-process execution as a special case). However, sometimes parallelizing isn't so simple: in these cases, consider using the tools in :mod:..singleproc.

**preprocess** ()

Allows subclasses to define a set-up procedure to be called before corpus processing begins.

**postprocess** (*error=False*)

Allows subclasses to define a finishing procedure to be called after corpus processing if finished.

**create\_pool** (*processes*)

Should return an instance of the pool to be used for document processing. Should generally be a subclass of DocumentProcessorPool.

Always called after preprocess().

**retrieve\_processing\_status** ()

**update\_processing\_status** (*docs\_completed, archive\_name, filename*)

**execute** ()

**skip\_invalid** (*fn*)

Decorator to apply to document map executor process\_document() methods where you want to skip doing any processing if any of the input documents are invalid and just pass through the error information.

Be careful not to confuse this with the process\_document() methods on datatypes. You don't need a decorator on them to skip invalid documents, as it's not called on them anyway.

**skip\_invalids** (*fn*)

Decorator to apply to document map executor process\_documents() methods where you want to skip doing any processing if any of the input documents are invalid and just pass through the error information.

**invalid\_doc\_on\_error** (*fn*)

Decorator to apply to process\_document() methods that causes all exceptions to be caught and an InvalidDocument to be returned as the result, instead of letting the error propagate up and call a halt to the whole corpus processing.

**invalid\_docs\_on\_error** (*fn*)

Decorator to apply to process\_documents() methods that causes all exceptions to be caught and an InvalidDocument to be returned as the result for every input document.

**class ProcessOutput** (*archive, filename, data*)

Bases: object

Wrapper for all result data coming out from a worker.

**class InputQueueFeeder** (*input\_queue, iterator, complete\_callback=None*)

Bases: threading.Thread

Background thread to read input documents from an iterator and feed them onto an input queue for worker processes/threads.

**get\_next\_output\_document** ()

**check\_invalid** (*archive, filename*)

Checks whether a given document was invalid in the input. Once the check has been performed, the item is removed from the list, for efficiency, so this should only be called once per document.

**run** ()

**check\_for\_error** ()

Can be called from the main thread to check whether an error has occurred in this thread and raise a suitable exception if so

**shutdown** (*timeout=3.0*)

Cancel the feeder, if it's still feeding and stop the thread. Call only after you're sure you no longer need anything from any of the queues. Waits for the thread to end.

Call from the main thread (that created the feeder) only.

**class DocumentProcessorPool** (*processes*)

Bases: `object`

Base class for pools that provide an easy implementation of parallelization for document map modules. Defines the core interface for pools.

If you're using multiprocessing, you'll want to use the multiprocessing-specific subclass.

**notify\_no\_more\_inputs** ()

**static create\_queue** (*maxsize=None*)

May be overridden by subclasses to provide different implementations of a Queue. By default, uses the multiprocessing queue type. Whatever is returned, it should implement the interface of `Queue.Queue`.

**shutdown** ()

**empty\_all\_queues** ()

**class DocumentMapProcessMixin** (*input\_queue, output\_queue, exception\_queue, docs\_per\_batch=1*)

Bases: `object`

Mixin/base class that should be implemented by all worker processes for document map pools.

**set\_up** ()

Called when the process starts, before it starts accepting documents.

**process\_document** (*archive, filename, \*docs*)

**process\_documents** (*doc\_tuples*)

Batched version of `process_document()`. Default implementation just calls `process_document()` on each document, but if you want to group documents together and process multiple at once, you can override this method and make sure the *docs\_per\_batch* is set > 1.

Each item in the list of doc tuples should be a tuple of the positional args to `process_document()` – i.e. `archive_name, filename, doc_from_corpus1, [doc_from_corpus2, ...]`

**tear\_down** ()

Called from within the process after processing is complete, before exiting.

**notify\_no\_more\_inputs** ()

Called when there aren't any more inputs to come.

**exception WorkerStartupError** (*\*args, \*\*kwargs*)

Bases: `exceptions.Exception`

**exception WorkerShutdownError** (*\*args, \*\*kwargs*)

Bases: `exceptions.Exception`

## Submodules

### pimlico.core.modules.base module

This module provides base classes for Pimlico modules.

The procedure for creating a new module is the same whether you're contributing a module to the core set in the Pimlico codebase or a standalone module in your own codebase, or for a specific pipeline.

A Pimlico module is identified by the full Python-path to the Python package that contains it. This package should be laid out as follows:

- The module's metadata is defined by a class in `info.py` called `ModuleInfo`, which should inherit from `BaseModuleInfo` or one of its subclasses.
- The module's functionality is provided by a class in `execute.py` called `ModuleExecutor`, which should inherit from `BaseModuleExecutor`.

The exec Python module will not be imported until an instance of the module is to be run. This means that you can import dependencies and do any necessary initialization at the point where it's executed, without worrying about incurring the associated costs (and dependencies) every time a pipeline using the module is loaded.

```
class BaseModuleInfo (module_name, pipeline, inputs={}, options={}, optional_outputs=[], doc-
                        string='', include_outputs=[], alt_expanded_from=None, alt_param_settings=[],
                        module_variables={})
```

Bases: `object`

Abstract base class for all pipeline modules' metadata.

**module\_type\_name** = `None`

**module\_readable\_name** = `None`

**module\_options** = {}

**module\_inputs** = []

Specifies a list of (name, datatype class) pairs for outputs that are always written

**module\_optional\_outputs** = []

Whether the module should be executed Typically True for almost all modules, except input modules (though some of them may also require execution) and filters

**module\_executable** = `True`

If specified, this `ModuleExecutor` class will be used instead of looking one up in the exec Python module

**module\_executor\_override** = `None`

Usually `None`. In the case of stages of a multi-stage module, stores a pointer to the main module.

**main\_module** = `None`

**module\_outputs** = []

Specifies a list of (name, datatype class) pairs for outputs that are written only if they're specified in the "output" option or used by another module

**load\_executor** ()

Loads a `ModuleExecutor` for this Pimlico module. Usually, this just involves calling `load_module_executor()`, but the default executor loading may be overridden for a particular module type by overriding this function. It should always return a subclass of `ModuleExecutor`, unless there's an error.

**classmethod get\_key\_info\_table** ()

When generating module docs, the table at the top of the page is produced by calling this method. It should

return a list of two-item lists (title + value). Make sure to include the super-class call if you override this to add in extra module-specific info.

**metadata\_filename**

**get\_metadata()**

**set\_metadata\_value** (*attr*, *val*)

**set\_metadata\_values** (*val\_dict*)

**status**

**execution\_history\_path**

**add\_execution\_history\_record** (*line*)

Output a single line to the file that stores the history of module execution, so we can trace what we've done.

**execution\_history**

Get the entire recorded execution history for this module. Returns an empty string if no history has been recorded.

**input\_names**

**output\_names**

**classmethod process\_module\_options** (*opt\_dict*)

Parse the options in a dictionary (probably from a config file), checking that they're valid for this model type.

**Parameters** **opt\_dict** – dict of options, keyed by option name

**Returns** dict of options

**classmethod extract\_input\_options** (*opt\_dict*, *module\_name=None*, *previous\_module\_name=None*, *module\_expansions={}*)

Given the config options for a module instance, pull out the ones that specify where the inputs come from and match them up with the appropriate input names.

The inputs returned are just names as they come from the config file. They are split into module name and output name, but they are not in any way matched up with the modules they connect to or type checked.

**Parameters**

- **module\_name** – name of the module being processed, for error output. If not given, the name isn't included in the error.
- **previous\_module\_name** – name of the previous module in the order given in the config file, allowing a single-input module to default to connecting to this if the input connection wasn't given
- **module\_expansions** – dictionary mapping module names to a list of expanded module names, where expansion has been performed as a result of alternatives in the parameters. Provided here so that the unexpanded names may be used to refer to the whole list of module names, where a module takes multiple inputs on one input parameter

**Returns** dictionary of inputs

**static get\_extra\_outputs\_from\_options** (*options*)

Normally, which optional outputs get produced by a module depend on the 'output' option given in the config file, plus any outputs that get used by subsequent modules. By overriding this method, module types can add extra outputs into the list of those to be included, conditional on other options.

It also receives the processed dictionary of inputs, so that the additional outputs can depend on what is fed into the input.

E.g. the `corenlp` module include the ‘annotations’ output if annotators are specified, so that the user doesn’t need to give both options.

**provide\_further\_outputs** ()

Called during instantiation, once inputs and options are available, to add a further list of module outputs that are dependent on inputs or options.

**get\_module\_output\_dir** (*short\_term\_store=False*)

Gets the path to the base output dir to be used by this module, relative to the storage base dir. When outputting data, the storage base dir will always be the short term store path, but when looking for the output data other base paths might be explored, including the long term store.

**Parameters** **short\_term\_store** – if True, return absolute path to output dir in short-term store (used for output)

**Returns** path, relative to store base path, or if `short_term_store=True` absolute path to output dir

**get\_absolute\_output\_dir** (*output\_name*)

The simplest way to get hold of the directory to use to output data to for a given output. This is the usual way to get an output directory for an output writer.

The directory is an absolute path to a location in the Pimlico short-term store.

**Parameters** **output\_name** – the name of an output

**Returns** the absolute path to the output directory to use for the named output

**get\_output\_dir** (*output\_name, short\_term\_store=False*)

**Parameters**

- **short\_term\_store** – return an absolute path in the short-term store. If False (default), return a relative path, specified relative to the root of the Pimlico store used. This allows multiple stores to be searched for output
- **output\_name** – the name of an output

**Returns** the path to the output directory to use for the named output, which may be relative to the root of the Pimlico store in use (default) or an absolute path in the short-term store, depending on *short\_term\_store*

**get\_output\_datatype** (*output\_name=None, additional\_names=[]*)

**instantiate\_output\_datatype** (*output\_name, output\_datatype, \*\*kwargs*)

Subclasses may want to override this to provide special behaviour for instantiating particular outputs’ datatypes.

Additional kwargs will be pass through to the datatype’s init.

**get\_output** (*output\_name=None, additional\_names=None, \*\*kwargs*)

Get a datatype instance corresponding to one of the outputs of the module.

Additional kwargs will be pass through to the datatype’s init.

**is\_multiple\_input** (*input\_name=None*)

Returns True if the named input (or default input if no name is given) is a MultipleInputs input, False otherwise. If it is, `get_input()` will return a list, otherwise it will return a single datatype.

**get\_input\_module\_connection** (*input\_name=None, always\_list=False*)

Get the ModuleInfo instance and output name for the output that connects up with a named input (or the

first input) on this module instance. Used by `get_input()` – most of the time you probably want to use that to get the instantiated datatype for an input.

If the input type was specified with `MultipleInputs`, meaning that we’re expecting an unbounded number of inputs, this is a list. Otherwise, it’s a single (module, output\_name) pair. If `always_list=True`, in this latter case we return a single-item list.

**get\_input\_datatype** (*input\_name=None, always\_list=False*)

Get a list of datatype classes corresponding to one of the inputs to the module. If an input name is not given, the first input is returned.

If the input type was specified with `MultipleInputs`, meaning that we’re expecting an unbounded number of inputs, this is a list. Otherwise, it’s a single datatype.

**get\_input** (*input\_name=None, always\_list=False, \*\*kwargs*)

Get a datatype instances corresponding to one of the inputs to the module. Looks up the corresponding output from another module and uses that module’s metadata to get that output’s instance. If an input name is not given, the first input is returned.

If the input type was specified with `MultipleInputs`, meaning that we’re expecting an unbounded number of inputs, this is a list. Otherwise, it’s a single datatype instance. If `always_list=True`, in this latter case we return a single-item list.

Additional kwargs will be passed through to the datatype’s init call.

**input\_ready** (*input\_name=None*)

Check whether the datatype is (or datatypes are) ready to go, corresponding to the named input.

**Parameters** `input_name` – input to check

**Returns** True if input is ready

**all\_inputs\_ready** ()

Check `input_ready()` on all inputs.

**Returns** True if all input datatypes are ready to be used

**classmethod is\_filter** ()

**missing\_module\_data** ()

Reports missing data not associated with an input dataset.

Calling `missing_data()` reports any problems with input data associated with a particular input to this module. However, modules may also rely on data that does not come from one of their inputs. This happens primarily (perhaps solely) when a module option points to a data source. This might be the case with any module, but is particularly common among input reader modules, which have no inputs, but read data according to their options.

**Returns** list of problems

**missing\_data** (*input\_names=None, assume\_executed=[], assume\_failed=[], al-  
low\_preliminary=False*)

Check whether all the input data for this module is available. If not, return a list strings indicating which outputs of which modules are not available. If it’s all ready, returns an empty list.

To check specific inputs, give a list of input names. To check all inputs, don’t specify `input_names`. To check the default input, give `input_names=[None]`. If not checking a specific input, also checks non-input data (see `missing_module_data()`).

If `assume_executed` is given, it should be a list of module names which may be assumed to have been executed at the point when this module is executed. Any outputs from those modules will be excluded from the input checks for this module, on the assumption that they will have become available, even if they’re not currently available, by the time they’re needed.

If *assume\_executed* is given, it should be a list of module names which should be assumed to have failed. If we rely on data from the output of one of them, instead of checking whether it's available we simply assume it's not.

Why do this? When running multiple modules in sequence, if one fails it is possible that its output datasets look like complete datasets. For example, a partially written iterable corpus may look like a perfectly valid corpus, which happens to be smaller than it should be. After the execution failure, we may check other modules to see whether it's possible to run them. Then we need to know not to trust the output data from the failed module, even if it looks valid.

If *allow\_preliminary=True*, for any inputs that are multiple inputs and have multiple connections to previous modules, consider them to be satisfied if at least one of their inputs is ready. The normal behaviour is to require all of them to be ready, but in a preliminary run this requirement is relaxed.

**classmethod** `is_input()`

**dependencies**

**Returns** list of names of modules that this one depends on for its inputs.

**get\_transitive\_dependencies()**

Transitive closure of *dependencies*.

**Returns** list of names of modules that this one recursively (transitively) depends on for its inputs.

**typecheck\_inputs()**

**get\_software\_dependencies()**

Check that all software required to execute this module is installed and locatable. This is separate to metadata config checks, so that you don't need to satisfy the dependencies for all modules in order to be able to run one of them. You might, for example, want to run different modules on different machines. This is called when a module is about to be executed and each of the dependencies is checked.

Returns a list of instances of subclasses of `:class:~pimlico.core.dependencies.base.SoftwareDependency`, representing the libraries that this module depends on.

Take care when providing dependency classes that you don't put any import statements at the top of the Python module that will make loading the dependency type itself dependent on runtime dependencies. You'll want to run import checks by putting import statements within this method.

You should call the super method for checking superclass dependencies.

**get\_input\_software\_dependencies()**

Collects library dependencies from the input datatypes to this module, which will need to be satisfied for the module to be run.

Unlike *get\_software\_dependencies()*, it shouldn't need to be overridden by subclasses, since it just collects the results of getting dependencies from the datatypes.

**get\_output\_software\_dependencies()**

Collects library dependencies from the output datatypes to this module, which will need to be satisfied for the module to be run.

Unlike *get\_input\_software\_dependencies()*, it may not be the case that all of these dependencies strictly need to be satisfied before the module can be run. It could be that a datatype can be written without satisfying all the dependencies needed to read it. However, we assume that dependencies of all output datatypes must be satisfied in order to run the module that writes them, since this is usually the case, and these are checked before running the module.

Unlike *get\_software\_dependencies()*, it shouldn't need to be overridden by subclasses, since it just collects the results of getting dependencies from the datatypes.

**check\_ready\_to\_run()**

Called before a module is run, or if the ‘check’ command is called. This will only be called after all library dependencies have been confirmed ready (see :method:get\_software\_dependencies).

Essentially, this covers any module-specific checks that used to be in check\_runtime\_dependencies() other than library installation (e.g. checking models exist).

Always call the super class’ method if you override.

Returns a list of (name, description) pairs, where the name identifies the problem briefly and the description explains what’s missing and (ideally) how to fix it.

**reset\_execution()**

Remove all output data and metadata from this module to make a fresh start, as if it’s never been executed.

May be overridden if a module has some side effect other than creating/modifying things in its output directory(/ies), but overridden methods should always call the super method. Occasionally this is necessary, but most of the time the base implementation is enough.

**get\_detailed\_status()**

Returns a list of strings, containing detailed information about the module’s status that is specific to the module type. This may include module-specific information about execution status, for example.

Subclasses may override this to supply useful (human-readable) information specific to the module type. They should call the super method.

**classmethod module\_package\_name()**

The package name for the module, which is used to identify it in config files. This is the package containing the info.py in which the ModuleInfo is defined.

**get\_execution\_dependency\_tree()**

Tree of modules that will be executed when this one is executed. Where this module depends on filters, the tree goes back through them to find what they depend on (since they will be executed simultaneously)

**get\_all\_executed\_modules()**

Returns a list of all the modules that will be executed when this one is (including itself). This is the current module (if executable), plus any filters used to produce its inputs.

**lock\_path****lock()**

Mark the module as locked, so that it cannot be executed. Called when execution begins, to ensure that you don’t end up executing the same module twice simultaneously.

**unlock()**

Remove the execution lock on this module.

**is\_locked()**

**Returns** True if the module is currently locked from execution

**get\_new\_log\_filename(name='error')**

Returns an absolute path that can be used to output a log file for this module. This is used for outputting error logs. It will always return a filename that doesn’t currently exist, so can be used multiple times to output multiple logs.

**collect\_unexecuted\_dependencies(modules)**

Given a list of modules, checks through all the modules that they depend on to put together a list of modules that need to be executed so that the given list will be left in an executed state. The list includes the modules themselves, if they’re not fully executed, and unexecuted dependencies of any unexecuted modules (recursively).

**Parameters** **modules** – list of ModuleInfo instances



**Returns** list of ModuleInfo instances that need to be executed

**collect\_runnable\_modules** (*pipeline, preliminary=False*)

Look for all unexecuted modules in the pipeline to find any that are ready to be executed. Keep collecting runnable modules, including those that will become runnable once we’ve run earlier ones in the list, to produce a list of a sequence of modules that could be set running now.

**Parameters** **pipeline** – pipeline config

**Returns** ordered list of runnable modules. Note that it must be run in this order, as some might depend on earlier ones in the list

**satisfies\_typecheck** (*provided\_type, type\_requirements*)

Interface to Pimlico’s standard type checking (see *check\_type*) that returns a boolean to say whether type checking succeeded or not.

**check\_type** (*provided\_type, type\_requirements*)

Type-checking algorithm for making sure outputs from modules connect up with inputs that they satisfy the requirements for.

**class BaseModuleExecutor** (*module\_instance\_info, stage=None, debug=False, force\_rerun=False*)

Bases: object

Abstract base class for executors for Pimlico modules. These are classes that actually do the work of executing the module on given inputs, writing to given output locations.

**execute** ()

Run the actual module execution.

May return None, in which case it’s assumed to have fully completed. If a string is returned, it’s used as an alternative module execution status. Used, e.g., by multi-stage modules that need to be run multiple times.

**exception ModuleInfoLoadError** (*\*args, \*\*kwargs*)

Bases: exceptions.Exception

**exception ModuleExecutorLoadError**

Bases: exceptions.Exception

**exception ModuleTypeError**

Bases: exceptions.Exception

**exception TypeCheckError**

Bases: exceptions.Exception

**exception DependencyError** (*message, stderr=None, stdout=None*)

Bases: exceptions.Exception

Raised when a module’s dependencies are not satisfied. Generally, this means a dependency library needs to be installed, either on the local system or (more often) by calling the appropriate make target in the lib directory.

**load\_module\_executor** (*path\_or\_info*)

Utility for loading the executor class for a module from its full path. More or less just a wrapper around an import, with some error checking. Locates the executor by a standard procedure that involves checking for an “execute” python module alongside the info’s module.

Note that you shouldn’t generally use this directly, but instead call the *load\_executor()* method on a module info (which will call this, unless special behaviour has been defined).

**Parameters** **path** – path to Python package containing the module

**Returns** class

**load\_module\_info** (*path*)

Utility to load the metadata for a Pimlico pipeline module from its package Python path.

**Parameters** *path* –

**Returns**

## pimlico.core.modules.execute module

Runtime execution of modules

This module provides the functionality to check that Pimlico modules are ready to execute and execute them. It is used by the *run* command.

**check\_and\_execute\_modules** (*pipeline, module\_names, force\_rerun=False, debug=False, log=None, all\_deps=False, check\_only=False, exit\_on\_error=False, preliminary=False, email=None*)

Main method called by the *run* command that first checks a pipeline, checks all pre-execution requirements of the modules to be executed and then executes each of them. The most common case is to execute just one module, but a sequence may be given.

**Parameters**

- **exit\_on\_error** – drop out if a `ModuleExecutionError` occurs in any individual module, instead of continuing to the next module that can be run
- **pipeline** – loaded `PipelineConfig`
- **module\_names** – list of names of modules to execute in the order they should be run
- **force\_rerun** – execute modules, even if they’re already marked as complete
- **debug** – output debugging info
- **log** – logger, if you have one you want to reuse
- **all\_deps** – also include unexecuted dependencies of the given modules
- **check\_only** – run all checks, but stop before executing. Used for *check* command

**Returns**

**check\_modules\_ready** (*pipeline, modules, log, preliminary=False*)

Check that a module is ready to be executed. Always called before execution begins.

**Parameters**

- **pipeline** – loaded `PipelineConfig`
- **modules** – loaded `ModuleInfo` instances, given in the order they’re going to be executed. For each module, it’s assumed that those before it in the list have already been run when it is run.
- **log** – logger to output to

**Returns** If *preliminary=True*, list of problems that were ignored by allowing preliminary run. Otherwise, `None` – we raise an exception when we first encounter a problem

**execute\_modules** (*pipeline, modules, log, force\_rerun=False, debug=False, exit\_on\_error=False, preliminary=False, email=None*)

**format\_execution\_dependency\_tree** (*tree*)

**send\_final\_report\_email** (*pipeline, error\_modules, success\_modules, skipped\_modules, all\_modules*)

**send\_module\_report\_email** (*pipeline, module, short\_error, long\_error*)

**exception ModuleExecutionError** (\*args, \*\*kwargs)

Bases: `exceptions.Exception`

**exception ModuleNotReadyError** (\*args, \*\*kwargs)

Bases: `pimlico.core.modules.execute.ModuleExecutionError`

**exception ModuleAlreadyCompletedError** (\*args, \*\*kwargs)

Bases: `pimlico.core.modules.execute.ModuleExecutionError`

**exception StopProcessing**

Bases: `exceptions.Exception`

## pimlico.core.modules.inputs module

Base classes and utilities for input modules in a pipeline.

**class InputModuleInfo** (*module\_name*, *pipeline*, *inputs*={}, *options*={}, *optional\_outputs*=[], *docstring*=',', *include\_outputs*=[], *alt\_expanded\_from*=None, *alt\_param\_settings*=[], *module\_variables*={})

Bases: `pimlico.core.modules.base.BaseModuleInfo`

Base class for input modules. These don't get executed in general, they just provide a way to iterate over input data.

You probably don't want to subclass this. It's usually simplest to define a datatype for reading the input data and then just specify its class as the module's type. This results in a subclass of this module info being created dynamically to read that data.

Note that `module_executable` is typically set to `False` and the base class does this. However, some input modules need to be executed before the input is usable, for example to collect stats about the input data.

**module\_type\_name** = 'input'

**module\_executable** = `False`

**instantiate\_output\_datatype** (*output\_name*, *output\_datatype*, \*\*kwargs)

**input\_module\_factory** (*datatype*)

Create an input module class to load a given datatype.

**class ReaderOutputType** (*reader\_options*, *pipeline*, \*\*kwargs)

Bases: `pimlico.datatypes.base.IterableCorpus`

A datatype for reading in input according to input module options and allowing it to be iterated over by other modules.

Typically used together with `iterable_input_reader_factory()` as the output datatype.

`__len__` should be overridden to take the processed input module options and return the length of the corpus (number of documents).

`__iter__` should use the processed input module options and return an iterator over the corpus' documents (e.g. a generator function). Each item yielded should be a pair (`doc_name`, `data`) and `data` should be in the appropriate internal format associated with the document type.

`data_ready` should be overridden to use the processed input module options and return `True` if the data is ready to be read in.

In all cases, the input options are available as `self.reader_options`.

**datatype\_name** = 'reader\_iterator'

**data\_point\_type = None**

Must be overridden by subclasses

**emulated\_datatype**

Subclass information should be ignored for type checking. Should be treated exactly as an IterableCorpus alias of IterableCorpus

**data\_ready()**

**iterable\_input\_reader\_factory**(*input\_module\_options*, *output\_type*, *module\_type\_name=None*,  
*module\_readable\_name=None*)

Factory for creating an input reader module type. This is a non-executable module that has no inputs. It reads its data from some external location, using the given module options. The resulting dataset is an IterableCorpus subtype, with the given document type.

*output\_type* is a datatype that performs the actual iteration over the data and is instantiated with the processed options as its first argument. This is typically created by subclassing ReaderOutputType and providing *len*, *iter* and *data\_ready* methods.

**How is this different from “input\_module\_factory”?** This method is used in your module code to prepare a ModuleInfo class for reading a particular type of input data and presenting it as a Pimlico dataset of the given type. *input\_module\_factory*, on the other hand, is used by Pimlico when you specify a datatype as a module type in a config file.

Note that, in future versions, reading datasets output by another Pimlico pipeline will be the only purpose for that special notation. The possibility of specifying *input\_module\_options* to create an input reader will disappear, so the use of *input\_module\_options* should be phased out and replaced with input reader modules, such as those created by this factory.

## pimlico.core.modules.multistage module

**class MultistageModuleInfo**(*module\_name*, *pipeline*, *\*\*kwargs*)

Bases: *pimlico.core.modules.base.BaseModuleInfo*

Base class for multi-stage modules. You almost certainly don’t want to override this yourself, but use the factory method instead. It exists mainly for providing a way of identifying multi-stage modules.

**module\_executable = True**

**stages = None**

**typecheck\_inputs()**

Overridden to check internal output-input connections as well as the main module’s inputs.

**get\_software\_dependencies()**

**get\_input\_software\_dependencies()**

**check\_ready\_to\_run()**

**get\_detailed\_status()**

**reset\_execution()**

**classmethod get\_key\_info\_table()**

Add the stages into the key info table.

**get\_next\_stage()**

If there are more stages to be executed, returns a pair of the module info and stage definition. Otherwise, returns (None, None)

**status**

`is_locked()`

**multistage\_module**(*multistage\_module\_type\_name*, *module\_stages*, *use\_stage\_option\_names=False*,  
*module\_readable\_name=None*)

Factory to build a multi-stage module type out of a series of stages, each of which specifies a module type for the stage. The stages should be a list of *ModuleStage* objects.

**class ModuleStage**(*name*, *module\_info\_cls*, *connections=None*, *output\_connections=None*, *option\_connections=None*, *use\_stage\_option\_names=False*)

Bases: `object`

A single stage in a multi-stage module.

If no explicit input connections are given, the default input to this module is connected to the default output from the previous.

Connections can be given as a list of `ModuleConnections`.

Output connections specify that one of this module's outputs should be used as an output from the multi-stage module. Optional outputs for the multi-stage module are not currently supported (though could in theory be added later). This should be a list of `ModuleOutputConnections`. If none are given for any of the stages, the module will have a single output, which is the default output from the last stage.

Option connections allow you to specify the names that are used for the multistage module's options that get passed through to this stage's module options. Simply specify a dict for `option_connections` where the keys are names module options for this stage and the values are the names that should be used for the multistage module's options.

You may map multiple options from different stages to the same option name for the multistage module. This will result in the same option value being passed through to both stages. Note that help text, option type, option processing, etc will be taken from the first stage's option (in case the two options aren't identical).

Options not explicitly mapped to a name will use the name `<stage_name>_<option_name>`. If `use_stage_option_names=True`, this prefix will not be added: the stage's option names will be used directly as the option name of the multistage module. Note that there is a danger of clashing option names with this behaviour, so only do it if you know the stages have distinct option names (or should share their values where the names overlap).

**class ModuleConnection**

Bases: `object`

**class InternalModuleConnection**(*input\_name*, *output\_name=None*, *previous\_module=None*)

Bases: `pimlico.core.modules.multistage.ModuleConnection`

Connection between the output of one module in the multi-stage module and the input to another.

May specify the name of the previous module that a connection should be made to. If this is not given, the previous module in the sequence will be assumed.

If `output_name=None`, connects to the default output of the previous module.

**class ModuleInputConnection**(*stage\_input\_name=None*, *main\_input\_name=None*)

Bases: `pimlico.core.modules.multistage.ModuleConnection`

Connection of a sub-module's input to an input to the multi-stage module.

If `main_input_name` is not given, the name for the input to the multistage module will be identical to the stage input name. This might lead to unintended behaviour if multiple inputs end up with the same name, so you can specify a different name if necessary to avoid clashes.

If multiple inputs (e.g. from different stages) are connected to the same main input name, they will take input from the same previous module output. Nothing clever is done to unify the type requirements, however: the first stage's type requirement is used for the main module's input.

If *stage\_input\_name* is not given, the module's default input will be connected.

**class ModuleOutputConnection** (*stage\_output\_name=None, main\_output\_name=None*)

Bases: `object`

Specifies the connection of a sub-module's output to the multi-stage module's output. Works in a similar way to *ModuleInputConnection*.

**exception MultistageModulePreparationError**

Bases: `exceptions.Exception`

## pimlico.core.modules.options module

Utilities and type processors for module options.

**opt\_type\_help** (*help\_text*)

Decorator to add help text to functions that are designed to be used as module option processors. The help text will be used to describe the type in documentation.

**format\_option\_type** (*t*)

**str\_to\_bool** (*string*)

Convert a string value to a boolean in a sensible way. Suitable for specifying booleans as options.

**Parameters** **string** – input string

**Returns** boolean value

**choose\_from\_list** (*options, name=None*)

Utility for option processors to limit the valid values to a list of possibilities.

**comma\_separated\_list** (*item\_type=<type 'str'>, length=None*)

Option processor type that accepts comma-separated lists of strings. Each value is then parsed according to the given *item\_type* (default: `string`).

**comma\_separated\_strings** (*string*)

**json\_string** (*string*)

**process\_module\_options** (*opt\_def, opt\_dict, module\_type\_name*)

Utility for processing runtime module options. Called from module base class.

**Parameters**

- **opt\_def** – dictionary defining available options
- **opt\_dict** – dictionary of option values
- **module\_type\_name** – name for error output

**Returns** dictionary of processed options

**exception ModuleOptionParseError**

Bases: `exceptions.Exception`

## Module contents

Core functionality for loading and executing different types of pipeline module.

## pimlico.core.visualize package

### Submodules

#### pimlico.core.visualize.deps module

```
class GraphvizDependency (**kwargs)
    Bases: pimlico.core.dependencies.base.SystemCommandDependency
    installation_instructions ()
```

#### pimlico.core.visualize.status module

```
build_graph_with_status (pipeline, all=False)
```

### Module contents

Visualization of pipelines using Graphviz.

This is not fully implemented yet. In fact, I've barely started. But you'll find some indication of where I'm going with it on the [Pimlico Wishlist](#).

---

**Note:** Do not import anything from subpackages unless you're doing graph visualization, as they will trigger a check for Graphviz and try to install it.

---



---

**Note:** A bug in pygraphviz means that automatic installation on Ubuntu (and perhaps other systems) gets in a twist and leaves an unresolved dependency. If you have this problem and system-wide install is an option, just install with *sudo apt-get install python-pygraphviz*.

---

### Submodules

#### pimlico.core.config module

Reading of pipeline config from a file into the data structure used to run and manipulate the pipeline's data.

```
class PipelineConfig (name, pipeline_config, local_config, filename=None, variant='main', available_variants=[], log=None, all_filenames=None, module_aliases={}, local_config_sources=None)
```

Bases: `object`

Main configuration for a pipeline, read in from a config file.

For details on how to write config files that get read by this class, see [Pipeline config](#).

#### modules

List of module names, in the order they were specified in the config file.

#### module\_dependencies

Dictionary mapping a module name to a list of the names of modules that it depends on for its inputs.

**module\_dependents**

Opposite of `module_dependencies`. Returns a mapping from module names to a list of modules the depend on the module.

**get\_dependent\_modules** (*module\_name*, *recurse=False*, *exclude=[]*)

Return a list of the names of modules that depend on the named module for their inputs.

If *exclude* is given, we don't perform a recursive call on any of the modules in the list. For each item we recurse on, we extend the *exclude* list in the recursive call to include everything found so far (in other recursive calls). This avoids unnecessary recursion in complex pipelines.

If *exclude=None*, it is also passed through to recursive calls as *None*. Its default value of *[]* avoids excessive recursion from the top-level call, by allowing things to be added to the exclusion list for recursive calls.

**Parameters** *recurse* – include all transitive dependents, not just those that immediately depend on the module.

**append\_module** (*module\_info*)

Add a `moduleinfo` to the end of the pipeline. This is mainly for use while loaded a pipeline from a config file.

**get\_module\_schedule** ()

Work out the order in which modules should be executed. This is an ordering that respects dependencies, so that modules are executed after their dependencies, but otherwise follows the order in which modules were specified in the config.

**Returns** list of module names

**reset\_all\_modules** ()

Resets the execution states of all modules, restoring the output dirs as if nothing's been run.

**path\_relative\_to\_config** (*path*)

Get an absolute path to a file/directory that's been specified relative to a config file (usually within the config file).

**Parameters** *path* – relative path

**Returns** absolute path

**static load** (*filename*, *local\_config=None*, *variant='main'*, *override\_local\_config={}*)

Main function that loads a pipeline from a config file.

**Parameters**

- **filename** – file to read config from
- **local\_config** – location of local config file, where we'll read system-wide config. Usually not specified, in which case standard locations are searched. When loading programmatically, you might want to give this
- **variant** – pipeline variant to load
- **override\_local\_config** – extra configuration values to override the system-wide config

**Returns****static load\_local\_config** (*filename=None*, *override={}*)

Load local config parameters. These are usually specified in a `.pimlico` file, but may be overridden by other config locations, on the command line, or elsewhere programmatically.

**static empty** (*local\_config=None*, *override\_local\_config={}*, *override\_pipeline\_config={}*)

Used to programmatically create an empty pipeline. It will contain no modules, but provides a gateway to system info, etc and can be used in place of a real Pimlico pipeline.



**Parameters**

- **local\_config** – filename to load local config from. If not given, the default locations are searched
- **override\_local\_config** – manually override certain local config parameters. Dict of parameter values

**Returns** the *PipelineConfig* instance

**find\_data\_path** (*path*, *default=None*)

Given a path to a data dir/file relative to a data store, tries taking it relative to various store base dirs. If it exists in a store, that absolute path is returned. If it exists in no store, return None. If the path is already an absolute path, nothing is done to it.

The stores searched are the long-term store and the short-term store, though in the future more valid data storage locations may be added.

**Parameters**

- **path** – path to data, relative to store base
- **default** – usually, return None if no data is found. If default="short", return path relative to short-term store in this case. If default="long", long-term store.

**Returns** absolute path to data, or None if not found in any store

**find\_all\_data\_paths** (*path*)

**get\_data\_search\_paths** (*path*)

Like *find\_all\_data\_paths()*, but returns a list of all absolute paths which this data path could correspond to, whether or not they exist.

**Parameters** **path** – relative path within Pimlico directory structures

**Returns** list of string

**get\_storage\_roots** ()

Returns a list of all the (pipeline-specific) storage root locations known to the pipeline.

Currently, this is always [*self.short\_term\_store*, *self.long\_term\_store*], but in future we may have a more flexible system that allows an unbounded number of storage locations.

**step**

**enable\_step** ()

Enable super-verbose, interactive step mode.

::seealso:

Module :mod:pimlico.cli.debug  
The debug module defines the behaviour of step mode.

**exception PipelineConfigParseError** (*\*args*, *\*\*kwargs*)

Bases: `exceptions.Exception`

General problems interpreting pipeline config

**exception PipelineStructureError**

Bases: `exceptions.Exception`

Fundamental structural problems in a pipeline.

**exception PipelineCheckError** (*cause*, *\*args*, *\*\*kwargs*)

Bases: `exceptions.Exception`

Error in the process of explicitly checking a pipeline for problems.

**preprocess\_config\_file** (*filename*, *variant*='main', *initial\_vars*={})

Workhorse of the initial part of config file reading. Deals with all of our custom stuff for pipeline configs, such as preprocessing directives and includes.

**Parameters**

- **filename** – file from which to read main config
- **variant** – name of a variant to load. The default (*main*) loads the main variant, which always exists
- **initial\_vars** – variable assignments to make available for substitution. This will be added to by any *vars* sections that are read.

**Returns** tuple: raw config dict; list of variants that could be loaded; final vars dict; list of filenames that were read, including included files; dict of docstrings for each config section

**check\_for\_cycles** (*pipeline*)

Basic cyclical dependency check, always run on pipeline before use.

**check\_release** (*release\_str*)

Check a release name against the current version of Pimlico to determine whether we meet the requirement.

**check\_pipeline** (*pipeline*)

Checks a pipeline over for metadata errors, cycles, module typing errors and other problems. Called every time a pipeline is loaded, to check the whole pipeline's metadata is in order.

Raises a *PipelineCheckError* if anything's wrong.

**get\_dependencies** (*pipeline*, *modules*, *recursive*=False, *sources*=False)

Get a list of software dependencies required by the subset of modules given.

If recursive=True, dependencies' dependencies are added to the list too.

**Parameters**

- **pipeline** –
- **modules** – list of modules to check. If None, checks all modules

**print\_missing\_dependencies** (*pipeline*, *modules*)

Check runtime dependencies for a subset of modules and output a table of missing dependencies.

**Parameters**

- **pipeline** –
- **modules** – list of modules to check. If None, checks all modules

**Returns** True if no missing dependencies, False otherwise

**print\_dependency\_leaf\_problems** (*dep*, *local\_config*)

## pimlico.core.logs module

**get\_log\_file** (*name*)

Returns the path to a log file that may be used to output helpful logging info. Typically used to output verbose error information if something goes wrong. The file can be found in the Pimlico log dir.

**Parameters** **name** – identifier to distinguish from other logs

**Returns** path

## pimlico.core.paths module

`abs_path_or_model_dir_path(path, model_type)`

## Module contents

### pimlico.datatypes package

#### Subpackages

#### pimlico.datatypes.coref package

#### Submodules

#### pimlico.datatypes.coref.corenlp module

Datatypes for coreference resolution output. Based on Stanford CoreNLP's coref output, so includes all the information provided by that.

```
class CorefDocumentType (options, metadata)
    Bases: pimlico.datatypes.jsondoc.JsonDocumentType
    process_document (doc)
```

```
class CorefCorpus (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.jsondoc.JsonDocumentCorpus
    datatype_name = 'corenlp_coref'
    data_point_type
        alias of CorefDocumentType
```

```
class CorefCorpusWriter (base_dir, gzip=False, append=False, trust_length=False, encoding='utf-8',
                        **kwargs)
    Bases: pimlico.datatypes.tar.TarredCorpusWriter
    document_to_raw_data (data)
```

```
class Entity (id, mentions)
    Bases: object
```

```
class Mention (id, sentence_num, start_index, end_index, text, type, position=None, animacy=None,
               is_representative_mention=None, number=None, gender=None)
    Bases: object
    static from_json (json)
    to_json_dict ()
```

#### pimlico.datatypes.coref.opennlp module

Datatypes for coreference resolution output. Based on OpenNLP's coref output, so includes all the information provided by that. This is a slight different set of information to CoreNLP. Currently, there's no way to convert between the two datatypes, but in future it will be easy to provide an adapter that carries across the information common to the two (which for most purposes will be sufficient).

```
class CorefDocumentType (options, metadata)
    Bases: pimlico.datatypes.jsondoc.JsonDocumentType

    process_document (doc)

class CorefCorpus (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.jsondoc.JsonDocumentCorpus

    datatype_name = 'opennlp_coref'

    data_point_type
        alias of CorefDocumentType

class CorefCorpusWriter (base_dir, readable=False, **kwargs)
    Bases: pimlico.datatypes.jsondoc.JsonDocumentCorpusWriter

    document_to_raw_data (data)

class Entity (id, mentions, category=None, gender=None, gender_prob=None, number=None, number_prob=None)
    Bases: object

    get_head_word (pronouns=['i', 'you', 'he', 'she', 'it', 'we', 'they', 'me', 'him', 'her', 'us', 'them', 'myself', 'yourself', 'himself', 'herself', 'itself', 'ourselves', 'themselves', 'my', 'your', 'his', 'its', "it's", 'our', 'their', 'mine', 'yours', 'ours', 'theirs', 'this', 'that', 'those', 'these'])
        Retrieve a head word from the entity's mentions if possible. Returns None if no suitable head word can be found: e.g., if all mentions are pronouns.

        Pronouns are filtered out using :data:pimlico.utils.linguistic.ENGLISH_PRONOUNS by default. You can override this with the pronouns kwargs. If pronouns=None, no filtering is done.

    to_json_dict ()

    static from_json (json)

    static from_java_object (obj)

class Mention (sentence_num, start_index, end_index, text, gender=None, gender_prob=None, number=None, number_prob=None, head_start_index=None, head_end_index=None, name_type=None)
    Bases: object

    static from_json (json)

    to_json_dict ()

    static from_java_object (obj)
```

## Module contents

```
OpenNLPCorefCorpus
    alias of CorefCorpus

OpenNLPCorefCorpusWriter
    alias of CorefCorpusWriter

CoreNLPCorefCorpus
    alias of CorefCorpus

CoreNLPCorefCorpusWriter
    alias of CorefCorpusWriter
```

## pimlico.datatypes.formatters package

### Submodules

#### pimlico.datatypes.formatters.features module

```
class FeatureListScoreFormatter (corpus)
    Bases: pimlico.cli.browser.formatter.DocumentBrowserFormatter

    DATATYPE
        alias of FeatureListScoreDocumentType

    format_document (doc)
```

#### pimlico.datatypes.formatters.tokenized module

```
class TokenizedDocumentFormatter (corpus, raw_data=False)
    Bases: pimlico.cli.browser.formatter.DocumentBrowserFormatter

    DATATYPE
        alias of TokenizedDocumentType

    format_document (doc)
```

### Module contents

## pimlico.datatypes.parse package

### Submodules

#### pimlico.datatypes.parse.candc module

```
class CandcOutputCorpus (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.tar.TarredCorpus

    datatype_name = 'candc_output'

    data_point_type
        alias of CandcOutputDocumentType

class CandcOutputCorpusWriter (base_dir, gzip=False, append=False, trust_length=False,
                               encoding='utf-8', **kwargs)
    Bases: pimlico.datatypes.tar.TarredCorpusWriter

    document_to_raw_data (data)
```

#### pimlico.datatypes.parse.dependency module

```
class StanfordDependencyParseCorpus (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.jsondoc.JsonDocumentCorpus

    datatype_name = 'stanford_dependency_parsers'
```

**data\_point\_type**  
alias of `StanfordDependencyParseDocumentType`

**class StanfordDependencyParseCorpusWriter** (*base\_dir, readable=False, \*\*kwargs*)  
Bases: `pimlico.datatypes.jsondoc.JsonDocumentCorpusWriter`

**document\_to\_raw\_data** (*data*)

**class CoNLLDependencyParseCorpus** (*base\_dir, pipeline, \*\*kwargs*)  
Bases: `pimlico.datatypes.word_annotations.WordAnnotationCorpus`

10-field CoNLL dependency parse format (conllx) – i.e. post parsing.

**Fields are:** id (int), word form, lemma, coarse POS, POS, features, head (int), dep relation, phead (int), pdeprel

The last two are usually not used.

**datatype\_name** = 'conll\_dependency\_parses'

**data\_point\_type**  
alias of `CoNLLDependencyParseDocumentType`

**class CoNLLDependencyParseCorpusWriter** (*base\_dir, \*\*kwargs*)  
Bases: `pimlico.datatypes.word_annotations.WordAnnotationCorpusWriter`

**document\_to\_raw\_data** (*data*)

**class CoNLLDependencyParseInputCorpus** (*base\_dir, pipeline, \*\*kwargs*)  
Bases: `pimlico.datatypes.word_annotations.WordAnnotationCorpus`

The version of the CoNLL format (conllx) that only has the first 6 columns, i.e. no dependency parse yet annotated.

**datatype\_name** = 'conll\_dependency\_parse\_inputs'

**data\_point\_type**  
alias of `CoNLLDependencyParseInputDocumentType`

**class CoNLLDependencyParseInputCorpusWriter** (*base\_dir, \*\*kwargs*)  
Bases: `pimlico.datatypes.word_annotations.WordAnnotationCorpusWriter`

**document\_to\_raw\_data** (*data*)

## Module contents

**TODO** Parse trees are temporary implementations that don't actually parse the data, but just split it into sentences.

**class TreeStringsDocumentType** (*options, metadata*)  
Bases: `pimlico.datatypes.documents.RawDocumentType`

**process\_document** (*doc*)

**class ConstituencyParseTreeCorpus** (*base\_dir, pipeline, \*\*kwargs*)  
Bases: `pimlico.datatypes.tar.TarredCorpus`

Note that this is not fully developed yet. At the moment, you'll just get, for each document, a list of the texts of each tree. In future, they will be better represented.

**datatype\_name** = 'parse\_trees'

**data\_point\_type**  
alias of `TreeStringsDocumentType`

```

class ConstituencyParseTreeCorpusWriter (base_dir,          gzip=False,          append=False,
                                         trust_length=False, encoding='utf-8', **kwargs)
    Bases: pimlico.datatypes.tar.TarredCorpusWriter

    document_to_raw_data (data)

class CandcOutputCorpus (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.tar.TarredCorpus

    data_point_type
        alias of CandcOutputDocumentType

    datatype_name = 'candc_output'

class CandcOutputCorpusWriter (base_dir,  gzip=False,  append=False,  trust_length=False,
                               encoding='utf-8', **kwargs)
    Bases: pimlico.datatypes.tar.TarredCorpusWriter

    document_to_raw_data (data)

class StanfordDependencyParseCorpus (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.jsondoc.JsonDocumentCorpus

    data_point_type
        alias of StanfordDependencyParseDocumentType

    datatype_name = 'stanford_dependency_pares'

class StanfordDependencyParseCorpusWriter (base_dir, readable=False, **kwargs)
    Bases: pimlico.datatypes.jsondoc.JsonDocumentCorpusWriter

    document_to_raw_data (data)

class CoNLLDependencyParseCorpus (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.word_annotations.WordAnnotationCorpus

    10-field CoNLL dependency parse format (conllx) – i.e. post parsing.

    Fields are: id (int), word form, lemma, coarse POS, POS, features, head (int), dep relation, phead (int), pdeprel

    The last two are usually not used.

    data_point_type
        alias of CoNLLDependencyParseDocumentType

    datatype_name = 'conll_dependency_pares'

class CoNLLDependencyParseCorpusWriter (base_dir, **kwargs)
    Bases: pimlico.datatypes.word_annotations.WordAnnotationCorpusWriter

    document_to_raw_data (data)

class CoNLLDependencyParseInputCorpus (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.word_annotations.WordAnnotationCorpus

    The version of the CoNLL format (conllx) that only has the first 6 columns, i.e. no dependency parse yet
    annotated.

    data_point_type
        alias of CoNLLDependencyParseInputDocumentType

    datatype_name = 'conll_dependency_parse_inputs'

class CoNLLDependencyParseInputCorpusWriter (base_dir, **kwargs)
    Bases: pimlico.datatypes.word_annotations.WordAnnotationCorpusWriter

    document_to_raw_data (data)

```

## Submodules

### pimlico.datatypes.arrays module

Wrappers around Numpy arrays and Scipy sparse matrices.

```
class NumpyArray (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.files.NamedFileCollection
```

```
    datatype_name = 'numpy_array'
```

```
    filenames = ['array.npy']
```

```
    array
```

```
    get_software_dependencies ()
```

```
class NumpyArrayWriter (base_dir, additional_name=None)
    Bases: pimlico.datatypes.base.PimlicoDatatypeWriter
```

```
    set_array (array)
```

```
class ScipySparseMatrix (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.base.PimlicoDatatype
```

Wrapper around Scipy sparse matrices. The matrix loaded is always in COO format – you probably want to convert to something else before using it. See scipy docs on sparse matrix conversions.

```
    datatype_name = 'scipy_sparse_array'
```

```
    filenames = ['array.mtx']
```

```
    array
```

```
    get_software_dependencies ()
```

```
class ScipySparseMatrixWriter (base_dir, additional_name=None)
    Bases: pimlico.datatypes.base.PimlicoDatatypeWriter
```

```
    set_matrix (mat)
```

### pimlico.datatypes.base module

Datatypes provide interfaces for reading (and in some cases writing) datasets. At their most basic, they define a way to iterate over a dataset linearly. Some datatypes may also provide other functionality, such as random access or compression.

As much as possible, Pimlico pipelines should use standard datatypes to connect up the output of modules with the input of others. Most datatypes have a lot in common, which should be reflected in their sharing common base classes. Custom datatypes will be needed for most datasets when they're used as inputs, but as far as possible, these should be converted into standard datatypes like *TarredCorpus*, early in the pipeline.

---

**Note:** The following classes were moved to *core* in version 0.6rc

---

```
class PimlicoDatatype (base_dir, pipeline, module=None, additional_name=None,
    use_main_metadata=False, **kwargs)
```

```
    Bases: object
```

The abstract superclass of all datatypes. Provides basic functionality for identifying where data should be stored and such.



Datatypes are used to specify the routines for reading the output from modules. They're also used to specify how to read pipeline inputs. Most datatypes that have data simply read it in when required. Some (in particular those used as inputs) need a preparation phase to be run, where the raw data itself isn't sufficient to implement the reading interfaces required. In this case, they should override `prepare_data()`.

Datatypes may require/allow options to be set when they're used to read pipeline inputs. These are specified, in the same way as module options, by `input_module_options` on the datatype class.

Datatypes may supply a set of additional datatypes. These should be guaranteed to be available if the main datatype is available. They must require no extra processing to be made available, unless that is done on the fly while reading the datatype (like a filter) or while the main datatype is being written.

Additional datatypes can be accessed in config files by specifying the main datatype (as a previous module, optionally with an output name) and the additional datatype name in the form *main\_datatype->additional\_name*. Multiple additional names may be given, causing the next name to be looked up as an additional datatype of the initially loaded additional datatype. E..g *main\_datatype->additional0->additional1*.

To avoid conflicts in the metadata between datatypes using the same directory, datatypes loaded as additional datatypes have their additional name available to them and use it as a prefix to the metadata filename.

If *use\_main\_metadata=True* on an additional datatype, the same metadata will be read as for the main datatype to which this is an additional datatype.

*module* is the `ModuleInfo` instance for the pipeline module that this datatype was produced by. It may be `None`, if the datatype wasn't instantiated by a module. It is not required to be set if you're instantiating a datatype in some context other than module output. It should generally be set for input datatypes, though, since they are treated as being created by a special input module.

**`requires_data_preparation = False`**

**`input_module_options = {}`**

Override to provide shell commands specific to this datatype. Should include the superclass' list.

**`shell_commands = []`**

List of additional datatypes provided by this one, given as (name, datatype class) pairs. For each of these, a call to *get\_additional\_datatype(name)* (once the main datatype is ready) should return a datatype instance that is also ready.

**`supplied_additional = []`**

Most datatype classes define their own type of corpus, which is often a subtype of some other. Some, however, emulate another type and it is that type that should be considered to be the type of the dataset, not the class itself.

For example, `TarredCorpusFilter` dynamically produces something that looks like a `TarredCorpus`, and further down the pipeline, if its type is need, it should be considered to be a `TarredCorpus`.

Most of the time, this can be left empty, but occasionally it needs to be set.

**`emulated_datatype = None`**

**`datatype_name = 'base_datatype'`**

**`metadata`**

Read in metadata from a file in the corpus directory.

Note that this is no longer cached in memory. We need to be sure that the metadata values returned are always up to date with what is on disk, so always re-read the file when we need to get a value from the metadata. Since the file is typically small, this is unlikely to cause a problem. If we decide to return to cacheing the metadata dictionary in future, we will need to make sure that we can never run into problems with out-of-date metadata being returned.

**get\_required\_paths()**

Returns a list of paths to files that should be available for the data to be read. The base `data_ready()` implementation checks that these are all available and, if the datatype is used as an input to a pipeline and requires a data preparation routine to be run, data preparation will not be executed until these files are available.

Paths may be absolute or relative. If relative, they refer to files within the data directory and `data_ready()` will fail if the data dir doesn't exist.

**Returns** list of absolute or relative paths

**get\_software\_dependencies()**

Check that all software required to read this datatype is installed and locatable. This is separate to metadata config checks, so that you don't need to satisfy the dependencies for all modules in order to be able to run one of them. You might, for example, want to run different modules on different machines. This is called when a module is about to be executed and each of the dependencies is checked.

Returns a list of instances of subclasses of `:class:~pimlico.core.dependencies.base.SoftwareDependency`, representing the libraries that this module depends on.

Take care when providing dependency classes that you don't put any import statements at the top of the Python module that will make loading the dependency type itself dependent on runtime dependencies. You'll want to run import checks by putting import statements within this method.

You should call the super method for checking superclass dependencies.

**prepare\_data(output\_dir, log)****classmethod create\_from\_options(base\_dir, pipeline, options={}, module=None)****data\_ready()**

Check whether the data corresponding to this datatype instance exists and is ready to be read.

Default implementation just checks whether the data dir exists. Subclasses might want to add their own checks, or even override this, if the data dir isn't needed.

**get\_detailed\_status()**

Returns a list of strings, containing detailed information about the data. Only called if `data_ready() == True`.

Subclasses may override this to supply useful (human-readable) information specific to the datatype. They should call the super method.

**classmethod datatype\_full\_class\_name()**

The fully qualified name of the class for this datatype, by which it is reference in config files. Generally, datatypes don't need to override this, but type requirements that take the place of datatypes for type checking need to provide it.

**instantiate\_additional\_datatype(name, additional\_name)**

Default implementation just assumes the datatype class can be instantiated using the default constructor, with the same base dir and pipeline as the main datatype. Options given to the main datatype are passed down to the additional datatype.

**classmethod check\_type(supplied\_type)**

Method used by datatype type-checking algorithm to determine whether a supplied datatype (given as a class, which is a subclass of `PimlicoDatatype`) is compatible with the present datatype, which is being treated as a type requirement.

Typically, the present class is a type requirement on a module input and `supplied_type` is the type provided by a previous module's output.

The default implementation simply checks whether *supplied\_type* is a subclass of the present class. Subclasses may wish to impose different or additional checks.

**Parameters** *supplied\_type* – type provided where the present class is required, or datatype instance

**Returns** True if the check is successful, False otherwise

**classmethod** *type\_checking\_name* ()

Supplies a name for this datatype to be used in type-checking error messages. Default implementation just provides the class name. Classes that override *check\_supplied\_type()* may want to override this too.

**classmethod** *full\_datatype\_name* ()

Returns a string/unicode name for the datatype that includes relevant sub-type information. The default implementation just uses the attribute *datatype\_name*, but subclasses may have more detailed information to add. For example, iterable corpus types also supply information about the data-point type.

**class** *DynamicOutputDatatype*

Bases: *object*

Types of module outputs may be specified as a subclass of *PimlicoDatatype*, or alternatively as an *instance* of *DynamicOutputType*. In this case, *get\_datatype()* is called when the output datatype is needed, passing in the module info instance for the module, so that a specialized datatype can be produced on the basis of options, input types, etc.

The dynamic type must provide certain pieces of information needed for typechecking.

**datatype\_name** = None

**get\_datatype** (*module\_info*)

**get\_base\_datatype\_class** ()

If it's possible to say before the instance of a *ModuleInfo* is available what base datatype will be produced, implement this to return the class. By default, it returns None.

If this information is available, it will be used in documentation.

**class** *DynamicInputDatatypeRequirement*

Bases: *object*

Types of module inputs may be given as a subclass of *PimlicoDatatype*, a tuple of datatypes, or an instance a *DynamicInputDatatypeRequirement* subclass. In this case, *check\_type(supplied\_type)* is called during typechecking to check whether the type that we've got conforms to the input type requirements.

Additionally, if *datatype\_doc\_info* is provided, it is used to represent the input type constraints in documentation.

**datatype\_doc\_info** = None

**check\_type** (*supplied\_type*)

**type\_checking\_name** ()

Supplies a name for this datatype to be used in type-checking error messages. Default implementation just provides the class name. Subclasses may want to override this too.

**class** *MultipleInputs* (*datatype\_requirements*)

Bases: *object*

An input datatype that can be used as an item in a module's inputs, which lets the module accept an unbounded number of inputs, all satisfying the same datatype requirements. When writing the inputs in a config file, they can be specified as a comma-separated list of the usual type of specification (module name, with optional output name). Each item in the list must point to a datatype that satisfies the type-checking.

The list may also include (or entirely consist of) a base module name from the pipeline that has been expanded into multiple modules according to alternative parameters (the type separated by vertical bars, see [Multiple parameter values](#)). Use the notation *\*name*, where *name* is the base module name, to denote all of the expanded module names as inputs. These are treated as if you'd written out all of the expanded module names separated by commas.

When `get_input()` is called on the module, instead of returning a single datatype, a list of datatypes is returned.

**class PimlicoDatatypeWriter** (*base\_dir*, *additional\_name=None*)

Bases: `object`

Abstract base class for data writer associated with Pimlico datatypes.

**require\_tasks** (*\*tasks*)

Add a name or multiple names to the list of output tasks that must be completed before writing is finished

**task\_complete** (*task*)

**incomplete\_tasks**

**write\_metadata** ()

**subordinate\_additional\_name** (*name*)

**class IterableCorpus** (*\*args*, *\*\*kwargs*)

Bases: `pimlico.datatypes.base.PimlicoDatatype`

Superclass of all datatypes which represent a dataset that can be iterated over document by document (or data-point by datapoint - what exactly we're iterating over may vary, though documents are most common).

The actual type of the data depends on the subclass: it could be, e.g. coref output, etc. Information about the type of individual documents is provided by *document\_type* and this is used in type checking.

At creation time, length should be provided in the metadata, denoting how many documents are in the dataset.

**datatype\_name** = 'iterable\_corpus'

**data\_point\_type**

alias of `RawDocumentType`

**shell\_commands** = [`<pimlico.datatypes.base.CountInvalidCmd object>`]

**get\_detailed\_status** ()

**classmethod check\_supplied\_type** (*supplied\_type*)

Override type checking to require that the supplied type have a document type that is compatible with (i.e. a subclass of) the document type of this class.

**classmethod type\_checking\_name** ()

**classmethod full\_datatype\_name** ()

**process\_document\_data\_with\_datatype** (*data*)

Applies the corpus' datatype's `process_document()` method to the raw data :param data: :return:

**class IterableCorpusWriter** (*base\_dir*, *additional\_name=None*)

Bases: `pimlico.datatypes.base.PimlicoDatatypeWriter`

**class InvalidDocument** (*module\_name*, *error\_info=None*)

Bases: `object`

Widely used in Pimlico to represent an empty document that is empty not because the original input document was empty, but because a module along the way had an error processing it. Document readers/writers should generally be robust to this and simply pass through the whole thing where possible, so that it's always possible to work out, where one of these pops up, where the error occurred.

**static load** (*text*)

**static invalid\_document\_or\_text** (*text*)

If the text represents an invalid document, parse it and return an InvalidDocument object. Otherwise, return the text as is.

**exception DatatypeLoadError**

Bases: `exceptions.Exception`

**exception DatatypeWriteError**

Bases: `exceptions.Exception`

**load\_datatype** (*path*)

Try loading a datatype class for a given path. Raises a DatatypeLoadError if it's not a valid datatype path.

### pimlico.datatypes.caevo module

**class CaevoCorpus** (*base\_dir, pipeline, \*\*kwargs*)

Bases: `pimlico.datatypes.tar.TarredCorpus`

Datatype for Caevo output. The output is stored exactly as it comes out from Caevo, in an XML format. This datatype reads in that XML and provides easy access to its components.

Since we simply store the XML that comes from Caevo, there's no corresponding corpus writer. The data is output using a `:class:pimlico.datatypes.tar.TarredCorpusWriter`.

**data\_point\_type**

alias of `CaevoDocumentType`

### pimlico.datatypes.core module

Some basic core datatypes that are commonly used for simple datatypes, file types, etc.

**class SingleTextDocument** (*base\_dir, pipeline, module=None, additional\_name=None, use\_main\_metadata=False, \*\*kwargs*)

Bases: `pimlico.datatypes.files.NamedFileCollection`

**datatype\_name** = 'single\_doc'

**filenames** = ['data.txt']

**read\_data** ()

**class SingleTextDocumentWriter** (*base\_dir, \*\*kwargs*)

Bases: `pimlico.datatypes.base.PimlicoDatatypeWriter`

**class Dict** (*base\_dir, pipeline, module=None, additional\_name=None, use\_main\_metadata=False, \*\*kwargs*)

Bases: `pimlico.datatypes.files.NamedFileCollection`

Simply stores a Python dict, pickled to disk.

**datatype\_name** = 'dict'

**filenames** = ['data']

**data**

**class DictWriter** (*base\_dir, \*\*kwargs*)

Bases: `pimlico.datatypes.base.PimlicoDatatypeWriter`

```
class StringList (base_dir, pipeline, module=None, additional_name=None, use_main_metadata=False,  
                 **kwargs)
```

Bases: `pimlico.datatypes.base.PimlicoDatatype`

Simply stores a Python list of strings, written out to disk in a readable form. Not the most efficient format, but if the list isn't humungous it's OK (e.g. storing vocabularies).

```
datatype_name = 'string_list'
```

```
data_ready ()
```

```
path
```

```
data
```

```
class StringListWriter (base_dir, **kwargs)
```

Bases: `pimlico.datatypes.base.PimlicoDatatypeWriter`

## pimlico.datatypes.dictionary module

This module implements the concept of Dictionary – a mapping between words and their integer ids.

The implementation is based on Gensim, because Gensim is wonderful and there's no need to reinvent the wheel. We don't use Gensim's data structure directly, because it's unnecessary to depend on the whole of Gensim just for one data structure.

```
class Dictionary (base_dir, pipeline, **kwargs)
```

Bases: `pimlico.datatypes.base.PimlicoDatatype`

Dictionary encapsulates the mapping between normalized words and their integer ids.

```
datatype_name = 'dictionary'
```

```
get_data ()
```

```
data_ready ()
```

```
get_detailed_status ()
```

```
class DictionaryWriter (base_dir)
```

Bases: `pimlico.datatypes.base.PimlicoDatatypeWriter`

```
add_documents (documents, prune_at=2000000)
```

```
filter (threshold=None, no_above=None, limit=None)
```

## pimlico.datatypes.documents module

Document types used to represent datatypes of individual documents in an IterableCorpus or subtype.

```
class DataPointType (options, metadata)
```

Bases: `object`

Base data-point type for iterable corpora. All iterable corpora should have data-point types that are subclasses of this.

```
input_module_options = {}
```

```
formatters = []
```

List of (name, cls\_path) pairs specifying a standard set of formatters that the user might want to choose from to view a dataset of this type. The user is not restricted to this set, but can easily choose these by

name, instead of specifying a class path themselves. The first in the list is the default used if no formatter is specified. Falls back to DefaultFormatter if empty

**class RawDocumentType** (*options, metadata*)

Bases: `pimlico.datatypes.documents.DataPointType`

Base document type. All document types for tarred corpora should be subclasses of this.

It may be used itself as well, where documents are just treated as raw data, though most of the time it will be appropriate to use subclasses to provide more information and processing operations specific to the datatype.

**process\_document** (*doc*)

**class RawTextDocumentType** (*options, metadata*)

Bases: `pimlico.datatypes.documents.RawDocumentType`

Subclass of RawDocumentType used to indicate that the document represents text (not just any old raw data) and that it hasn't been processed (tokenized, etc). Note that text that has been tokenized, parsed, etc does not used subclasses of this type, so they will not be considered compatible if this type is used as a requirement.

**input\_module\_options** = {'encoding': {'default': 'utf8', 'help': 'Encoding to assume for input files. Default: utf8'}}

**process\_document** (*doc*)

## pimlico.datatypes.embeddings module

**class Vocab** (*word, index, count=0*)

Bases: `object`

A single vocabulary item, used internally for collecting per-word frequency info. A simplified version of Gensim's Vocab.

**class Embeddings** (*base\_dir, pipeline, \*\*kwargs*)

Bases: `pimlico.datatypes.base.PimlicoDatatype`

Datatype to store embedding vectors, together with their words. Based on Gensim's `KeyedVectors` object, but adapted for use in Pimlico and so as not to depend on Gensim. (This means that this can be used more generally for storing embeddings, even when we're not depending on Gensim.)

Provides a method to map to Gensim's `KeyedVectors` type for compatibility.

Doesn't provide all of the functionality of `KeyedVectors`, since the main purpose of this is for storage of vectors and other functionality, like similarity computations, can be provided by utilities or by direct use of Gensim.

**vectors**

**normed\_vectors**

**vector\_size**

**index2vocab**

**index2word**

**vocab**

**word\_vec** (*word*)

Accept a single word as input. Returns the word's representation in vector space, as a 1D numpy array.

**word\_vecs** (*words*)

Accept multiple words as input. Returns the words' representations in vector space, as a 1D numpy array.

**to\_keyed\_vectors** ()

```
class EmbeddingsWriter (base_dir, **kwargs)
    Bases: pimlico.datatypes.base.PimlicoDatatypeWriter

    write_vectors (arr)
        Write out vectors from a Numpy array

    write_word_counts (word_counts)
        Write out vocab from a list of words with counts.

        Parameters word_counts – list of (unicode, int) pairs giving each word and its count. Vocab
        indices are determined by the order of words

    write_vocab_list (vocab_items)
        Write out vocab from a list of vocab items (see Vocab).

        Parameters vocab_items – list of “Vocab”s

    write_keyed_vectors (*kvecs)
        Write both vectors and vocabulary straight from Gensim’s KeyedVectors data structure. Can accept
        multiple objects, which will then be concatenated in the output.
```

### pimlico.datatypes.features module

```
class KeyValueListDocumentType (options, metadata)
    Bases: pimlico.datatypes.documents.RawDocumentType

    process_document (doc)

class KeyValueListCorpus (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.tar.TarredCorpus

    datatype_name = ‘key_value_lists’

    data_point_type
        alias of KeyValueListDocumentType

class KeyValueListCorpusWriter (base_dir, separator=’,’, fv_separator=‘=’, **kwargs)
    Bases: pimlico.datatypes.tar.TarredCorpusWriter

    document_to_raw_data (data)

class TermFeatureListDocumentType (options, metadata)
    Bases: pimlico.datatypes.features.KeyValueListDocumentType

    process_document (doc)

class TermFeatureListCorpus (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.features.KeyValueListCorpus

    Special case of KeyValueListCorpus, where one special feature “term” is always present and the other feature
    types are counts of the occurrence of a particular feature with this term in each data point.

    datatype_name = ‘term_feature_lists’

    data_point_type
        alias of TermFeatureListDocumentType

class TermFeatureListCorpusWriter (base_dir, **kwargs)
    Bases: pimlico.datatypes.features.KeyValueListCorpusWriter

    document_to_raw_data (data)
```



```
class IndexedTermFeatureListDataPointType (options, metadata)
```

Bases: `pimlico.datatypes.documents.DataPointType`

```
class IndexedTermFeatureListCorpus (*args, **kwargs)
```

Bases: `pimlico.datatypes.base.IterableCorpus`

Term-feature instances, indexed by a dictionary, so that all that's stored is the indices of the terms and features and the feature counts for each instance. This is iterable, but, unlike `TermFeatureListCorpus`, doesn't iterate over documents. Now that we've filtered extracted features down to a smaller vocab, we put everything in one big file, with one data point per line.

Since we're now storing indices, we can use a compact format that's fast to read from disk, making iterating over the dataset faster than if we had to read strings, look them up in the vocab, etc.

By default, the ints are stored as C longs, which use 4 bytes. If you know you don't need ints this big, you can choose 1 or 2 bytes, or even 8 (long long). By default, the ints are unsigned, but they may be signed.

**data\_point\_type**

alias of `IndexedTermFeatureListDataPointType`

**term\_dictionary**

**feature\_dictionary**

```
class IndexedTermFeatureListCorpusWriter (base_dir, term_dictionary, feature_dictionary,
                                           bytes=4, signed=False, index_input=False,
                                           **kwargs)
```

Bases: `pimlico.datatypes.base.IterableCorpusWriter`

`index_input=True` means that the input terms and feature names are already mapped to dictionary indices, so are assumed to be ints. Otherwise, inputs will be looked up in the appropriate dictionary to get an index.

**write\_dictionaries()**

**add\_data\_points(iterable)**

```
class FeatureListScoreDocumentType (options, metadata)
```

Bases: `pimlico.datatypes.documents.RawDocumentType`

Document type that stores a list of features, each associated with a floating-point score. The feature lists are simply lists of indices to a feature set for the whole corpus that includes all feature types and which is stored along with the dataset. These may be binary features (present or absent for each data point), or may have a weight associated with them. If they are binary, the returned data will have a weight of 1 associated with each.

A corpus of this type can be used to train, for example, a regression.

If scores and weights are passed in as Decimal objects, they will be stored as strings. If they are floats, they will be converted to Decimals via their string representation (avoiding some of the oddness of converting between binary and decimal representations). To avoid loss of precision, pass in all scores and weights as Decimal objects.

**formatters** = [('features', 'pimlico.datatypes.formatters.features.FeatureListScoreFormatter')]

**process\_document(doc)**

```
class FeatureListScoreCorpus (base_dir, pipeline, **kwargs)
```

Bases: `pimlico.datatypes.tar.TarredCorpus`

**datatype\_name** = 'scored\_weight\_feature\_lists'

**data\_point\_type**

alias of `FeatureListScoreDocumentType`

```
class FeatureListScoreCorpusWriter (base_dir, features, separator=':', index_input=False,
                                     **kwargs)
```

Bases: `pimlico.datatypes.tar.TarredCorpusWriter`

Input should be a list of data points. Each is a (score, feature list) pair, where score is a Decimal, or other numeric type. Feature list is a list of (feature name, weight) pairs, or just feature names. If weights are not given, they will default to 1 when read in (but no weight is stored).

If `index_input=True`, it is assumed that feature IDs will be given instead of feature names. Otherwise, the feature names will be looked up in the feature list. Any features not found in the feature type list will simply be skipped.

```
document_to_raw_data (data)
```

## pimlico.datatypes.files module

```
class File (base_dir, pipeline, module=None, additional_name=None, use_main_metadata=False,
            **kwargs)
```

Bases: `pimlico.datatypes.base.PimlicoDatatype`

Simple datatype that supplies a single file, providing the path to it. Use `FileCollection` with a single file where possible.

This is an abstract class: subclasses need to provide a way of getting to (e.g. storing) the filename in question.

This overlaps somewhat with `FileCollection`, but is mainly here for backwards compatibility. Future datatypes should prefer the use of `FileCollection`.

```
datatype_name = 'file'
```

```
data_ready ()
```

```
absolute_path
```

```
class NamedFileCollection (base_dir, pipeline, module=None, additional_name=None,
                           use_main_metadata=False, **kwargs)
```

Bases: `pimlico.datatypes.base.PimlicoDatatype`

Abstract base datatype for datatypes that store a fixed collection of files, which have fixed names (or at least names that can be determined from the class). Very many datatypes fall into this category. Overriding this base class provides them with some common functionality, including the possibility of creating a union of multiple datatypes.

The attribute `filenames` should specify a list of filenames contained by the datatype.

All files are contained in the datatypes data directory. If files are stored in subdirectories, this may be specified in the list of filenames using `/` s. (Always use forward slashes, regardless of the operating system.)

```
datatype_name = 'file_collection'
```

```
filenames = []
```

```
data_ready ()
```

```
get_absolute_path (filename)
```

```
class NamedFileCollectionWriter (base_dir)
```

Bases: `pimlico.datatypes.base.PimlicoDatatypeWriter`

```
filenames = []
```

```
write_file (filename, data)
```

```
get_absolute_path (filename)
```

**named\_file\_collection\_union** (\*file\_collection\_classes, \*\*kwargs)

Takes a number of subclasses of `FileCollection` and produces a new datatype that shares the functionality of all of them and is constituted of the union of the filenames.

The datatype name of the result will be produced automatically from the inputs, unless the `kwargs` name is given to specify a new one.

Note that the input classes' `__init__`'s will each be called once, with the standard `PimlicoDatatype` args. If this behaviour does not suit the datatypes you're using, override the `init` or define the union some other way.

**filename\_with\_range** (val)

Option processor for file paths with an optional start and end line at the end.

**class UnnamedFileCollection** (\*args, \*\*kwargs)

Bases: `pimlico.datatypes.base.IterableCorpus`

---

**Note:** Datatypes used for reading input data are being phased out and replaced by input reader modules. Use `pimlico.modules.input.text.raw_text_files` instead of this for reading raw text files at the start of your pipeline.

---

A file collection that's just a bunch of files with arbitrary names. The names are not necessarily known until the data is ready. They may be specified as a list in the metadata, or through datatype options, in the case of input datatypes.

This datatype is particularly useful for loading individual files or sets of files at the start of a pipeline. If you just want the raw data from each file, you can use this class as it is. It's an `IterableCorpus` with a raw data type. If you want to apply some special processing to each file, do so by overriding this class and specifying the `data_point_type`, providing a `DataPointType` subclass that does the necessary processing.

When using it as an input datatype to load arbitrary files, specify a list of absolute paths to the files you want to use. They must be absolute paths, but remember that you can make use of various *special substitutions in the config file* to give paths relative to your project root, or other locations.

The file paths may use `globs` to match multiple files. By default, it is assumed that every filename should exist and every glob should match at least one file. If this does not hold, the dataset is assumed to be not ready. You can override this by placing a `?` at the start of a filename/glob, indicating that it will be included if it exists, but is not depended on for considering the data ready to use.

The same postprocessing will be applied to every file. In cases where you need to apply different processing to different subsets of the files, define multiple input modules, with different data point types, for example, and then combine them using `pimlico.modules.corpora.concat`.

**datatype\_name** = 'unnamed\_file\_collection'

**input\_module\_options** = {'files': {'required': True, 'type': <function \_fn>, 'help': "Comma-separated list of absolute paths or globs to files to load."}}

**data\_ready** ()

**get\_paths** (error\_on\_missing=False)

**get\_paths\_from\_options** (error\_on\_missing=False)

Get a list of paths to all the files specified in the `files` option. If `error_on_missing=True`, non-optional paths or globs that do not correspond to an existing file cause an `IOError` to be raised.

**path\_name\_to\_doc\_name** (path)

**class UnnamedFileCollectionWriter** (\*args, \*\*kwargs)

Bases: `pimlico.datatypes.base.PimlicoDatatypeWriter`

Use as a context manager to write a bag of files out to the output directory.

Provide each file's raw data and a filename to use to the function `write_file()` within the `with` statement. The writer will keep track of what files you've output and store the list.

**write\_file** (*filename, data*)

**NamedFile** (*name*)

Datatype factory that produces something like a `File` datatype, pointing to a single file, but doesn't store its path, just refers to a particular file in the data dir.

**Parameters** **name** – name of the file

**Returns** datatype class

**class NamedFileWriter** (*base\_dir, filename, \*\*kwargs*)

Bases: `pimlico.datatypes.base.PimlicoDatatypeWriter`

**absolute\_path**

**write\_data** (*data*)

Write the given string data to the appropriate output file

**class RawTextFiles** (*\*args, \*\*kwargs*)

Bases: `pimlico.datatypes.files.UnnamedFileCollection`

Essentially the same as `RawTextDirectory`, but more flexible. Should generally be used in preference to `RawTextDirectory`.

Basic datatype for reading in all the files in a collection as raw text documents.

Generally, this may be appropriate to use as the input datatype at the start of a pipeline. You'll then want to pass it through a tarred corpus filter to get it into a suitable form for input to other modules.

**data\_point\_type**

alias of `RawTextDocumentType`

**class RawTextDirectory** (*\*args, \*\*kwargs*)

Bases: `pimlico.datatypes.base.IterableCorpus`

Basic datatype for reading in all the files in a directory and its subdirectories as raw text documents.

Generally, this may be appropriate to use as the input datatype at the start of a pipeline. You'll then want to pass it through a tarred corpus filter to get it into a suitable form for input to other modules.

**datatype\_name** = 'raw\_text\_directory'

**input\_module\_options** = {'path': {'required': True, 'help': 'Full path to the directory containing the files'}, 'encoding': 'utf-8'}

**data\_point\_type**

alias of `RawTextDocumentType`

**requires\_data\_preparation** = True

**prepare\_data** (*output\_dir, log*)

**walk** ()

**filter\_document** (*doc*)

Each document is passed through this filter before being yielded. Default implementation does nothing, but this makes it easy to implement custom postprocessing by overriding.

**get\_required\_paths** ()

## pimlico.datatypes.floats module

Similar to :mod:pimlico.datatypes.ints, but for lists of floats.

```
class FloatListsDocumentType (options, metadata)
    Bases: pimlico.datatypes.documents.RawDocumentType

    formatters = [('float_lists', 'pimlico.datatypes.floats.FloatListsFormatter')]

    process_document (data)

    read_rows (reader)
```

```
class FloatListsFormatter (corpus)
    Bases: pimlico.cli.browser.formatter.DocumentBrowserFormatter

    DATATYPE
        alias of FloatListsDocumentType

    format_document (doc)
```

```
class FloatListsDocumentCorpus (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.tar.TarredCorpus

    Corpus of float list data: each doc contains lists of float. Unlike IntegerTableDocumentCorpus, they are
    not all constrained to have the same length. The downside is that the storage format (and probably I/O speed)
    isn't quite as efficient. It's still better than just storing ints as strings or JSON objects.
```

The floats are stored as C double, which use 8 bytes. At the moment, we don't provide any way to change this. An alternative would be to use C floats, losing precision but (almost) halving storage size.

```
datatype_name = 'float_lists_corpus'

data_point_type
    alias of FloatListsDocumentType
```

```
class FloatListsDocumentCorpusWriter (base_dir, **kwargs)
    Bases: pimlico.datatypes.tar.TarredCorpusWriter

    document_to_raw_data (data)
```

```
class FloatListDocumentType (options, metadata)
    Bases: pimlico.datatypes.documents.RawDocumentType

    Like FloatListsDocumentType, but each document is treated as a single list of floats.

    process_document (data)

    read_floats (reader)
```

```
class FloatListDocumentCorpus (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.tar.TarredCorpus

    Corpus of float data: each doc contains a single sequence of floats.

    The floats are stored as C doubles, using 8 bytes each.

    datatype_name = 'float_list_corpus'

    data_point_type
        alias of FloatListDocumentType
```

```
class FloatListDocumentCorpusWriter (base_dir, **kwargs)
    Bases: pimlico.datatypes.tar.TarredCorpusWriter

    document_to_raw_data (data)
```

## pimlico.datatypes.ints module

```
class IntegerListsDocumentType (options, metadata)
    Bases: pimlico.datatypes.documents.RawDocumentType

    unpacker

    process_document (data)

    read_rows (reader)
```

```
class IntegerListsDocumentCorpus (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.tar.TarredCorpus

    Corpus of integer list data: each doc contains lists of ints. Unlike IntegerTableDocumentCorpus, they
    are not all constrained to have the same length. The downside is that the storage format (and probably I/O speed)
    isn't quite as good. It's still better than just storing ints as strings or JSON objects.
```

By default, the ints are stored as C longs, which use 4 bytes. If you know you don't need ints this big, you can choose 1 or 2 bytes, or even 8 (long long). By default, the ints are unsigned, but they may be signed.

```
datatype_name = 'integer_lists_corpus'

data_point_type
    alias of IntegerListsDocumentType
```

```
class IntegerListsDocumentCorpusWriter (base_dir, signed=False, bytes=8, **kwargs)
    Bases: pimlico.datatypes.tar.TarredCorpusWriter

    document_to_raw_data (data)
```

```
class IntegerListDocumentType (options, metadata)
    Bases: pimlico.datatypes.documents.RawDocumentType

    Like IntegerListsDocumentType, but each document is treated as a single list of integers.

    unpacker

    int_size

    process_document (data)

    read_ints (reader)
```

```
class IntegerListDocumentCorpus (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.tar.TarredCorpus

    Corpus of integer data: each doc contains a single sequence of ints.

    By default, the ints are stored as C longs, which use 4 bytes. If you know you don't need ints this big, you can
    choose 1 or 2 bytes, or even 8 (long long). By default, the ints are unsigned, but they may be signed.
```

```
datatype_name = 'integer_list_corpus'

data_point_type
    alias of IntegerListDocumentType
```

```
class IntegerListDocumentCorpusWriter (base_dir, signed=False, bytes=8, **kwargs)
    Bases: pimlico.datatypes.tar.TarredCorpusWriter

    document_to_raw_data (data)
```

## pimlico.datatypes.jsondoc module

**class JsonDocumentCorpus** (*base\_dir, pipeline, \*\*kwargs*)

Bases: *pimlico.datatypes.tar.TarredCorpus*

Very simple document corpus in which each document is a JSON object.

**datatype\_name** = 'json'

**data\_point\_type**

alias of *JsonDocumentType*

**class JsonDocumentCorpusWriter** (*base\_dir, readable=False, \*\*kwargs*)

Bases: *pimlico.datatypes.tar.TarredCorpusWriter*

If *readable=True*, JSON text output will be nicely formatted so that it's human-readable. Otherwise, it will be formatted to take up less space.

**document\_to\_raw\_data** (*data*)

## pimlico.datatypes.keras module

Datatypes for storing and loading Keras models.

**class KerasModelWriter** (*base\_dir, \*\*kwargs*)

Bases: *pimlico.datatypes.base.PimlicoDatatypeWriter*

Writer for storing both types of Keras model (since they provide the same storage interface).

**write\_model** (*model*)

**write\_architecture** (*model*)

**write\_weights** (*model*)

**class KerasModel** (*base\_dir, pipeline, module=None, additional\_name=None, use\_main\_metadata=False, \*\*kwargs*)

Bases: *pimlico.datatypes.base.PimlicoDatatype*

Datatype for both types of Keras models, stored using Keras' own storage mechanisms. This uses Keras' method of storing the model architecture as JSON and stores the weights using hdf5.

**datatype\_name** = 'keras\_model'

**custom\_objects** = {}

**get\_software\_dependencies** ()

**get\_custom\_objects** ()

**load\_model** ()

**class KerasModelBuilderClassWriter** (*base\_dir, build\_params, builder\_class\_path, \*\*kwargs*)

Bases: *pimlico.datatypes.base.PimlicoDatatypeWriter*

Writer for storing Keras models in the manner described in :cls:KerasModelBuilderClass.

**write\_weights** (*model*)

**class KerasModelBuilderClass** (*base\_dir, pipeline, \*\*kwargs*)

Bases: *pimlico.datatypes.base.PimlicoDatatype*

An alternative way to store Keras models.

Create a class whose init method build the model architecture. It should take a kwarg called *build\_params*, which is a JSON-encodable dictionary of parameters that determine how the model gets build (hyperparameters). When you initialize your model for training, create this hyperparameter dictionary and use it to instantiate the model class.

Use the `KerasModelBuilderClassWriter` to store the model during training. Create a writer, then start model training, storing the weights to the filename given by the *weights\_filename* attribute of the writer. The hyperparameter dictionary will also be stored.

The writer also stores the fully-qualified path of the model-builder class. When we read the datatype and want to rebuild the model, we import the class, instantiate it and then set its weights to those we've stored.

The model builder class must have the model stored in an attribute *model*.

**weights\_filename**

**load\_build\_params()**

**create\_builder\_class** (*override\_params=None*)

**load\_model** (*override\_params=None*)

Instantiate the model builder class with the stored parameters and set the weights on the model to those stored.

**Returns** model builder instance (keras model in attribute *model*)

## pimlico.datatypes.plotting module

**class PlotOutput** (*base\_dir, pipeline, module=None, additional\_name=None, use\_main\_metadata=False, \*\*kwargs*)

Bases: `pimlico.datatypes.base.PimlicoDatatype`

Output from matplotlib plotting.

Contains the dataset being plotted, a script to build the plot, and the output PDF.

Also supplies additional datatypes to point to the individual files.

**supplied\_additional** = [('pdf', <class 'pimlico.datatypes.files.NamedFile'>), ('code', <class 'pimlico.datatypes.files.N

**script\_path**

**plot()**

Runs the plotting script. Errors are not caught, so if there's a problem in the script they'll be raised.

**pdf\_path**

**data\_path**

**class PlotOutputWriter** (*base\_dir*)

Bases: `pimlico.datatypes.base.PimlicoDatatypeWriter`

## pimlico.datatypes.r module

**class RTabSeparatedValuesFile** (*base\_dir, pipeline, module=None, additional\_name=None, use\_main\_metadata=False, \*\*kwargs*)

Bases: `pimlico.datatypes.files.File`

Tabular data stored in a TSV file, suitable for reading in using R's *read.delim* function.

**datatype\_name** = 'r\_tsv'



**absolute\_path**

**get\_detailed\_status** ()

**class RTabSeparatedValuesFileWriter** (*base\_dir*, *headings=None*, *\*\*kwargs*)

Bases: *pimlico.datatypes.base.PimlicoDatatypeWriter*

Writer for TSV files suitable for reading with R.

If *headings* is specified, this is written as the first line of the file, so *headings=TRUE* should be used when reading into R.

Double quotes (") in the fields will be replaced by double-double quotes (""), which R interprets as a double quote. Fields containing tabs will be surrounded by normal double quotes. When you read the data into R, the default value of *quotes* (") should therefore be fine. No escaping is performed on single quotes (').

**absolute\_path**

**write\_row** (*row*)

If elements are not of string type, they will be coerced to a string for writing. If you want to format them differently, do it before calling this method and pass in strings.

## pimlico.datatypes.results module

**class NumericResult** (*base\_dir*, *pipeline*, *\*\*kwargs*)

Bases: *pimlico.datatypes.base.PimlicoDatatype*

Simple datatype to contain a numeric value and a label, representing the result of some process, such as evaluation of a model on a task.

For example, allows results to be plotted by passing them into a graph plotting module.

**datatype\_name** = 'numeric\_result'

**data**

Raw JSON data

**result**

The numeric result being stored

**label**

A label to identify this result (e.g. model name)

**class NumericResultWriter** (*base\_dir*, *\*\*kwargs*)

Bases: *pimlico.datatypes.base.PimlicoDatatypeWriter*

## pimlico.datatypes.sklearn module

Datatypes for storing and loading Scikit-learn models.

**class SklearnModelWriter** (*base\_dir*, *\*\*kwargs*)

Bases: *pimlico.datatypes.base.PimlicoDatatypeWriter*

**write\_model** (*model*)

**class SklearnModel** (*base\_dir*, *pipeline*, *\*\*kwargs*)

Bases: *pimlico.datatypes.base.PimlicoDatatype*

Datatype for storing Scikit-learn models.

Very simple storage mechanism: we just pickle the model to a file. Instead of the standard Python pickle package, we use [Joblib](#), which stores large data objects (especially Numpy arrays) more efficiently.

```
get_software_dependencies ()  
  
load_model ()
```

### pimlico.datatypes.spans module

```
class SentenceSpansDocumentType (options, metadata)  
    Bases: pimlico.datatypes.jsondoc.JsonDocumentType  
  
    process_document (doc)  
  
class SentenceSpansCorpus (base_dir, pipeline, **kwargs)  
    Bases: pimlico.datatypes.jsondoc.JsonDocumentCorpus  
  
    data_point_type  
        alias of SentenceSpansDocumentType  
  
class SentenceSpansCorpusWriter (base_dir, readable=False, **kwargs)  
    Bases: pimlico.datatypes.jsondoc.JsonDocumentCorpusWriter  
  
    document_to_raw_data (data)
```

### pimlico.datatypes.table module

```
get_struct (bytes, signed, row_length)  
  
class IntegerTableDocumentType (options, metadata)  
    Bases: pimlico.datatypes.documents.RawDocumentType  
  
    unpacker  
  
    row_size  
  
    process_document (data)  
  
    read_rows (reader)  
  
class IntegerTableDocumentCorpus (base_dir, pipeline, **kwargs)  
    Bases: pimlico.datatypes.tar.TarredCorpus  
  
    Corpus of tabular integer data: each doc contains rows of ints, where each row contains the same number  
    of values. This allows a more compact representation, which doesn't require converting the ints to strings or  
    scanning for line ends, so is quite a bit quicker and results in much smaller file sizes. The downside is that the  
    files are not human-readable.  
  
    By default, the ints are stored as C longs, which use 4 bytes. If you know you don't need ints this big, you can  
    choose 1 or 2 bytes, or even 8 (long long). By default, the ints are unsigned, but they may be signed.  
  
    datatype_name = 'integer_table_corpus'  
  
    data_point_type  
        alias of IntegerTableDocumentType  
  
class IntegerTableDocumentCorpusWriter (base_dir, row_length, signed=False, bytes=8,  
                                         **kwargs)  
    Bases: pimlico.datatypes.tar.TarredCorpusWriter  
  
    document_to_raw_data (data)
```

**pimlico.datatypes.tar module**

**class TarredCorpus** (*base\_dir, pipeline, \*\*kwargs*)

Bases: *pimlico.datatypes.base.IterableCorpus*

**datatype\_name** = 'tar'

**document\_preprocessors** = []

**data\_point\_type**

alias of RawDocumentType

**extract\_file** (*archive\_name, filename*)

Extract an individual file by archive name and filename. This is not an efficient way of extracting a lot of files. The typical use case of a tarred corpus is to iterate over its files, which is much faster.

**doc\_iter** (*subsample=None, start\_after=None, skip=None*)

**archive\_iter** (*subsample=None, start\_after=None, skip=None*)

**process\_document** (*data*)

Process the data read in for a single document. Allows easy implementation of datatypes using TarredCorpus to do all the archive handling, etc, just specifying a particular way of handling the data within documents.

By default, uses the document data processing provided by the document type.

Most of the time, you shouldn't need to override this, but just write a document type that does the necessary processing.

**list\_archive\_iter** ()

**data\_ready** ()

**class TarredCorpusWriter** (*base\_dir, gzip=False, append=False, trust\_length=False, encoding='utf-8', \*\*kwargs*)

Bases: *pimlico.datatypes.base.IterableCorpusWriter*

If *gzip=True*, each document is gzipped before adding it to the archive. Not the same as creating a tarball, since the docs are gzipped *before* adding them, not the whole archive together, but it means we can easily iterate over the documents, unzipping them as required.

A subtlety of TarredCorpusWriter and its subclasses is that, as soon as the writer has been initialized, it must be legitimate to initialize a datatype to read the corpus. Naturally, at this point there will be no documents in the corpus, but it allows us to do document processing on the fly by initializing writers and readers to be sure the pre/post-processing is identical to if we were writing the docs to disk and reading them in again.

If *append=True*, existing archives and their files are not overwritten, the new files are just added to the end. This is useful where we want to restart processing that was broken off in the middle. If *trust\_length=True*, when appending the initial length of the corpus is read from the metadata already written. Otherwise (default), the number of docs already written is actually counted during initialization. This is sensible when the previous writing process may have ended abruptly, so that the metadata is not reliable. If you know you can trust the metadata, however, setting *trust\_length=True* will speed things up.

**add\_document** (*archive\_name, doc\_name, data*)

**document\_to\_raw\_data** (*doc*)

Overridden by subclasses to provide the mapping from the structured data supplied to the writer to the actual raw string to be written to disk. Override this instead of *add\_document()*, so that filters can do the mapping on the fly without writing the output to disk.

**class AlignedTarredCorpora** (*corpora*)

Bases: *object*

Iterator for iterating over multiple corpora simultaneously that contain the same files, grouped into archives in the same way. This is the standard utility for taking multiple inputs to a Pimlico module that contain different data but for the same corpus (e.g. output of different tools).

**archive\_iter** (*subsample=None, start\_after=None*)

**exception CorpusAlignmentError**

Bases: `exceptions.Exception`

**exception TarredCorpusIterationError**

Bases: `exceptions.Exception`

## pimlico.datatypes.tokenized module

**class TokenizedCorpus** (*base\_dir, pipeline, \*\*kwargs*)

Bases: `pimlico.datatypes.tar.TarredCorpus`

Specialized datatype for a tarred corpus that's had tokenization applied. The datatype does very little - the main reason for its existence is to allow modules to require that a corpus has been tokenized before it's given as input.

Each document is a list of sentences. Each sentence is a list of words.

**datatype\_name** = 'tokenized'

**data\_point\_type**

alias of `TokenizedDocumentType`

**class TokenizedCorpusWriter** (*base\_dir, gzip=False, append=False, trust\_length=False, encoding='utf-8', \*\*kwargs*)

Bases: `pimlico.datatypes.tar.TarredCorpusWriter`

Simple writer that takes lists of tokens and outputs them with a sentence per line and tokens separated by spaces.

**document\_to\_raw\_data** (*doc*)

## pimlico.datatypes.word2vec module

**class Word2VecModel** (*base\_dir, pipeline, \*\*kwargs*)

Bases: `pimlico.datatypes.base.PimlicoDatatype`

Datatype for storing Gensim-trained word2vec embeddings.

See also:

Datatype `pimlico.datatypes.embeddings.Embeddings` Another, more generic way, to write the same data, which should generally be used in preference to this one. `Embeddings` does not depend on Gensim, but can be converted to Gensim's data structure easily.

**shell\_commands** = [`<pimlico.datatypes.word2vec.NearestNeighboursCommand object>`, `<pimlico.datatypes.word2vec.`

**data\_ready** ()

**load\_model** ()

**model**

**get\_software\_dependencies** ()

---

```
class Word2VecModelWriter (base_dir, verb_only=False, **kwargs)
    Bases: pimlico.datatypes.base.PimlicoDatatypeWriter
```

---

**Note:** Generally, it's preferable to use *pimlico.datatypes.embeddings.Embeddings*, which is more generic, so easier to connect up with general vector/embedding-handling modules.

---

```
write_word2vec_model (model)
write_keyed_vectors (vectors)
```

## pimlico.datatypes.word\_annotations module

```
class WordAnnotationsDocumentType (options, metadata)
    Bases: pimlico.datatypes.documents.RawDocumentType
```

```
sentence_boundary_re
word_boundary
word_re
process_document (raw_data)
```

```
class WordAnnotationCorpus (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.tar.TarredCorpus
```

```
datatype_name = 'word_annotations'
data_point_type
    alias of WordAnnotationsDocumentType
annotation_fields = None
```

```
read_annotation_fields ()
    Get the available annotation fields from the dataset's configuration. These are the actual fields that will be
    available in the dictionary produced corresponding to each word.
```

```
data_ready ()
```

```
class WordAnnotationCorpusWriter (sentence_boundary, word_boundary, word_format, non-
    word_chars, base_dir, **kwargs)
    Bases: pimlico.datatypes.tar.TarredCorpusWriter
```

Ensures that the correct metadata is provided for a word annotation corpus. Doesn't take care of the formatting of the data: that needs to be done by the writing code, or by a subclass.

```
class SimpleWordAnnotationCorpusWriter (base_dir, field_names, field_sep=u'|', **kwargs)
    Bases: pimlico.datatypes.word_annotations.WordAnnotationCorpusWriter
```

Takes care of writing word annotations in a simple format, where each line contains a sentence, words are separated by spaces and a series of annotation fields for each word are separated by |s (or a given separator). This corresponds to the standard tag format for C&C.

```
document_to_raw_data (data)
```

```
class AddAnnotationField (input_name, add_fields)
    Bases: pimlico.datatypes.base.DynamicOutputDatatype
```

```
get_datatype (module_info)
classmethod get_base_datatype_class ()
```

**class WordAnnotationCorpusWithRequiredFields** (*required\_fields*)

Bases: *pimlico.datatypes.base.DynamicInputDatatypeRequirement*

Dynamic (functional) type that can be used in place of a module's input type. In typechecking, checks whether the input module is a WordAnnotationCorpus (or subtype) and whether its fields include all of those required.

**check\_type** (*supplied\_type*)

**exception AnnotationParseError**

Bases: *exceptions.Exception*

## pimlico.datatypes.xml module

Input module for extracting documents from XML files. Gigaword, for example, is stored in this way.

Depends on BeautifulSoup (see "bs4" target in lib dir Makefile).

**class XmlDocumentIterator** (*\*args, \*\*kwargs*)

Bases: *pimlico.datatypes.base.IterableCorpus*

**requires\_data\_preparation** = True

**input\_module\_options** = {'path': {'required': True, 'help': 'Path to the data'}, 'filter\_on\_doc\_attr': {'type': <functi

**data\_point\_type**

alias of *RawTextDocumentType*

**get\_software\_dependencies** ()

**prepare\_data** (*output\_dir, log*)

**data\_ready** ()

## Module contents

**OpenNLPCorefCorpus**

alias of *CorefCorpus*

**OpenNLPCorefCorpusWriter**

alias of *CorefCorpusWriter*

**CoreNLPCorefCorpus**

alias of *CorefCorpus*

**CoreNLPCorefCorpusWriter**

alias of *CorefCorpusWriter*

**class ConstituencyParseTreeCorpus** (*base\_dir, pipeline, \*\*kwargs*)

Bases: *pimlico.datatypes.tar.TarredCorpus*

Note that this is not fully developed yet. At the moment, you'll just get, for each document, a list of the texts of each tree. In future, they will be better represented.

**data\_point\_type**

alias of *TreeStringsDocumentType*

**datatype\_name** = 'parse\_trees'

**class ConstituencyParseTreeCorpusWriter** (*base\_dir, gzip=False, append=False, trust\_length=False, encoding='utf-8', \*\*kwargs*)

Bases: *pimlico.datatypes.tar.TarredCorpusWriter*

```

    document_to_raw_data (data)

class TreeStringsDocumentType (options, metadata)
    Bases: pimlico.datatypes.documents.RawDocumentType

    process_document (doc)

class CandcOutputCorpus (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.tar.TarredCorpus

    data_point_type
        alias of CandcOutputDocumentType

    datatype_name = 'candc_output'

class CandcOutputCorpusWriter (base_dir, gzip=False, append=False, trust_length=False,
                                encoding='utf-8', **kwargs)
    Bases: pimlico.datatypes.tar.TarredCorpusWriter

    document_to_raw_data (data)

class StanfordDependencyParseCorpus (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.jsondoc.JsonDocumentCorpus

    data_point_type
        alias of StanfordDependencyParseDocumentType

    datatype_name = 'stanford_dependency_pares'

class StanfordDependencyParseCorpusWriter (base_dir, readable=False, **kwargs)
    Bases: pimlico.datatypes.jsondoc.JsonDocumentCorpusWriter

    document_to_raw_data (data)

class CoNLLDependencyParseCorpus (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.word_annotations.WordAnnotationCorpus

    10-field CoNLL dependency parse format (conllx) – i.e. post parsing.

    Fields are: id (int), word form, lemma, coarse POS, POS, features, head (int), dep relation, phead (int), pdeprel

    The last two are usually not used.

    data_point_type
        alias of CoNLLDependencyParseDocumentType

    datatype_name = 'conll_dependency_pares'

class CoNLLDependencyParseCorpusWriter (base_dir, **kwargs)
    Bases: pimlico.datatypes.word_annotations.WordAnnotationCorpusWriter

    document_to_raw_data (data)

class CoNLLDependencyParseInputCorpus (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.word_annotations.WordAnnotationCorpus

    The version of the CoNLL format (conllx) that only has the first 6 columns, i.e. no dependency parse yet
    annotated.

    data_point_type
        alias of CoNLLDependencyParseInputDocumentType

    datatype_name = 'conll_dependency_parse_inputs'

class CoNLLDependencyParseInputCorpusWriter (base_dir, **kwargs)
    Bases: pimlico.datatypes.word_annotations.WordAnnotationCorpusWriter

```

```
document_to_raw_data (data)
```

```
class NumpyArray (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.files.NamedFileCollection

    array

    datatype_name = 'numpy_array'

    filenames = ['array.npy']

    get_software_dependencies ()
```

```
class NumpyArrayWriter (base_dir, additional_name=None)
    Bases: pimlico.datatypes.base.PimlicoDatatypeWriter

    set_array (array)
```

```
class ScipySparseMatrix (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.base.PimlicoDatatype

    Wrapper around Scipy sparse matrices. The matrix loaded is always in COO format – you probably want to
    convert to something else before using it. See scipy docs on sparse matrix conversions.

    array

    datatype_name = 'scipy_sparse_array'

    filenames = ['array.mtx']

    get_software_dependencies ()
```

```
class ScipySparseMatrixWriter (base_dir, additional_name=None)
    Bases: pimlico.datatypes.base.PimlicoDatatypeWriter

    set_matrix (mat)
```

```
class PimlicoDatatype (base_dir, pipeline, module=None, additional_name=None,
                      use_main_metadata=False, **kwargs)
    Bases: object

    The abstract superclass of all datatypes. Provides basic functionality for identifying where data should be stored
    and such.

    Datatypes are used to specify the routines for reading the output from modules. They're also used to specify
    how to read pipeline inputs. Most datatypes that have data simply read it in when required. Some (in particular
    those used as inputs) need a preparation phase to be run, where the raw data itself isn't sufficient to implement
    the reading interfaces required. In this case, they should override prepare_data().

    Datatypes may require/allow options to be set when they're used to read pipeline inputs. These are specified, in
    the same way as module options, by input_module_options on the datatype class.

    Datatypes may supply a set of additional datatypes. These should be guaranteed to be available if the main
    datatype is available. They must require no extra processing to be made available, unless that is done on the fly
    while reading the datatype (like a filter) or while the main datatype is being written.

    Additional datatypes can be accessed in config files by specifying the main datatype (as a previous module,
    optionally with an output name) and the additional datatype name in the form main_datatype->additional_name.
    Multiple additional names may be given, causing the next name to be looked up as an additional datatype of the
    initially loaded additional datatype. E.g main_datatype->additional0->additional1.

    To avoid conflicts in the metadata between datatypes using the same directory, datatypes loaded as additional
    datatypes have their additional name available to them and use it as a prefix to the metadata filename.

    If use_main_metadata=True on an additional datatype, the same metadata will be read as for the main datatype
    to which this is an additional datatype.
```



*module* is the ModuleInfo instance for the pipeline module that this datatype was produced by. It may be None, if the datatype wasn't instantiated by a module. It is not required to be set if you're instantiating a datatype in some context other than module output. It should generally be set for input datatypes, though, since they are treated as being created by a special input module.

**classmethod** `check_type` (*supplied\_type*)

Method used by datatype type-checking algorithm to determine whether a supplied datatype (given as a class, which is a subclass of PimlicoDatatype) is compatible with the present datatype, which is being treated as a type requirement.

Typically, the present class is a type requirement on a module input and *supplied\_type* is the type provided by a previous module's output.

The default implementation simply checks whether *supplied\_type* is a subclass of the present class. Subclasses may wish to impose different or additional checks.

**Parameters** `supplied_type` – type provided where the present class is required, or datatype instance

**Returns** True if the check is successful, False otherwise

**classmethod** `create_from_options` (*base\_dir*, *pipeline*, *options*={}, *module*=None)

**data\_ready** ()

Check whether the data corresponding to this datatype instance exists and is ready to be read.

Default implementation just checks whether the data dir exists. Subclasses might want to add their own checks, or even override this, if the data dir isn't needed.

**classmethod** `datatype_full_class_name` ()

The fully qualified name of the class for this datatype, by which it is reference in config files. Generally, datatypes don't need to override this, but type requirements that take the place of datatypes for type checking need to provide it.

**datatype\_name** = 'base\_datatype'

**emulated\_datatype** = None

**classmethod** `full_datatype_name` ()

Returns a string/unicode name for the datatype that includes relevant sub-type information. The default implementation just uses the attribute *datatype\_name*, but subclasses may have more detailed information to add. For example, iterable corpus types also supply information about the data-point type.

**get\_detailed\_status** ()

Returns a list of strings, containing detailed information about the data. Only called if `data_ready()` == True.

Subclasses may override this to supply useful (human-readable) information specific to the datatype. They should called the super method.

**get\_required\_paths** ()

Returns a list of paths to files that should be available for the data to be read. The base `data_ready()` implementation checks that these are all available and, if the datatype is used as an input to a pipeline and requires a data preparation routine to be run, data preparation will not be executed until these files are available.

Paths may be absolute or relative. If relative, they refer to files within the data directory and `data_ready()` will fail if the data dir doesn't exist.

**Returns** list of absolute or relative paths

**get\_software\_dependencies** ()

Check that all software required to read this datatype is installed and locatable. This is separate to metadata

config checks, so that you don't need to satisfy the dependencies for all modules in order to be able to run one of them. You might, for example, want to run different modules on different machines. This is called when a module is about to be executed and each of the dependencies is checked.

Returns a list of instances of subclasses of `:class:~pimlico.core.dependencies.base.SoftwareDependency`, representing the libraries that this module depends on.

Take care when providing dependency classes that you don't put any import statements at the top of the Python module that will make loading the dependency type itself dependent on runtime dependencies. You'll want to run import checks by putting import statements within this method.

You should call the super method for checking superclass dependencies.

**input\_module\_options** = {}

**instantiate\_additional\_datatype** (*name*, *additional\_name*)

Default implementation just assumes the datatype class can be instantiated using the default constructor, with the same base dir and pipeline as the main datatype. Options given to the main datatype are passed down to the additional datatype.

**metadata**

Read in metadata from a file in the corpus directory.

Note that this is no longer cached in memory. We need to be sure that the metadata values returned are always up to date with what is on disk, so always re-read the file when we need to get a value from the metadata. Since the file is typically small, this is unlikely to cause a problem. If we decide to return to cacheing the metadata dictionary in future, we will need to make sure that we can never run into problems with out-of-date metadata being returned.

**prepare\_data** (*output\_dir*, *log*)

**requires\_data\_preparation** = False

**shell\_commands** = []

**supplied\_additional** = []

**classmethod type\_checking\_name** ()

Supplies a name for this datatype to be used in type-checking error messages. Default implementation just provides the class name. Classes that override `check_supplied_type()` may want to override this too.

**class PimlicoDatatypeWriter** (*base\_dir*, *additional\_name=None*)

Bases: `object`

Abstract base class for data writer associated with Pimlico datatypes.

**incomplete\_tasks**

**require\_tasks** (*\*tasks*)

Add a name or multiple names to the list of output tasks that must be completed before writing is finished

**subordinate\_additional\_name** (*name*)

**task\_complete** (*task*)

**write\_metadata** ()

**class IterableCorpus** (*\*args*, *\*\*kwargs*)

Bases: `pimlico.datatypes.base.PimlicoDatatype`

Superclass of all datatypes which represent a dataset that can be iterated over document by document (or datapoint by datapoint - what exactly we're iterating over may vary, though documents are most common).

The actual type of the data depends on the subclass: it could be, e.g. coref output, etc. Information about the type of individual documents is provided by `document_type` and this is used in type checking.

At creation time, length should be provided in the metadata, denoting how many documents are in the dataset.

**classmethod** `check_supplied_type` (*supplied\_type*)

Override type checking to require that the supplied type have a document type that is compatible with (i.e. a subclass of) the document type of this class.

**data\_point\_type**

alias of *RawDocumentType*

**datatype\_name** = 'iterable\_corpus'

**classmethod** `full_datatype_name` ()

**get\_detailed\_status** ()

**process\_document\_data\_with\_datatype** (*data*)

Applies the corpus' datatype's `process_document()` method to the raw data :param data: :return:

**shell\_commands** = [*<pimlico.datatypes.base.CountInvalidCmd object>*]

**classmethod** `type_checking_name` ()

**class** `IterableCorpusWriter` (*base\_dir*, *additional\_name=None*)

Bases: *pimlico.datatypes.base.PimlicoDatatypeWriter*

**class** `DynamicOutputDatatype`

Bases: `object`

Types of module outputs may be specified as a subclass of *PimlicoDatatype*, or alternatively as an *instance* of `DynamicOutputType`. In this case, `get_datatype()` is called when the output datatype is needed, passing in the module info instance for the module, so that a specialized datatype can be produced on the basis of options, input types, etc.

The dynamic type must provide certain pieces of information needed for typechecking.

**datatype\_name** = `None`

**get\_base\_datatype\_class** ()

If it's possible to say before the instance of a `ModuleInfo` is available what base datatype will be produced, implement this to return the class. By default, it returns `None`.

If this information is available, it will be used in documentation.

**get\_datatype** (*module\_info*)

**class** `DynamicInputDatatypeRequirement`

Bases: `object`

Types of module inputs may be given as a subclass of *PimlicoDatatype*, a tuple of datatypes, or an instance a `DynamicInputDatatypeRequirement` subclass. In this case, `check_type(supplied_type)` is called during typechecking to check whether the type that we've got conforms to the input type requirements.

Additionally, if `datatype_doc_info` is provided, it is used to represent the input type constraints in documentation.

**check\_type** (*supplied\_type*)

**datatype\_doc\_info** = `None`

**type\_checking\_name** ()

Supplies a name for this datatype to be used in type-checking error messages. Default implementation just provides the class name. Subclasses may want to override this too.

**class** `InvalidDocument` (*module\_name*, *error\_info=None*)

Bases: `object`

Widely used in Pimlico to represent an empty document that is empty not because the original input document was empty, but because a module along the way had an error processing it. Document readers/writers should generally be robust to this and simply pass through the whole thing where possible, so that it's always possible to work out, where one of these pops up, where the error occurred.

**static invalid\_document\_or\_text** (*text*)

If the text represents an invalid document, parse it and return an InvalidDocument object. Otherwise, return the text as is.

**static load** (*text*)

**exception DatatypeLoadError**

Bases: `exceptions.Exception`

**exception DatatypeWriteError**

Bases: `exceptions.Exception`

**load\_datatype** (*path*)

Try loading a datatype class for a given path. Raises a DatatypeLoadError if it's not a valid datatype path.

**class MultipleInputs** (*datatype\_requirements*)

Bases: `object`

An input datatype that can be used as an item in a module's inputs, which lets the module accept an unbounded number of inputs, all satisfying the same datatype requirements. When writing the inputs in a config file, they can be specified as a comma-separated list of the usual type of specification (module name, with optional output name). Each item in the list must point to a datatype that satisfies the type-checking.

The list may also include (or entirely consist of) a base module name from the pipeline that has been expanded into multiple modules according to alternative parameters (the type separated by vertical bars, see [Multiple parameter values](#)). Use the notation *\*name*, where *name* is the base module name, to denote all of the expanded module names as inputs. These are treated as if you'd written out all of the expanded module names separated by commas.

When `get_input()` is called on the module, instead of returning a single datatype, a list of datatypes is returned.

**class CaervoCorpus** (*base\_dir*, *pipeline*, *\*\*kwargs*)

Bases: `pimlico.datatypes.tar.TarredCorpus`

Datatype for Caervo output. The output is stored exactly as it comes out from Caervo, in an XML format. This datatype reads in that XML and provides easy access to its components.

Since we simply store the XML that comes from Caervo, there's no corresponding corpus writer. The data is output using a `:class:pimlico.datatypes.tar.TarredCorpusWriter`.

**data\_point\_type**

alias of `CaervoDocumentType`

**class Dictionary** (*base\_dir*, *pipeline*, *\*\*kwargs*)

Bases: `pimlico.datatypes.base.PimlicoDatatype`

Dictionary encapsulates the mapping between normalized words and their integer ids.

**data\_ready** ()

**datatype\_name** = 'dictionary'

**get\_data** ()

**get\_detailed\_status** ()

**class DictionaryWriter** (*base\_dir*)

Bases: `pimlico.datatypes.base.PimlicoDatatypeWriter`

```

    add_documents (documents, prune_at=2000000)

    filter (threshold=None, no_above=None, limit=None)

class KeyValueCollection (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.tar.TarredCorpus

    data_point_type
        alias of KeyValueCollectionDocumentType

    datatype_name = 'key_value_lists'

class KeyValueCollectionWriter (base_dir, separator=',', fv_separator='=', **kwargs)
    Bases: pimlico.datatypes.tar.TarredCorpusWriter

    document_to_raw_data (data)

class KeyValueCollectionDocumentType (options, metadata)
    Bases: pimlico.datatypes.documents.RawDocumentType

    process_document (doc)

class TermFeatureListCorpus (base_dir, pipeline, **kwargs)
    Bases: pimlico.datatypes.features.KeyValueCollection

    Special case of KeyValueCollection, where one special feature “term” is always present and the other feature
    types are counts of the occurrence of a particular feature with this term in each data point.

    data_point_type
        alias of TermFeatureListDocumentType

    datatype_name = 'term_feature_lists'

class TermFeatureListCorpusWriter (base_dir, **kwargs)
    Bases: pimlico.datatypes.features.KeyValueCollectionWriter

    document_to_raw_data (data)

class TermFeatureListDocumentType (options, metadata)
    Bases: pimlico.datatypes.features.KeyValueCollectionDocumentType

    process_document (doc)

class IndexedTermFeatureListCorpus (*args, **kwargs)
    Bases: pimlico.datatypes.base.IterableCorpus

    Term-feature instances, indexed by a dictionary, so that all that’s stored is the indices of the terms and features
    and the feature counts for each instance. This is iterable, but, unlike TermFeatureListCorpus, doesn’t iterate
    over documents. Now that we’ve filtered extracted features down to a smaller vocab, we put everything in one
    big file, with one data point per line.

    Since we’re now storing indices, we can use a compact format that’s fast to read from disk, making iterating
    over the dataset faster than if we had to read strings, look them up in the vocab, etc.

    By default, the ints are stored as C longs, which use 4 bytes. If you know you don’t need ints this big, you can
    choose 1 or 2 bytes, or even 8 (long long). By default, the ints are unsigned, but they may be signed.

    data_point_type
        alias of IndexedTermFeatureListDataPointType

    feature_dictionary

    term_dictionary

```

```
class IndexedTermFeatureListCorpusWriter (base_dir, term_dictionary, feature_dictionary,
                                          bytes=4, signed=False, index_input=False,
                                          **kwargs)
```

Bases: *pimlico.datatypes.base.IterableCorpusWriter*

index\_input=True means that the input terms and feature names are already mapped to dictionary indices, so are assumed to be ints. Otherwise, inputs will be looked up in the appropriate dictionary to get an index.

**add\_data\_points** (*iterable*)

**write\_dictionaries** ()

```
class IndexedTermFeatureListDataPointType (options, metadata)
```

Bases: *pimlico.datatypes.documents.DataPointType*

```
class FeatureListScoreCorpus (base_dir, pipeline, **kwargs)
```

Bases: *pimlico.datatypes.tar.TarredCorpus*

**data\_point\_type**

alias of *FeatureListScoreDocumentType*

**datatype\_name** = 'scored\_weight\_feature\_lists'

```
class FeatureListScoreCorpusWriter (base_dir, features, separator=':', index_input=False,
                                    **kwargs)
```

Bases: *pimlico.datatypes.tar.TarredCorpusWriter*

Input should be a list of data points. Each is a (score, feature list) pair, where score is a Decimal, or other numeric type. Feature list is a list of (feature name, weight) pairs, or just feature names. If weights are not given, they will default to 1 when read in (but no weight is stored).

If index\_input=True, it is assumed that feature IDs will be given instead of feature names. Otherwise, the feature names will be looked up in the feature list. Any features not found in the feature type list will simply be skipped.

**document\_to\_raw\_data** (*data*)

```
class FeatureListScoreDocumentType (options, metadata)
```

Bases: *pimlico.datatypes.documents.RawDocumentType*

Document type that stores a list of features, each associated with a floating-point score. The feature lists are simply lists of indices to a feature set for the whole corpus that includes all feature types and which is stored along with the dataset. These may be binary features (present or absent for each data point), or may have a weight associated with them. If they are binary, the returned data will have a weight of 1 associated with each.

A corpus of this type can be used to train, for example, a regression.

If scores and weights are passed in as Decimal objects, they will be stored as strings. If they are floats, they will be converted to Decimals via their string representation (avoiding some of the oddness of converting between binary and decimal representations). To avoid loss of precision, pass in all scores and weights as Decimal objects.

**formatters** = [('features', 'pimlico.datatypes.formatters.features.FeatureListScoreFormatter')]

**process\_document** (*doc*)

```
class JsonDocumentCorpus (base_dir, pipeline, **kwargs)
```

Bases: *pimlico.datatypes.tar.TarredCorpus*

Very simple document corpus in which each document is a JSON object.

**data\_point\_type**

alias of *JsonDocumentType*

**datatype\_name** = 'json'

**class JsonDocumentCorpusWriter** (*base\_dir*, *readable=False*, *\*\*kwargs*)

Bases: *pimlico.datatypes.tar.TarredCorpusWriter*

If *readable=True*, JSON text output will be nicely formatted so that it's human-readable. Otherwise, it will be formatted to take up less space.

**document\_to\_raw\_data** (*data*)

**class TarredCorpus** (*base\_dir*, *pipeline*, *\*\*kwargs*)

Bases: *pimlico.datatypes.base.IterableCorpus*

**archive\_iter** (*subsample=None*, *start\_after=None*, *skip=None*)

**data\_point\_type**

alias of *RawDocumentType*

**data\_ready** ()

**datatype\_name** = 'tar'

**doc\_iter** (*subsample=None*, *start\_after=None*, *skip=None*)

**document\_preprocessors** = []

**extract\_file** (*archive\_name*, *filename*)

Extract an individual file by archive name and filename. This is not an efficient way of extracting a lot of files. The typical use case of a tarred corpus is to iterate over its files, which is much faster.

**list\_archive\_iter** ()

**process\_document** (*data*)

Process the data read in for a single document. Allows easy implementation of datatypes using TarredCorpus to do all the archive handling, etc, just specifying a particular way of handling the data within documents.

By default, uses the document data processing provided by the document type.

Most of the time, you shouldn't need to override this, but just write a document type that does the necessary processing.

**class TarredCorpusWriter** (*base\_dir*, *gzip=False*, *append=False*, *trust\_length=False*, *encoding='utf-8'*, *\*\*kwargs*)

Bases: *pimlico.datatypes.base.IterableCorpusWriter*

If *gzip=True*, each document is gzipped before adding it to the archive. Not the same as creating a tarball, since the docs are gzipped *before* adding them, not the whole archive together, but it means we can easily iterate over the documents, unzipping them as required.

A subtlety of TarredCorpusWriter and its subclasses is that, as soon as the writer has been initialized, it must be legitimate to initialize a datatype to read the corpus. Naturally, at this point there will be no documents in the corpus, but it allows us to do document processing on the fly by initializing writers and readers to be sure the pre/post-processing is identical to if we were writing the docs to disk and reading them in again.

If *append=True*, existing archives and their files are not overwritten, the new files are just added to the end. This is useful where we want to restart processing that was broken off in the middle. If *trust\_length=True*, when appending the initial length of the corpus is read from the metadata already written. Otherwise (default), the number of docs already written is actually counted during initialization. This is sensible when the previous writing process may have ended abruptly, so that the metadata is not reliable. If you know you can trust the metadata, however, setting *trust\_length=True* will speed things up.

**add\_document** (*archive\_name*, *doc\_name*, *data*)

**document\_to\_raw\_data** (*doc*)

Overridden by subclasses to provide the mapping from the structured data supplied to the writer to the

actual raw string to be written to disk. Override this instead of `add_document()`, so that filters can do the mapping on the fly without writing the output to disk.

**class `AlignedTarredCorpora`** (*corpora*)

Bases: `object`

Iterator for iterating over multiple corpora simultaneously that contain the same files, grouped into archives in the same way. This is the standard utility for taking multiple inputs to a Pimlico module that contain different data but for the same corpus (e.g. output of different tools).

**`archive_iter`** (*subsample=None, start\_after=None*)

**exception `CorpusAlignmentError`**

Bases: `exceptions.Exception`

**exception `TarredCorpusIterationError`**

Bases: `exceptions.Exception`

**class `TokenizedCorpus`** (*base\_dir, pipeline, \*\*kwargs*)

Bases: `pimlico.datatypes.tar.TarredCorpus`

Specialized datatype for a tarred corpus that's had tokenization applied. The datatype does very little - the main reason for its existence is to allow modules to require that a corpus has been tokenized before it's given as input.

Each document is a list of sentences. Each sentence is a list of words.

**`data_point_type`**

alias of `TokenizedDocumentType`

**`datatype_name`** = `'tokenized'`

**class `TokenizedCorpusWriter`** (*base\_dir, gzip=False, append=False, trust\_length=False, encoding='utf-8', \*\*kwargs*)

Bases: `pimlico.datatypes.tar.TarredCorpusWriter`

Simple writer that takes lists of tokens and outputs them with a sentence per line and tokens separated by spaces.

**`document_to_raw_data`** (*doc*)

**class `WordAnnotationCorpus`** (*base\_dir, pipeline, \*\*kwargs*)

Bases: `pimlico.datatypes.tar.TarredCorpus`

**`annotation_fields`** = `None`

**`data_point_type`**

alias of `WordAnnotationsDocumentType`

**`data_ready`** ()

**`datatype_name`** = `'word_annotations'`

**`read_annotation_fields`** ()

Get the available annotation fields from the dataset's configuration. These are the actual fields that will be available in the dictionary produced corresponding to each word.

**class `WordAnnotationCorpusWriter`** (*sentence\_boundary, word\_boundary, word\_format, non-word\_chars, base\_dir, \*\*kwargs*)

Bases: `pimlico.datatypes.tar.TarredCorpusWriter`

Ensures that the correct metadata is provided for a word annotation corpus. Doesn't take care of the formatting of the data: that needs to be done by the writing code, or by a subclass.

**class `SimpleWordAnnotationCorpusWriter`** (*base\_dir, field\_names, field\_sep=u'|', \*\*kwargs*)

Bases: `pimlico.datatypes.word_annotations.WordAnnotationCorpusWriter`



Takes care of writing word annotations in a simple format, where each line contains a sentence, words are separated by spaces and a series of annotation fields for each word are separated by |s (or a given separator). This corresponds to the standard tag format for C&C.

**document\_to\_raw\_data** (*data*)

**class AddAnnotationField** (*input\_name, add\_fields*)

Bases: *pimlico.datatypes.base.DynamicOutputDatatype*

**classmethod get\_base\_datatype\_class** ()

**get\_datatype** (*module\_info*)

**class WordAnnotationCorpusWithRequiredFields** (*required\_fields*)

Bases: *pimlico.datatypes.base.DynamicInputDatatypeRequirement*

Dynamic (functional) type that can be used in place of a module's input type. In typechecking, checks whether the input module is a WordAnnotationCorpus (or subtype) and whether its fields include all of those required.

**check\_type** (*supplied\_type*)

**exception AnnotationParseError**

Bases: *exceptions.Exception*

**class WordAnnotationsDocumentType** (*options, metadata*)

Bases: *pimlico.datatypes.documents.RawDocumentType*

**process\_document** (*raw\_data*)

**sentence\_boundary\_re**

**word\_boundary**

**word\_re**

**class XmlDocumentIterator** (*\*args, \*\*kwargs*)

Bases: *pimlico.datatypes.base.IterableCorpus*

**data\_point\_type**

alias of *RawTextDocumentType*

**data\_ready** ()

**get\_software\_dependencies** ()

**input\_module\_options** = {'path': {'required': True, 'help': 'Path to the data'}, 'filter\_on\_doc\_attr': {'type': <function>...}

**prepare\_data** (*output\_dir, log*)

**requires\_data\_preparation** = True

**class DataPointType** (*options, metadata*)

Bases: *object*

Base data-point type for iterable corpora. All iterable corpora should have data-point types that are subclasses of this.

**formatters** = []

**input\_module\_options** = {}

**class RawDocumentType** (*options, metadata*)

Bases: *pimlico.datatypes.documents.DataPointType*

Base document type. All document types for tarred corpora should be subclasses of this.

It may be used itself as well, where documents are just treated as raw data, though most of the time it will be appropriate to use subclasses to provide more information and processing operations specific to the datatype.

**process\_document** (*doc*)

**class RawTextDocumentType** (*options, metadata*)

Bases: *pimlico.datatypes.documents.RawDocumentType*

Subclass of RawDocumentType used to indicate that the document represents text (not just any old raw data) and that it hasn't been processed (tokenized, etc). Note that text that has been tokenized, parsed, etc does not used subclasses of this type, so they will not be considered compatible if this type is used as a requirement.

**input\_module\_options** = {'encoding': {'default': 'utf8', 'help': 'Encoding to assume for input files. Default: utf8'}}

**process\_document** (*doc*)

## pimlico.test package

### Submodules

#### pimlico.test.pipeline module

Pipeline tests

Pimlico modules and datatypes cannot always be easily tested with unit tests and where they can it's often not easy to work out how to write the tests in a neatly packaged way. Instead, modules can package up tests in the form of a small pipeline that comes with a tiny dataset to use as input. The pipeline can be run in a test environment, where software dependencies are installed and local config is prepared to store output and so on.

This way of providing tests also has the advantage that modules at the same time provide a demo (or several) of how to use them – how pipeline config should look and what sort of input data to use.

**class TestPipeline** (*pipeline, run\_modules, log*)

Bases: *object*

**static load\_pipeline** (*path, storage\_root*)

Load a test pipeline from a config file.

Path may be absolute, or given relative to Pimlico test data directory (PIMLICO\_ROOT/test/data)

**get\_uninstalled\_dependencies** ()

**test\_all\_modules** ()

**test\_input\_module** (*module\_name*)

**test\_module\_execution** (*module\_name*)

**run\_test\_pipeline** (*path, module\_names, log, no\_clean=False*)

Run a test pipeline, loading the pipeline config from a given path (which may be relative to the Pimlico test data directory) and running each of the named modules, including any of those modules' dependencies.

Any software dependencies not already available that can be installed automatically will be installed in the current environment. If there are unsatisfied dependencies that can't be automatically installed, an error will be raised.

If any of the modules name explicitly is an input dataset, it is loaded and data\_ready() is checked. If it is an IterableCorpus, it is tested simply by iterating over the full corpus.

**run\_test\_suite** (*pipelines\_and\_modules, log, no\_clean=False*)

Parameters **pipeline\_and\_modules** – list of (pipeline, modules) pairs, where pipeline is a path to a config file and modules a list of module names to test

**clear\_storage\_dir()**

**exception TestPipelineRunError**

Bases: `exceptions.Exception`

## Module contents

### pimlico.utils package

#### Subpackages

#### pimlico.utils.docs package

#### Submodules

#### pimlico.utils.docs.commandgen module

Tool to generate Pimlico command docs. Based on Sphinx's apidoc tool.

**generate\_docs** (*output\_dir*)

Generate RST docs for Pimlico commands and output to a directory.

**generate\_docs\_for\_command** (*command\_cls*, *output\_dir*)

**generate\_contents\_page** (*commands*, *command\_descs*, *output\_dir*)

**cap\_first** (*txt*)

**strip\_common\_indent** (*code*)

#### pimlico.utils.docs.modulegen module

Tool to generate Pimlico module docs. Based on Sphinx's apidoc tool.

**generate\_docs\_for\_pymod** (*module*, *output\_dir*)

Generate RST docs for Pimlico modules on a given Python path and output to a directory.

**generate\_docs\_for\_pimlico\_mod** (*module\_path*, *output\_dir*, *submodules*=[])

**input\_datatype\_list** (*types*)

**input\_datatype\_text** (*datatype*)

**output\_datatype\_text** (*datatype*)

**generate\_contents\_page** (*modules*, *output\_dir*, *index\_name*, *title*, *content*)

#### pimlico.utils.docs.rest module

**make\_table** (*grid*, *header*=None)

**table\_div** (*col\_widths*, *header\_flag*=False)

**normalize\_cell** (*string*, *length*)

## Module contents

**trim\_docstring** (*docstring*)

## Submodules

### pimlico.utils.communicate module

**timeout\_process** (*\*args, \*\*kws*)

Context manager for use in a *with* statement. If the with block hasn't completed after the given number of seconds, the process is killed.

**Parameters** *proc* – process to kill if timeout is reached before end of block

**Returns**

**terminate\_process** (*proc, kill\_time=None*)

Ends a process started with subprocess. Tries killing, then falls back on terminating if it doesn't work.

**Parameters**

- **kill\_time** – time to allow the process to be killed before falling back on terminating
- **proc** – Popen instance

**Returns**

**class StreamCommunicationPacket** (*data*)

Bases: `object`

**length**

**encode** ()

**static read** (*stream*)

**exception StreamCommunicationError**

Bases: `exceptions.Exception`

### pimlico.utils.core module

**multiwith** (*\*args, \*\*kws*)

Taken from contextlib's nested(). We need the variable number of context managers that this function allows.

**is\_identifier** (*ident*)

Determines if string is valid Python identifier.

**remove\_duplicates** (*lst, key=<function <lambda>>*)

Remove duplicate values from a list, keeping just the first one, using a particular key function to compare them.

**infinite\_cycle** (*iterable*)

Iterate infinitely over the given iterable.

Watch out for calling this on a generator or iter: they can only be iterated over once, so you'll get stuck in an infinite loop with no more items yielded once you've gone over it once.

You may also specify a callable, in which case it will be called each time to get a new iterable/iterator. This is useful in the case of generator functions.

**Parameters** *iterable* – iterable or generator to loop over indefinitely

**import\_member** (*path*)

Import a class, function, or other module member by its fully-qualified Python name.

**Parameters** *path* – path to member, including full package path and class/function/etc name

**Returns** *cls*

**split\_seq** (*seq, separator, ignore\_empty\_final=False*)

Iterate over a sequence and group its values into lists, separated in the original sequence by the given value. If *on* is callable, it is called on each element to test whether it is a separator. Otherwise, elements that are equal to *on* are treated as separators.

**Parameters**

- **seq** – sequence to divide up
- **separator** – separator or separator test function
- **ignore\_empty\_final** – by default, if there's a separator at the end, the last sequence yielded is empty. If `ignore_empty_final=True`, in this case the last empty sequence is dropped

**Returns** iterator over subsequences

**split\_seq\_after** (*seq, separator*)

Somewhat like `split_seq`, but starts a new subsequence after each separator, without removing the separators. Each subsequence therefore ends with a separator, except the last one if there's no separator at the end.

**Parameters**

- **seq** – sequence to divide up
- **separator** – separator or separator test function

**Returns** iterator over subsequences

**chunk\_list** (*lst, length*)

Divides a list into chunks of max *length* length.

**class cached\_property** (*func*)

Bases: `object`

A property that is only computed once per instance and then replaces itself with an ordinary attribute. Deleting the attribute resets the property.

Often useful in Pimlico datatypes, where it can be time-consuming to load data, but we can't do it once when the datatype is first loaded, since the data might not be ready at that point. Instead, we can access the data, or particular parts of it, using properties and easily cache the result.

Taken from: <https://github.com/bottlepy/bottle>

## pimlico.utils.email module

Email sending utilities

Configure email sending functionality by adding the following fields to your Pimlico local config file:

**email\_sender** From-address for all sent emails

**email\_recipients** To-addresses, separated by commas. All notification emails will be sent to all recipients

**email\_host** (optional) Hostname of your SMTP server. Defaults to *localhost*

*email\_username* (optional) Username to authenticate with your SMTP server. If not given, it is assumed that no authentication is required

*email\_password* (optional) Password to authenticate with your SMTP server. Must be supplied if *username* is given

**class EmailConfig** (*sender=None, recipients=None, host=None, username=None, password=None*)

Bases: `object`

**classmethod from\_local\_config** (*local\_config*)

**send\_pimlico\_email** (*subject, content, local\_config, log*)

Primary method for sending emails from Pimlico. Tries to send an email with the given content, using the email details found in the local config. If something goes wrong, an error is logged on the given log.

#### Parameters

- **subject** – email subject
- **content** – email text (may be unicode)
- **local\_config** – local config dictionary
- **log** – logger to log errors to (and info if the sending works)

**send\_text\_email** (*email\_config, subject, content=None*)

**exception EmailError**

Bases: `exceptions.Exception`

## pimlico.utils.filesystem module

**dirsize** (*path*)

Recursively compute the size of the contents of a directory.

**Parameters** *path* –

**Returns** size in bytes

**format\_file\_size** (*bytes*)

**copy\_dir\_with\_progress** (*source\_dir, target\_dir*)

Utility for copying a large directory and displaying a progress bar showing how much is copied.

#### Parameters

- **source\_dir** –
- **target\_dir** –

#### Returns

**new\_filename** (*directory, initial\_filename='tmp\_file'*)

Generate a filename that doesn't already exist.

**retry\_open** (*filename, errnos=[13], retry\_schedule=[2, 10, 30, 120, 300], \*\*kwargs*)

Try opening a file, using the builtin `open()` function. If an `IOError` is raised and its *errno* is in the given list, wait a moment then retry. Keeps doing this, waiting a bit longer each time, hoping that the problem will go away.

Once too many attempts have been made, outputs a message and waits for user input. This means the user can fix the problem (e.g. renew credentials) and pick up where execution left off. If they choose not to, the original error will be raised

Default list of *errnos* is just `[13]` – permission denied.

Use `retry_schedule` to customize the lengths of time waited between retries. Default: 2s, 10s, 30s, 2m, 5m, then give up.

Additional kwargs are pass on to `open()`.

**extract\_from\_archive** (*archive\_filename, members, target\_dir, preserve\_dirs=True*)

Extract a file or files from an archive, which may be a tarball or a zip file (determined by the file extension).

**extract\_archive** (*archive\_filename, target\_dir, preserve\_dirs=True*)

Extract all files from an archive, which may be a tarball or a zip file (determined by the file extension).

### pimlico.utils.format module

**multiline\_tablate** (*table, widths, \*\*kwargs*)

**title\_box** (*title\_text*)

Make a nice big pretty title surrounded by a box.

### pimlico.utils.linguistic module

**strip\_punctuation** (*s, split\_words=True*)

### pimlico.utils.logging module

**get\_console\_logger** (*name, debug=False*)

Convenience function to make it easier to create new loggers.

#### Parameters

- **name** – logging system logger name
- **debug** – whether to use DEBUG level. By default, uses INFO

#### Returns

### pimlico.utils.network module

**get\_unused\_local\_port** ()

Find a local port that's not currently being used, which we'll be able to bind a service to once this function returns.

**get\_unused\_local\_ports** (*n*)

Find a number of local ports not currently in use. Binds each port found before looking for the next one. If you just called `get_unused_local_port()` multiple times, you'd get to same answer coming back.

### pimlico.utils.pipes module

**qget** (*queue, \*args, \*\*kwargs*)

Wrapper that calls the `get()` method of a queue, catching `EINTR` interrupts and retrying. Recent versions of Python have this built in, but with earlier versions you can end up having processes die while waiting on queue output because an `EINTR` has received (which isn't necessarily a problem).

#### Parameters

- **queue** –
- **args** – args to pass to queue’s `get()`
- **kwargs** – kwargs to pass to queue’s `get()`

#### Returns

**class** `OutputQueue` (*out*)

Bases: `object`

Direct a readable output (e.g. pipe from a subprocess) to a queue. Returns the queue. Output is added to the queue one line at a time. To perform a non-blocking read call `get_nowait()` or `get(timeout=T)`

**get\_nowait** ()

**get** (*timeout=None*)

**get\_available** ()

Don’t block. Just return everything that’s available in the queue.

### pimlico.utils.pos module

**pos\_tag\_to\_ptb** (*tag*)

see :doc:pos\_pos\_tags\_to\_ptb

**pos\_tags\_to\_ptb** (*tags*)

Takes a list of POS tags and checks they’re all in the PTB tagset. If they’re not, tries mapping them according to CCGBank’s special version of the tagset. If that doesn’t work, raises a `NonPTBTagError`.

**exception** `NonPTBTagError`

Bases: `exceptions.Exception`

### pimlico.utils.probability module

**limited\_shuffle** (*iterable, buffer\_size, rand\_generator=None*)

Some algorithms require the order of data to be randomized. An obvious solution is to put it all in a list and shuffle, but if you don’t want to load it all into memory that’s not an option. This method iterates over the data, keeping a buffer and choosing at random from the buffer what to put next. It’s less shuffled than the simpler solution, but limits the amount of memory used at any one time to the buffer size.

**limited\_shuffle\_numpy** (*iterable, buffer\_size, randint\_buffer\_size=1000*)

Identical behaviour to `limited_shuffle()`, but uses Numpy’s random sampling routines to generate a large number of random integers at once. This can make execution a bit bursty, but overall tends to speed things up, as we get the random sampling over in one big call to Numpy.

**batched\_randint** (*low, high=None, batch\_size=1000*)

Infinite iterable that produces random numbers in the given range by calling Numpy now and then to generate lots of random numbers at once and then yielding them one by one. Faster than sampling one at a time.

#### Parameters

- **a** – lowest number in range
- **b** – highest number in range
- **batch\_size** – number of ints to generate in one go



**sequential\_document\_sample** (*corpus*, *start=None*, *shuffle=None*, *sample\_rate=None*)

Wrapper around a `pimlico.datatypes.tar.TarredCorpus` to draw infinite samples of documents from the corpus, by iterating over the corpus (looping infinitely), yielding documents at random. If *sample\_rate* is given, it should be a float between 0 and 1, specifying the rough proportion of documents to sample. A lower value spreads out the documents more on average.

Optionally, the samples are shuffled within a limited scope. Set *shuffle* to the size of this scope (higher will shuffle more, but need to buffer more samples in memory). Otherwise (*shuffle=0*), they will appear in the order they were in the original corpus.

If *start* is given, that number of documents will be skipped before drawing any samples. Set *start=0* to start at the beginning of the corpus. By default (*start=None*) a random point in the corpus will be skipped to before beginning.

**sequential\_sample** (*iterable*, *start=0*, *shuffle=None*, *sample\_rate=None*)

Draw infinite samples from an iterable, by iterating over it (looping infinitely), yielding items at random. If *sample\_rate* is given, it should be a float between 0 and 1, specifying the rough proportion of documents to sample. A lower value spreads out the documents more on average.

Optionally, the samples are shuffled within a limited scope. Set *shuffle* to the size of this scope (higher will shuffle more, but need to buffer more samples in memory). Otherwise (*shuffle=0*), they will appear in the order they were in the original corpus.

If *start* is given, that number of documents will be skipped before drawing any samples. Set *start=0* to start at the beginning of the corpus. Note that setting this to a high number can result in a slow start-up, if iterating over the items is slow.

---

**Note:** If you're sampling documents from a `TarredCorpus`, it's better to use `sequential_document_sample()`, since it makes use of `TarredCorpus`'s built-in features to do the skipping and sampling more efficiently.

---

**subsample** (*iterable*, *sample\_rate*)

Subsample the given iterable at a given rate, between 0 and 1.

## pimlico.utils.progress module

**get\_progress\_bar** (*maxval*, *counter=False*, *title=None*, *start=True*)

Simple utility to build a standard progress bar, so I don't have to think about this each time I need one. Starts the progress bar immediately.

*start* is no longer used, included only for backwards compatibility.

**class SafeProgressBar** (*maxval=None*, *widgets=None*, *term\_width=None*, *poll=1*, *left\_justify=True*, *fd=<open file '<stderr>', mode 'w'>*)

Bases: `progressbar.ProgressBar`

Override basic progress bar to wrap `update()` method with a couple of extra features.

1. You don't need to call `start()` – it will be called when the first update is received. This is good for processes that have a bit of a start-up lag, or where starting to iterate might generate some other output.
2. An error is not raised if you update with a value higher than *maxval*. It's the most annoying thing ever if you run a long process and the whole thing fails near the end because you slightly miscalculated *maxval*.

**update** (*value=None*)

**increment** ()

**slice\_progress** (*iterable*, *num\_items*, *title=None*)

**class ProgressBarIter** (*iterable, title=None*)  
Bases: `object`

### pimlico.utils.strings module

**truncate** (*s, length, ellipsis=u'...'*)

**similarities** (*targets, reference*)

Compute string similarity of each of a list of targets to a given reference string. Uses *difflib.SequenceMatcher* to compute similarity.

#### Parameters

- **reference** – compare all strings to this one
- **targets** – list of targets to measure similarity of

**Returns** list of similarity values

**sorted\_by\_similarity** (*targets, reference*)

Return target list sorted by similarity to the reference string. See `:func:similarities` for similarity measurement.

### pimlico.utils.system module

Lowish-level system operations

**set\_proc\_title** (*title*)

Tries to set the current process title. This is very system-dependent and may not always work.

If it's available, we use the *setproctitle* package, which is the most reliable way to do this. If not, we try doing it by loading *libc* and calling *prctl* ourselves. This is not reliable and only works on Unix systems. If neither of these works, we give up and return `False`.

If you want to increase the chances of this working (e.g. your process titles don't seem to be getting set by Pimlico and you'd like them to), try installing *setproctitle*, either system-wide or in Pimlico's virtualenv.

@return: True if the process succeeds, False if there's an error

### pimlico.utils.timeout module

**timeout** (*func, args=(), kwargs={}, timeout\_duration=1, default=None*)

### pimlico.utils.web module

**download\_file** (*url, target\_file*)

## Module contents

## Module contents

The Pimlico Processing Toolkit (PIpelled Modular LInguistic CORpus processing) is a toolkit for building pipelines made up of linguistic processing tasks to run on large datasets (corpora). It provides a wrappers around many existing, widely used NLP (Natural Language Processing) tools.

**install\_core\_dependencies** ()

## 1.6 Future plans

Various things I plan to add to Pimlico in the futures. For a summary, see [Pimlico Wishlist](#).

### 1.6.1 Pimlico Wishlist

Things I plan to add to Pimlico.

- Further modules:
  - [CherryPicker](#) for coreference resolution
  - [Berkeley Parser](#) for fast constituency parsing
  - [Reconcile](#) coref. Seems to incorporate upstream NLP tasks. Would want to interface such that we can reuse output from other modules and just do coref.
- **Pipeline graph visualizations:** [Outputting pipeline diagrams](#). Maybe an interactive GUI to help with viewing large pipelines
- See [issue list on Github](#) for other specific plans
- Big redesign of datatype implementation is [documented as a Github project](#)

### Todos

The following to-dos appear elsewhere in the docs. They are generally bits of the documentation I've not written yet, but am aware are needed.

---

#### Todo

Describe how module dependencies are defined for different types of deps

---

(The original entry is located in `/home/docs/checkouts/readthedocs.org/user_builds/pimlico/checkouts/v0.7/docs/core/dependencies.rst`, line 73.)

---

#### Todo

Include some examples from the core modules of how deps are defined and some special cases of software fetching

---

(The original entry is located in `/home/docs/checkouts/readthedocs.org/user_builds/pimlico/checkouts/v0.7/docs/core/dependencies.rst`, line 80.)

---

#### Todo

Write documentation for this

---

(The original entry is located in `/home/docs/checkouts/readthedocs.org/user_builds/pimlico/checkouts/v0.7/docs/core/module_structure`, line 9.)

---

#### Todo

Document variants

---

(The original entry is located in `/home/docs/checkouts/readthedocs.org/user_builds/pimlico/checkouts/v0.7/docs/core/variants.rst`, line 5.)

---

### **Todo**

Write a guide to building document map modules.

For now, the skeletons below are a useful starting point, but there should be a more fulsome explanation here of what document map modules are all about and how to use them.

---

(The original entry is located in `/home/docs/checkouts/readthedocs.org/user_builds/pimlico/checkouts/v0.7/docs/guides/map_module.rst`, line 5.)

---

### **Todo**

Use a dataset that everyone can get to in the example

---

(The original entry is located in `/home/docs/checkouts/readthedocs.org/user_builds/pimlico/checkouts/v0.7/docs/guides/setup.rst`, line 84.)

---

### **Todo**

Add more output convertors: currently only provides tokenization

---

(The original entry is located in `/home/docs/checkouts/readthedocs.org/user_builds/pimlico/checkouts/v0.7/docs/modules/pimlico.module.rst`, line 26.)

---

### **Todo**

Document this module

---

(The original entry is located in `/home/docs/checkouts/readthedocs.org/user_builds/pimlico/checkouts/v0.7/docs/modules/pimlico.module.rst`, line 12.)

---

### **Todo**

Document this module

---

(The original entry is located in `/home/docs/checkouts/readthedocs.org/user_builds/pimlico/checkouts/v0.7/docs/modules/pimlico.module.rst`, line 12.)

---

### **Todo**

Document this module

---

(The original entry is located in `/home/docs/checkouts/readthedocs.org/user_builds/pimlico/checkouts/v0.7/docs/modules/pimlico.module.rst`, line 12.)

---

### **Todo**

Document this module

---

(The original entry is located in /home/docs/checkouts/readthedocs.org/user\_builds/pimlico/checkouts/v0.7/docs/modules/pimlico.mod  
line 12.)

---

#### **Todo**

Document this module

---

(The original entry is located in /home/docs/checkouts/readthedocs.org/user\_builds/pimlico/checkouts/v0.7/docs/modules/pimlico.mod  
line 12.)

---

#### **Todo**

Document this module

---

(The original entry is located in /home/docs/checkouts/readthedocs.org/user\_builds/pimlico/checkouts/v0.7/docs/modules/pimlico.mod  
line 12.)

---

#### **Todo**

Replace check\_runtime\_dependencies() with get\_software\_dependencies()

---

(The original entry is located in /home/docs/checkouts/readthedocs.org/user\_builds/pimlico/checkouts/v0.7/docs/modules/pimlico.mod  
line 17.)

---

#### **Todo**

Document this module

---

(The original entry is located in /home/docs/checkouts/readthedocs.org/user\_builds/pimlico/checkouts/v0.7/docs/modules/pimlico.mod  
line 12.)

---

#### **Todo**

Replace check\_runtime\_dependencies() with get\_software\_dependencies()

---

(The original entry is located in /home/docs/checkouts/readthedocs.org/user\_builds/pimlico/checkouts/v0.7/docs/modules/pimlico.mod  
line 17.)

---

#### **Todo**

Document this module

---

(The original entry is located in /home/docs/checkouts/readthedocs.org/user\_builds/pimlico/checkouts/v0.7/docs/modules/pimlico.mod  
line 12.)

---

#### **Todo**

Document this module

---

(The original entry is located in /home/docs/checkouts/readthedocs.org/user\_builds/pimlico/checkouts/v0.7/docs/modules/pimlico.modul  
line 12.)

## 1.6.2 Berkeley Parser

<https://github.com/slavpetrov/berkeleyparser>

Java constituency parser. Pre-trained models are also provided in the Github repo.

Probably no need for a Java wrapper here. The parser itself accepts input on stdin and outputs to stdout, so just use a subprocess with pipes.

## 1.6.3 Cherry Picker

### Coreference resolver

<http://www.hlt.utdallas.edu/~altaf/cherrypicker/>

Requires NER, POS tagging and constituency parsing to be done first. Tools for all of these are included in the Cherry Picker codebase, but we just need a wrapper around the Cherry Picker tool itself to be able to feed these annotations in from other modules and perform coref.

Write a Java wrapper and interface with it using Py4J, as with OpenNLP.

## 1.6.4 Outputting pipeline diagrams

Once pipeline config files get big, it can be difficult to follow what's going on in them, especially if the structure is more complex than just a linear pipeline. A useful feature would be the ability to display/output a visualization of the pipeline as a flow graph.

It looks like the easiest way to do this will be to construct a DOT graph using Graphviz/Pydot and then output the diagram using Graphviz.

<http://www.graphviz.org>

<https://pypi.python.org/pypi/pydot>

Building the graph should be pretty straightforward, since the mapping from modules to nodes is fairly direct.

We could also add extra information to the nodes, like current execution status.

- genindex
- search

### C

- `pimlico.cli`, 83
- `pimlico.cli.browser`, 75
- `pimlico.cli.browser.formatter`, 73
- `pimlico.cli.browser.tool`, 74
- `pimlico.cli.check`, 77
- `pimlico.cli.clean`, 78
- `pimlico.cli.debug`, 75
- `pimlico.cli.debug.stepper`, 75
- `pimlico.cli.loaddump`, 78
- `pimlico.cli.locations`, 79
- `pimlico.cli.main`, 79
- `pimlico.cli.newmodule`, 80
- `pimlico.cli.pyshell`, 81
- `pimlico.cli.reset`, 81
- `pimlico.cli.run`, 81
- `pimlico.cli.shell`, 77
- `pimlico.cli.shell.base`, 76
- `pimlico.cli.shell.commands`, 76
- `pimlico.cli.shell.runner`, 77
- `pimlico.cli.status`, 82
- `pimlico.cli.subcommands`, 82
- `pimlico.cli.testemail`, 82
- `pimlico.cli.util`, 83
- `pimlico.core`, 111
- `pimlico.core.config`, 107
- `pimlico.core.dependencies`, 87
- `pimlico.core.dependencies.base`, 83
- `pimlico.core.dependencies.core`, 85
- `pimlico.core.dependencies.java`, 85
- `pimlico.core.dependencies.python`, 86
- `pimlico.core.dependencies.versions`, 87
- `pimlico.core.external`, 88
- `pimlico.core.external.java`, 87
- `pimlico.core.logs`, 110
- `pimlico.core.modules`, 106
- `pimlico.core.modules.base`, 95
- `pimlico.core.modules.execute`, 102
- `pimlico.core.modules.inputs`, 103

- `pimlico.core.modules.map`, 92
- `pimlico.core.modules.map.filter`, 89
- `pimlico.core.modules.map.multiproc`, 89
- `pimlico.core.modules.map.singleproc`, 91
- `pimlico.core.modules.map.threaded`, 91
- `pimlico.core.modules.multistage`, 104
- `pimlico.core.modules.options`, 106
- `pimlico.core.paths`, 111
- `pimlico.core.visualize`, 107
- `pimlico.core.visualize.deps`, 107
- `pimlico.core.visualize.status`, 107

### d

- `pimlico.datatypes`, 138
- `pimlico.datatypes.arrays`, 116
- `pimlico.datatypes.base`, 116
- `pimlico.datatypes.caevo`, 121
- `pimlico.datatypes.core`, 121
- `pimlico.datatypes.coref`, 112
- `pimlico.datatypes.coref.corenlp`, 111
- `pimlico.datatypes.coref.opennlp`, 111
- `pimlico.datatypes.dictionary`, 122
- `pimlico.datatypes.documents`, 122
- `pimlico.datatypes.embeddings`, 123
- `pimlico.datatypes.features`, 124
- `pimlico.datatypes.files`, 126
- `pimlico.datatypes.floats`, 129
- `pimlico.datatypes.formatters`, 113
- `pimlico.datatypes.formatters.features`, 113
- `pimlico.datatypes.formatters.tokenized`, 113
- `pimlico.datatypes.ints`, 130
- `pimlico.datatypes.jsondoc`, 131
- `pimlico.datatypes.keras`, 131
- `pimlico.datatypes.parse`, 114
- `pimlico.datatypes.parse.candc`, 113
- `pimlico.datatypes.parse.dependency`, 113
- `pimlico.datatypes.plotting`, 132
- `pimlico.datatypes.r`, 132

pimlico.datatypes.results, 133  
pimlico.datatypes.sklearn, 133  
pimlico.datatypes.spans, 134  
pimlico.datatypes.table, 134  
pimlico.datatypes.tar, 135  
pimlico.datatypes.tokenized, 136  
pimlico.datatypes.word2vec, 136  
pimlico.datatypes.word\_annotations, 137  
pimlico.datatypes.xml, 138

## m

pimlico.modules, 36  
pimlico.modules.caevo, 36  
pimlico.modules.caevo.output, 37  
pimlico.modules.candc, 38  
pimlico.modules.corenlp, 38  
pimlico.modules.corpora, 39  
pimlico.modules.corpora.concat, 39  
pimlico.modules.corpora.corpus\_stats, 40  
pimlico.modules.corpora.format, 40  
pimlico.modules.corpora.list\_filter, 41  
pimlico.modules.corpora.split, 41  
pimlico.modules.corpora.subset, 42  
pimlico.modules.corpora.tar, 43  
pimlico.modules.corpora.tar\_filter, 44  
pimlico.modules.corpora.vocab\_builder, 44  
pimlico.modules.corpora.vocab\_mapper, 45  
pimlico.modules.embeddings, 45  
pimlico.modules.embeddings.dependencies, 46  
pimlico.modules.embeddings.word2vec, 46  
pimlico.modules.features, 47  
pimlico.modules.features.term\_feature\_compiler, 47  
pimlico.modules.features.term\_feature\_matrix\_builder, 48  
pimlico.modules.features.vocab\_builder, 48  
pimlico.modules.features.vocab\_mapper, 49  
pimlico.modules.input, 49  
pimlico.modules.input.embeddings, 50  
pimlico.modules.input.embeddings.fasttext, 50  
pimlico.modules.input.embeddings.fasttext.gensim, 50  
pimlico.modules.input.text, 51  
pimlico.modules.input.text.raw\_text\_files, 51  
pimlico.modules.malt, 52

pimlico.modules.malt.conll\_parser\_input, 52  
pimlico.modules.malt.parse, 52  
pimlico.modules.opennlp, 53  
pimlico.modules.opennlp.coreference, 53  
pimlico.modules.opennlp.coreference\_pipeline, 54  
pimlico.modules.opennlp.ner, 55  
pimlico.modules.opennlp.parse, 56  
pimlico.modules.opennlp.pos, 56  
pimlico.modules.opennlp.tokenize, 57  
pimlico.modules.r, 57  
pimlico.modules.r.script, 58  
pimlico.modules.regex, 58  
pimlico.modules.regex.annotated\_text, 58  
pimlico.modules.sklearn, 59  
pimlico.modules.sklearn.matrix\_factorization, 59  
pimlico.modules.utility, 60  
pimlico.modules.utility.alias, 60  
pimlico.modules.utility.copy\_file, 61  
pimlico.modules.visualization, 61  
pimlico.modules.visualization.bar\_chart, 62  
pimlico.modules.visualization.embeddings\_plot, 62

## p

pimlico, 158

## t

pimlico.test, 151  
pimlico.test.pipeline, 150

## u

pimlico.utils, 158  
pimlico.utils.communicate, 152  
pimlico.utils.core, 152  
pimlico.utils.docs, 152  
pimlico.utils.docs.commandgen, 151  
pimlico.utils.docs.modulegen, 151  
pimlico.utils.docs.rest, 151  
pimlico.utils.email, 153  
pimlico.utils.filesystem, 154  
pimlico.utils.format, 155  
pimlico.utils.linguistic, 155  
pimlico.utils.logging, 155  
pimlico.utils.network, 155  
pimlico.utils.pipes, 155  
pimlico.utils.pos, 156  
pimlico.utils.probability, 156  
pimlico.utils.progress, 157  
pimlico.utils.strings, 158



`pimlico.utils.system`, [158](#)  
`pimlico.utils.timeout`, [158](#)  
`pimlico.utils.web`, [158](#)



## A

- `abs_path_or_model_dir_path()` (in module `pimlico.core.paths`), 111
- `absolute_path` (File attribute), 126
- `absolute_path` (NamedFileWriter attribute), 128
- `absolute_path` (RTabSeparatedValuesFile attribute), 132
- `absolute_path` (RTabSeparatedValuesFileWriter attribute), 133
- `add_arguments()` (BrowseCmd method), 80
- `add_arguments()` (DepsCmd method), 77
- `add_arguments()` (DumpCmd method), 78
- `add_arguments()` (InputsCmd method), 79
- `add_arguments()` (InstallCmd method), 77
- `add_arguments()` (LoadCmd method), 79
- `add_arguments()` (LongStoreCmd method), 79
- `add_arguments()` (OutputCmd method), 79
- `add_arguments()` (PimlicoCLISubcommand method), 82
- `add_arguments()` (PythonShellCmd method), 81
- `add_arguments()` (ResetCmd method), 81
- `add_arguments()` (RunCmd method), 81
- `add_arguments()` (ShellCLICmd method), 77
- `add_arguments()` (StatusCmd method), 82
- `add_arguments()` (UnlockCmd method), 80
- `add_arguments()` (VariantsCmd method), 79
- `add_arguments()` (VisualizeCmd method), 80
- `add_data_points()` (IndexedTermFeatureListCorpusWriter method), 125, 146
- `add_document()` (TarredCorpusWriter method), 135, 147
- `add_documents()` (DictionaryWriter method), 122, 144
- `add_execution_history_record()` (BaseModuleInfo method), 96
- `AddAnnotationField` (class in `pimlico.datatypes`), 149
- `AddAnnotationField` (class in `pimlico.datatypes.word_annotations`), 137
- `AlignedTarredCorpora` (class in `pimlico.datatypes`), 148
- `AlignedTarredCorpora` (class in `pimlico.datatypes.tar`), 135
- `all_dependencies()` (SoftwareDependency method), 84
- `all_inputs_ready()` (BaseModuleInfo method), 98
- `annotation_fields` (WordAnnotationCorpus attribute), 137, 148
- `AnnotationParseError`, 138, 149
- `append_module()` (PipelineConfig method), 108
- `archive_iter()` (AlignedTarredCorpora method), 136, 148
- `archive_iter()` (DocumentMapOutputTypeWrapper method), 89
- `archive_iter()` (TarredCorpus method), 135, 147
- `archive_iter_decorator()` (in module `pimlico.cli.debug.stepper`), 75
- `array` (NumpyArray attribute), 116, 140
- `array` (ScipySparseMatrix attribute), 116, 140
- `ask()` (in module `pimlico.cli.newmodule`), 80
- `available()` (SoftwareDependency method), 83

## B

- `BaseModuleExecutor` (class in `pimlico.core.modules.base`), 101
- `BaseModuleInfo` (class in `pimlico.core.modules.base`), 95
- `batched_randint()` (in module `pimlico.utils.probability`), 156
- `browse_cmd()` (in module `pimlico.cli.browser.tool`), 74
- `browse_data()` (in module `pimlico.cli.browser.tool`), 74
- `BrowseCmd` (class in `pimlico.cli.main`), 80
- `build_graph_with_status()` (in module `pimlico.core.visualize.status`), 107

## C

- `cached_property` (class in `pimlico.utils.core`), 153
- `CaevoCorpus` (class in `pimlico.datatypes`), 144
- `CaevoCorpus` (class in `pimlico.datatypes.caevo`), 121
- `call_java()` (in module `pimlico.core.external.java`), 87
- `CandcOutputCorpus` (class in `pimlico.datatypes`), 139
- `CandcOutputCorpus` (class in `pimlico.datatypes.parse`), 115
- `CandcOutputCorpus` (class in `pimlico.datatypes.parse.candc`), 113
- `CandcOutputCorpusWriter` (class in `pimlico.datatypes`), 139

CandcOutputCorpusWriter (class in pimlico.datatypes.parse), 115  
 CandcOutputCorpusWriter (class in pimlico.datatypes.parse.candc), 113  
 cap\_first() (in module pimlico.utils.docs.commandgen), 151  
 check\_and\_execute\_modules() (in module pimlico.core.modules.execute), 102  
 check\_and\_install() (in module pimlico.core.dependencies.base), 84  
 check\_for\_cycles() (in module pimlico.core.config), 110  
 check\_for\_error() (InputQueueFeeder method), 94  
 check\_invalid() (InputQueueFeeder method), 93  
 check\_java() (in module pimlico.core.dependencies.java), 86  
 check\_java\_dependency() (in module pimlico.core.dependencies.java), 85  
 check\_modules\_ready() (in module pimlico.core.modules.execute), 102  
 check\_pipeline() (in module pimlico.core.config), 110  
 check\_ready\_to\_run() (BaseModuleInfo method), 99  
 check\_ready\_to\_run() (MultistageModuleInfo method), 104  
 check\_release() (in module pimlico.core.config), 110  
 check\_supplied\_type() (pimlico.datatypes.base.IterableCorpus class method), 120  
 check\_supplied\_type() (pimlico.datatypes.IterableCorpus class method), 143  
 check\_type() (DynamicInputDatatypeRequirement method), 119, 143  
 check\_type() (in module pimlico.core.modules.base), 101  
 check\_type() (pimlico.datatypes.base.PimlicoDatatype class method), 118  
 check\_type() (pimlico.datatypes.PimlicoDatatype class method), 141  
 check\_type() (WordAnnotationCorpusWithRequiredFields method), 138, 149  
 choose\_from\_list() (in module pimlico.core.modules.options), 106  
 chunk\_list() (in module pimlico.utils.core), 153  
 CleanCmd (class in pimlico.cli.clean), 78  
 clear\_output\_queues() (Py4JInterface method), 88  
 clear\_storage\_dir() (in module pimlico.test.pipeline), 151  
 cmdloop() (DataShell method), 76  
 collect\_runnable\_modules() (in module pimlico.core.modules.base), 101  
 collect\_unexecuted\_dependencies() (in module pimlico.core.modules.base), 100  
 comma\_separated\_list() (in module pimlico.core.modules.options), 106  
 comma\_separated\_strings() (in module pimlico.core.modules.options), 106  
 command\_desc (CleanCmd attribute), 78  
 command\_desc (DumpCmd attribute), 78  
 command\_desc (InputsCmd attribute), 79  
 command\_desc (LoadCmd attribute), 79  
 command\_desc (LongStoreCmd attribute), 79  
 command\_desc (NewModuleCmd attribute), 80  
 command\_desc (PimlicoCLISubcommand attribute), 82  
 command\_desc (UnlockCmd attribute), 80  
 command\_desc (VisualizeCmd attribute), 80  
 command\_help (BrowseCmd attribute), 80  
 command\_help (CleanCmd attribute), 78  
 command\_help (DepsCmd attribute), 77  
 command\_help (DumpCmd attribute), 78  
 command\_help (EmailCmd attribute), 82  
 command\_help (InputsCmd attribute), 79  
 command\_help (InstallCmd attribute), 77  
 command\_help (LoadCmd attribute), 79  
 command\_help (LongStoreCmd attribute), 79  
 command\_help (NewModuleCmd attribute), 80  
 command\_help (OutputCmd attribute), 79  
 command\_help (PimlicoCLISubcommand attribute), 82  
 command\_help (PythonShellCmd attribute), 81  
 command\_help (ResetCmd attribute), 81  
 command\_help (RunCmd attribute), 81  
 command\_help (ShellCLICmd attribute), 77  
 command\_help (StatusCmd attribute), 82  
 command\_help (UnlockCmd attribute), 80  
 command\_help (VariantsCmd attribute), 79  
 command\_help (VisualizeCmd attribute), 80  
 command\_name (BrowseCmd attribute), 80  
 command\_name (CleanCmd attribute), 78  
 command\_name (DepsCmd attribute), 77  
 command\_name (DumpCmd attribute), 78  
 command\_name (EmailCmd attribute), 82  
 command\_name (InputsCmd attribute), 79  
 command\_name (InstallCmd attribute), 77  
 command\_name (LoadCmd attribute), 79  
 command\_name (LongStoreCmd attribute), 79  
 command\_name (NewModuleCmd attribute), 80  
 command\_name (OutputCmd attribute), 79  
 command\_name (PimlicoCLISubcommand attribute), 82  
 command\_name (PythonShellCmd attribute), 81  
 command\_name (ResetCmd attribute), 81  
 command\_name (RunCmd attribute), 81  
 command\_name (ShellCLICmd attribute), 77  
 command\_name (StatusCmd attribute), 82  
 command\_name (UnlockCmd attribute), 80  
 command\_name (VariantsCmd attribute), 79  
 command\_name (VisualizeCmd attribute), 80  
 commands (MetadataCmd attribute), 76  
 commands (PythonCmd attribute), 77  
 commands (ShellCommand attribute), 76  
 compare\_dotted\_versions() (in module pimlico.core.dependencies.versions), 87

- CoNLLEDependencyParseCorpus (class in pimlico.datatypes), 139
- CoNLLEDependencyParseCorpus (class in pimlico.datatypes.parse), 115
- CoNLLEDependencyParseCorpus (class in pimlico.datatypes.parse.dependency), 114
- CoNLLEDependencyParseCorpusWriter (class in pimlico.datatypes), 139
- CoNLLEDependencyParseCorpusWriter (class in pimlico.datatypes.parse), 115
- CoNLLEDependencyParseCorpusWriter (class in pimlico.datatypes.parse.dependency), 114
- CoNLLEDependencyParseInputCorpus (class in pimlico.datatypes), 139
- CoNLLEDependencyParseInputCorpus (class in pimlico.datatypes.parse), 115
- CoNLLEDependencyParseInputCorpus (class in pimlico.datatypes.parse.dependency), 114
- CoNLLEDependencyParseInputCorpusWriter (class in pimlico.datatypes), 139
- CoNLLEDependencyParseInputCorpusWriter (class in pimlico.datatypes.parse), 115
- CoNLLEDependencyParseInputCorpusWriter (class in pimlico.datatypes.parse.dependency), 114
- ConstituencyParseTreeCorpus (class in pimlico.datatypes), 138
- ConstituencyParseTreeCorpus (class in pimlico.datatypes.parse), 114
- ConstituencyParseTreeCorpusWriter (class in pimlico.datatypes), 138
- ConstituencyParseTreeCorpusWriter (class in pimlico.datatypes.parse), 114
- copy\_dir\_with\_progress() (in module pimlico.utils.filesystem), 154
- CORE\_PIMLICO\_DEPENDENCIES (in module pimlico.core.dependencies.core), 85
- CorefCorpus (class in pimlico.datatypes.coref.corenlp), 111
- CorefCorpus (class in pimlico.datatypes.coref.opennlp), 112
- CorefCorpusWriter (class in pimlico.datatypes.coref.corenlp), 111
- CorefCorpusWriter (class in pimlico.datatypes.coref.opennlp), 112
- CorefDocumentType (class in pimlico.datatypes.coref.corenlp), 111
- CorefDocumentType (class in pimlico.datatypes.coref.opennlp), 111
- CoreNLPCorefCorpus (in module pimlico.datatypes), 138
- CoreNLPCorefCorpus (in module pimlico.datatypes.coref), 112
- CoreNLPCorefCorpusWriter (in module pimlico.datatypes), 138
- CoreNLPCorefCorpusWriter (in module pimlico.datatypes.coref), 112
- CorpusAlignmentError, 136, 148
- CorpusState (class in pimlico.cli.browser.tool), 74
- create\_builder\_class() (KerasModelBuilderClass method), 132
- create\_from\_options() (pimlico.datatypes.base.PimlicoDatatype class method), 118
- create\_from\_options() (pimlico.datatypes.PimlicoDatatype class method), 141
- create\_pool() (DocumentMapModuleExecutor method), 93
- create\_pool() (MultiprocessingMapModuleExecutor method), 90
- create\_pool() (SingleThreadMapModuleExecutor method), 91
- create\_pool() (ThreadingMapModuleExecutor method), 92
- create\_pop\_up() (InputPopupLauncher method), 75
- create\_pop\_up() (MessagePopupLauncher method), 75
- create\_queue() (DocumentProcessorPool static method), 94
- create\_queue() (MultiprocessingMapPool static method), 90
- create\_queue() (ThreadingMapPool static method), 91
- custom\_objects (KerasModel attribute), 131
- ## D
- data (Dict attribute), 121
- data (NumericResult attribute), 133
- data (StringList attribute), 122
- data\_path (PlotOutput attribute), 132
- data\_point\_type (CaevoCorpus attribute), 121, 144
- data\_point\_type (CandcOutputCorpus attribute), 113, 115, 139
- data\_point\_type (CoNLLEDependencyParseCorpus attribute), 114, 115, 139
- data\_point\_type (CoNLLEDependencyParseInputCorpus attribute), 114, 115, 139
- data\_point\_type (ConstituencyParseTreeCorpus attribute), 114, 138
- data\_point\_type (CorefCorpus attribute), 111, 112
- data\_point\_type (FeatureListScoreCorpus attribute), 125, 146
- data\_point\_type (FloatListDocumentCorpus attribute), 129
- data\_point\_type (FloatListsDocumentCorpus attribute), 129
- data\_point\_type (IndexedTermFeatureListCorpus attribute), 125, 145
- data\_point\_type (IntegerListDocumentCorpus attribute), 130

- `data_point_type` (IntegerListsDocumentCorpus attribute), 130
- `data_point_type` (IntegerTableDocumentCorpus attribute), 134
- `data_point_type` (IterableCorpus attribute), 120, 143
- `data_point_type` (JsonDocumentCorpus attribute), 131, 146
- `data_point_type` (KeyValueListCorpus attribute), 124, 145
- `data_point_type` (RawTextDirectory attribute), 128
- `data_point_type` (RawTextFiles attribute), 128
- `data_point_type` (ReaderOutputType attribute), 103
- `data_point_type` (SentenceSpansCorpus attribute), 134
- `data_point_type` (StanfordDependencyParseCorpus attribute), 113, 115, 139
- `data_point_type` (TarredCorpus attribute), 135, 147
- `data_point_type` (TermFeatureListCorpus attribute), 124, 145
- `data_point_type` (TokenizedCorpus attribute), 136, 148
- `data_point_type` (WordAnnotationCorpus attribute), 137, 148
- `data_point_type` (XmlDocumentIterator attribute), 138, 149
- `data_ready()` (Dictionary method), 122, 144
- `data_ready()` (DocumentMapOutputTypeWrapper method), 89
- `data_ready()` (File method), 126
- `data_ready()` (NamedFileCollection method), 126
- `data_ready()` (PimlicoDatatype method), 118, 141
- `data_ready()` (ReaderOutputType method), 104
- `data_ready()` (StringList method), 122
- `data_ready()` (TarredCorpus method), 135, 147
- `data_ready()` (UnnamedFileCollection method), 127
- `data_ready()` (Word2VecModel method), 136
- `data_ready()` (WordAnnotationCorpus method), 137, 148
- `data_ready()` (XmlDocumentIterator method), 138, 149
- `DataPointType` (class in `pimlico.datatypes`), 149
- `DataPointType` (class in `pimlico.datatypes.documents`), 122
- `DataShell` (class in `pimlico.cli.shell.base`), 76
- `DATATYPE` (DefaultFormatter attribute), 73
- `DATATYPE` (DocumentBrowserFormatter attribute), 73
- `DATATYPE` (FeatureListScoreFormatter attribute), 113
- `DATATYPE` (FloatListsFormatter attribute), 129
- `DATATYPE` (InvalidDocumentFormatter attribute), 74
- `DATATYPE` (TokenizedDocumentFormatter attribute), 113
- `datatype_doc_info` (DynamicInputDatatypeRequirement attribute), 119, 143
- `datatype_full_class_name()` (`pimlico.datatypes.base.PimlicoDatatype` class method), 118
- `datatype_full_class_name()` (`pimlico.datatypes.PimlicoDatatype` class method), 141
- `datatype_name` (CandcOutputCorpus attribute), 113, 115, 139
- `datatype_name` (CoNLLDependencyParseCorpus attribute), 114, 115, 139
- `datatype_name` (CoNLLDependencyParseInputCorpus attribute), 114, 115, 139
- `datatype_name` (ConstituencyParseTreeCorpus attribute), 114, 138
- `datatype_name` (CorefCorpus attribute), 111, 112
- `datatype_name` (Dict attribute), 121
- `datatype_name` (Dictionary attribute), 122, 144
- `datatype_name` (DynamicOutputDatatype attribute), 119, 143
- `datatype_name` (FeatureListScoreCorpus attribute), 125, 146
- `datatype_name` (File attribute), 126
- `datatype_name` (FloatListDocumentCorpus attribute), 129
- `datatype_name` (FloatListsDocumentCorpus attribute), 129
- `datatype_name` (IntegerListDocumentCorpus attribute), 130
- `datatype_name` (IntegerListsDocumentCorpus attribute), 130
- `datatype_name` (IntegerTableDocumentCorpus attribute), 134
- `datatype_name` (IterableCorpus attribute), 120, 143
- `datatype_name` (JsonDocumentCorpus attribute), 131, 146
- `datatype_name` (KerasModel attribute), 131
- `datatype_name` (KeyValueListCorpus attribute), 124, 145
- `datatype_name` (NamedFileCollection attribute), 126
- `datatype_name` (NumericResult attribute), 133
- `datatype_name` (NumpyArray attribute), 116, 140
- `datatype_name` (PimlicoDatatype attribute), 117, 141
- `datatype_name` (RawTextDirectory attribute), 128
- `datatype_name` (ReaderOutputType attribute), 103
- `datatype_name` (RTabSeparatedValuesFile attribute), 132
- `datatype_name` (ScipySparseMatrix attribute), 116, 140
- `datatype_name` (SingleTextDocument attribute), 121
- `datatype_name` (StanfordDependencyParseCorpus attribute), 113, 115, 139
- `datatype_name` (StringList attribute), 122
- `datatype_name` (TarredCorpus attribute), 135, 147
- `datatype_name` (TermFeatureListCorpus attribute), 124, 145
- `datatype_name` (TokenizedCorpus attribute), 136, 148
- `datatype_name` (UnnamedFileCollection attribute), 127
- `datatype_name` (WordAnnotationCorpus attribute), 137, 148
- `DatatypeLoadError`, 121, 144
- `DatatypeWriteError`, 121, 144
- `default()` (`DataShell` method), 76

- DefaultFormatter (class in pimlico.cli.browser.formatter), 73
- dependencies (BaseModuleInfo attribute), 99
- dependencies() (Py4JSoftwareDependency method), 86
- dependencies() (SoftwareDependency method), 84
- DependencyCheckerError, 88
- DependencyError, 101
- DepsCmd (class in pimlico.cli.check), 77
- Dict (class in pimlico.datatypes.core), 121
- Dictionary (class in pimlico.datatypes), 144
- Dictionary (class in pimlico.datatypes.dictionary), 122
- DictionaryWriter (class in pimlico.datatypes), 144
- DictionaryWriter (class in pimlico.datatypes.dictionary), 122
- DictWriter (class in pimlico.datatypes.core), 121
- dirsize() (in module pimlico.utils.filesystem), 154
- do\_EOF() (DataShell method), 76
- doc\_iter() (TarredCorpus method), 135, 147
- document\_preprocessors (TarredCorpus attribute), 135, 147
- document\_to\_raw\_data() (CandcOutputCorpusWriter method), 113, 115, 139
- document\_to\_raw\_data() (CoNLLDependencyParseCorpusWriter method), 114, 115, 139
- document\_to\_raw\_data() (CoNLLDependencyParseInputCorpusWriter method), 114, 115, 139
- document\_to\_raw\_data() (ConstituencyParseTreeCorpusWriter method), 115, 138
- document\_to\_raw\_data() (CorefCorpusWriter method), 111, 112
- document\_to\_raw\_data() (FeatureListScoreCorpusWriter method), 126, 146
- document\_to\_raw\_data() (FloatListDocumentCorpusWriter method), 129
- document\_to\_raw\_data() (FloatListsDocumentCorpusWriter method), 129
- document\_to\_raw\_data() (IntegerListDocumentCorpusWriter method), 130
- document\_to\_raw\_data() (IntegerListsDocumentCorpusWriter method), 130
- document\_to\_raw\_data() (IntegerTableDocumentCorpusWriter method), 134
- document\_to\_raw\_data() (JsonDocumentCorpusWriter method), 131, 147
- document\_to\_raw\_data() (KeyValueListCorpusWriter method), 124, 145
- document\_to\_raw\_data() (SentenceSpansCorpusWriter method), 134
- document\_to\_raw\_data() (SimpleWordAnnotationCorpusWriter method), 137, 149
- document\_to\_raw\_data() (StanfordDependencyParseCorpusWriter method), 114, 115, 139
- document\_to\_raw\_data() (TarredCorpusWriter method), 135, 147
- document\_to\_raw\_data() (TermFeatureListCorpusWriter method), 124, 145
- document\_to\_raw\_data() (TokenizedCorpusWriter method), 136, 148
- DocumentBrowserFormatter (class in pimlico.cli.browser.formatter), 73
- DocumentMapModuleExecutor (class in pimlico.core.modules.map), 92
- DocumentMapModuleInfo (class in pimlico.core.modules.map), 92
- DocumentMapOutputTypeWrapper (class in pimlico.core.modules.map.filter), 89
- DocumentMapProcessMixin (class in pimlico.core.modules.map), 94
- DocumentProcessorPool (class in pimlico.core.modules.map), 94
- download\_file() (in module pimlico.utils.web), 158
- DumpCmd (class in pimlico.cli.loadump), 78
- DynamicInputDatatypeRequirement (class in pimlico.datatypes), 143
- DynamicInputDatatypeRequirement (class in pimlico.datatypes.base), 119
- DynamicOutputDatatype (class in pimlico.datatypes), 143
- DynamicOutputDatatype (class in pimlico.datatypes.base), 119
- ## E
- EmailCmd (class in pimlico.cli.testemail), 82
- EmailConfig (class in pimlico.utils.email), 154
- EmailError, 154
- Embeddings (class in pimlico.datatypes.embeddings), 123
- EmbeddingsWriter (class in pimlico.datatypes.embeddings), 123
- empty() (PipelineConfig static method), 108
- empty\_all\_queues() (DocumentProcessorPool method), 94
- empty\_all\_queues() (MultiprocessingMapPool method), 90
- emptyline() (DataShell method), 76
- emulated\_datatype (PimlicoDatatype attribute), 117, 141
- emulated\_datatype (ReaderOutputType attribute), 104
- enable\_step() (PipelineConfig method), 109
- enable\_step\_for\_pipeline() (in module pimlico.cli.debug.stepper), 75
- encode() (StreamCommunicationPacket method), 152
- Entity (class in pimlico.datatypes.coref.corenlp), 111
- Entity (class in pimlico.datatypes.coref.opennlp), 112
- execute() (BaseModuleExecutor method), 101
- execute() (DocumentMapModuleExecutor method), 93
- execute() (MetadataCmd method), 77
- execute() (PythonCmd method), 77
- execute() (ShellCommand method), 76



- pimlico.core.modules.execute), 102
- 
- execution\_history (BaseModuleInfo attribute), 96
- 
- execution\_history\_path (BaseModuleInfo attribute), 96
- 
- extract\_archive() (in module pimlico.utils.filesystem), 155
- 
- extract\_file() (TarredCorpus method), 135, 147
- 
- extract\_from\_archive() (in module pimlico.utils.filesystem), 155
- 
- extract\_input\_options() (pimlico.core.modules.base.BaseModuleInfo class method), 96
- 
- ## F
- feature\_dictionary (IndexedTermFeatureListCorpus attribute), 125, 145
- 
- FeatureListScoreCorpus (class in pimlico.datatypes), 146
- 
- FeatureListScoreCorpus (class in pimlico.datatypes.features), 125
- 
- FeatureListScoreCorpusWriter (class in pimlico.datatypes), 146
- 
- FeatureListScoreCorpusWriter (class in pimlico.datatypes.features), 125
- 
- FeatureListScoreDocumentType (class in pimlico.datatypes), 146
- 
- FeatureListScoreDocumentType (class in pimlico.datatypes.features), 125
- 
- FeatureListScoreFormatter (class in pimlico.datatypes.formatters.features), 113
- 
- File (class in pimlico.datatypes.files), 126
- 
- filename\_with\_range() (in module pimlico.datatypes.files), 127
- 
- filenames (Dict attribute), 121
- 
- filenames (NamedFileCollection attribute), 126
- 
- filenames (NamedFileCollectionWriter attribute), 126
- 
- filenames (NumpyArray attribute), 116, 140
- 
- filenames (ScipySparseMatrix attribute), 116, 140
- 
- filenames (SingleTextDocument attribute), 121
- 
- filter() (DictionaryWriter method), 122, 145
- 
- filter\_document() (DocumentBrowserFormatter method), 73
- 
- filter\_document() (InvalidDocumentFormatter method), 74
- 
- filter\_document() (RawTextDirectory method), 128
- 
- find\_all\_data\_paths() (PipelineConfig method), 109
- 
- find\_data\_path() (PipelineConfig method), 109
- 
- FloatListDocumentCorpus (class in pimlico.datatypes.floats), 129
- 
- FloatListDocumentCorpusWriter (class in pimlico.datatypes.floats), 129
- 
- FloatListDocumentType (class in pimlico.datatypes.floats), 129
- 
- FloatListsDocumentCorpusWriter (class in pimlico.datatypes.floats), 129
- 
- FloatListsDocumentType (class in pimlico.datatypes.floats), 129
- 
- FloatListsFormatter (class in pimlico.datatypes.floats), 129
- 
- fmt\_frame\_info() (in module pimlico.cli.debug), 76
- 
- format\_document() (DefaultFormatter method), 74
- 
- format\_document() (DocumentBrowserFormatter method), 73
- 
- format\_document() (FeatureListScoreFormatter method), 113
- 
- format\_document() (FloatListsFormatter method), 129
- 
- format\_document() (InvalidDocumentFormatter method), 74
- 
- format\_document() (TokenizedDocumentFormatter method), 113
- 
- format\_execution\_dependency\_tree() (in module pimlico.core.modules.execute), 102
- 
- format\_execution\_error() (in module pimlico.cli.util), 83
- 
- format\_file\_size() (in module pimlico.utils.filesystem), 154
- 
- format\_option\_type() (in module pimlico.core.modules.options), 106
- 
- formatters (DataPointType attribute), 122, 149
- 
- formatters (FeatureListScoreDocumentType attribute), 125, 146
- 
- formatters (FloatListsDocumentType attribute), 129
- 
- from\_java\_object() (Entity static method), 112
- 
- from\_java\_object() (Mention static method), 112
- 
- from\_json() (Entity static method), 112
- 
- from\_json() (Mention static method), 111, 112
- 
- from\_local\_config() (pimlico.utils.email.EmailConfig class method), 154
- 
- full\_datatype\_name() (pimlico.datatypes.base.IterableCorpus class method), 120
- 
- full\_datatype\_name() (pimlico.datatypes.base.PimlicoDatatype class method), 119
- 
- full\_datatype\_name() (pimlico.datatypes.IterableCorpus class method), 143
- 
- full\_datatype\_name() (pimlico.datatypes.PimlicoDatatype class method), 141
- 
- ## G
- gateway\_client\_to\_running\_server() (in module pimlico.core.external.java), 88
- 
- generate\_contents\_page() (in module pimlico.utils.docs.commandgen), 151
- 
- generate\_contents\_page() (in module pimlico.utils.docs.modulegen), 151



- generate\_docs() (in module pimlico.utils.docs.commandgen), 151
- generate\_docs\_for\_command() (in module pimlico.utils.docs.commandgen), 151
- generate\_docs\_for\_pimlico\_mod() (in module pimlico.utils.docs.modulegen), 151
- generate\_docs\_for\_pymod() (in module pimlico.utils.docs.modulegen), 151
- get() (OutputQueue method), 156
- get\_absolute\_output\_dir() (BaseModuleInfo method), 97
- get\_absolute\_path() (NamedFileCollection method), 126
- get\_absolute\_path() (NamedFileCollectionWriter method), 126
- get\_all\_executed\_modules() (BaseModuleInfo method), 100
- get\_available() (OutputQueue method), 156
- get\_base\_datatype\_class() (DynamicOutputDatatype method), 119, 143
- get\_base\_datatype\_class() (pimlico.datatypes.AddAnnotationField class method), 149
- get\_base\_datatype\_class() (pimlico.datatypes.word\_annotations.AddAnnotationField class method), 137
- get\_classpath() (in module pimlico.core.dependencies.java), 86
- get\_classpath\_components() (JavaDependency method), 85
- get\_console\_logger() (in module pimlico.utils.logging), 155
- get\_custom\_objects() (KerasModel method), 131
- get\_data() (Dictionary method), 122, 144
- get\_data\_search\_paths() (PipelineConfig method), 109
- get\_datatype() (AddAnnotationField method), 137, 149
- get\_datatype() (DynamicOutputDatatype method), 119, 143
- get\_dependencies() (in module pimlico.core.config), 110
- get\_dependent\_modules() (PipelineConfig method), 108
- get\_detailed\_status() (BaseModuleInfo method), 100
- get\_detailed\_status() (Dictionary method), 122, 144
- get\_detailed\_status() (DocumentMapModuleInfo method), 92
- get\_detailed\_status() (IterableCorpus method), 120, 143
- get\_detailed\_status() (MultistageModuleInfo method), 104
- get\_detailed\_status() (PimlicoDatatype method), 118, 141
- get\_detailed\_status() (RTabSeparatedValuesFile method), 133
- get\_execution\_dependency\_tree() (BaseModuleInfo method), 100
- get\_extra\_outputs\_from\_options() (BaseModuleInfo static method), 96
- get\_head\_word() (Entity method), 112
- get\_input() (BaseModuleInfo method), 98
- get\_input\_datatype() (BaseModuleInfo method), 98
- get\_input\_decorator() (in module pimlico.cli.debug.stepper), 75
- get\_input\_module\_connection() (BaseModuleInfo method), 97
- get\_input\_software\_dependencies() (BaseModuleInfo method), 99
- get\_input\_software\_dependencies() (MultistageModuleInfo method), 104
- get\_installed\_version() (PythonPackageDependency method), 86
- get\_installed\_version() (PythonPackageOnPip method), 87
- get\_installed\_version() (SoftwareDependency method), 84
- get\_key\_info\_table() (pimlico.core.modules.base.BaseModuleInfo class method), 95
- get\_key\_info\_table() (pimlico.core.modules.multistage.MultistageModuleInfo class method), 104
- get\_log\_file() (in module pimlico.core.logs), 110
- get\_metadata() (BaseModuleInfo method), 96
- get\_module\_classpath() (in module pimlico.core.dependencies.java), 86
- get\_module\_output\_dir() (BaseModuleInfo method), 97
- get\_module\_schedule() (PipelineConfig method), 108
- get\_names() (DataShell method), 76
- get\_new\_log\_filename() (BaseModuleInfo method), 100
- get\_next\_output\_document() (InputQueueFeeder method), 93
- get\_next\_stage() (MultistageModuleInfo method), 104
- get\_nowait() (OutputQueue method), 156
- get\_output() (BaseModuleInfo method), 97
- get\_output\_datatype() (BaseModuleInfo method), 97
- get\_output\_dir() (BaseModuleInfo method), 97
- get\_output\_software\_dependencies() (BaseModuleInfo method), 99
- get\_paths() (UnnamedFileCollection method), 127
- get\_paths\_from\_options() (UnnamedFileCollection method), 127
- get\_pipeline() (in module pimlico.cli.pyshell), 81
- get\_pop\_up\_parameters() (InputPopupLauncher method), 75
- get\_pop\_up\_parameters() (MessagePopupLauncher method), 75
- get\_progress\_bar() (in module pimlico.utils.progress), 157
- get\_redirect\_func() (in module pimlico.core.external.java), 88
- get\_required\_paths() (PimlicoDatatype method), 117, 141
- get\_required\_paths() (RawTextDirectory method), 128

get\_software\_dependencies() (BaseModuleInfo method), 99

get\_software\_dependencies() (KerasModel method), 131

get\_software\_dependencies() (MultistageModuleInfo method), 104

get\_software\_dependencies() (NumpyArray method), 116, 140

get\_software\_dependencies() (PimlicoDatatype method), 118, 141

get\_software\_dependencies() (ScipySparseMatrix method), 116, 140

get\_software\_dependencies() (SklearnModel method), 134

get\_software\_dependencies() (Word2VecModel method), 136

get\_software\_dependencies() (XmlDocumentIterator method), 138, 149

get\_storage\_roots() (PipelineConfig method), 109

get\_struct() (in module pimlico.datatypes.table), 134

get\_transitive\_dependencies() (BaseModuleInfo method), 99

get\_uninstalled\_dependencies() (TestPipeline method), 150

get\_unused\_local\_port() (in module pimlico.utils.network), 155

get\_unused\_local\_ports() (in module pimlico.utils.network), 155

get\_writer() (DocumentMapModuleInfo method), 92

get\_writers() (DocumentMapModuleInfo method), 92

GraphvizDependency (class in pimlico.core.visualize.deps), 107

## H

help\_text (MetadataCmd attribute), 77

help\_text (PythonCmd attribute), 77

help\_text (ShellCommand attribute), 76

## I

import\_member() (in module pimlico.utils.core), 152

import\_package() (PythonPackageDependency method), 86

incomplete\_tasks (PimlicoDatatypeWriter attribute), 120, 142

increment() (SafeProgressBar method), 157

index2vocab (Embeddings attribute), 123

index2word (Embeddings attribute), 123

IndexedTermFeatureListCorpus (class in pimlico.datatypes), 145

IndexedTermFeatureListCorpus (class in pimlico.datatypes.features), 125

IndexedTermFeatureListCorpusWriter (class in pimlico.datatypes), 145

IndexedTermFeatureListCorpusWriter (class in pimlico.datatypes.features), 125

IndexedTermFeatureListDataPointType (class in pimlico.datatypes), 146

IndexedTermFeatureListDataPointType (class in pimlico.datatypes.features), 124

infinite\_cycle() (in module pimlico.utils.core), 152

input\_corpora (DocumentMapModuleInfo attribute), 92

input\_datatype\_list() (in module pimlico.utils.docs.modulegen), 151

input\_datatype\_text() (in module pimlico.utils.docs.modulegen), 151

input\_module\_factory() (in module pimlico.core.modules.inputs), 103

input\_module\_options (DataPointType attribute), 122, 149

input\_module\_options (PimlicoDatatype attribute), 117, 142

input\_module\_options (RawTextDirectory attribute), 128

input\_module\_options (RawTextDocumentType attribute), 123, 150

input\_module\_options (UnnamedFileCollection attribute), 127

input\_module\_options (XmlDocumentIterator attribute), 138, 149

input\_names (BaseModuleInfo attribute), 96

input\_ready() (BaseModuleInfo method), 98

InputDialog (class in pimlico.cli.browser.tool), 74

InputModuleInfo (class in pimlico.core.modules.inputs), 103

InputPopupLauncher (class in pimlico.cli.browser.tool), 75

InputQueueFeeder (class in pimlico.core.modules.map), 93

InputsCmd (class in pimlico.cli.locations), 79

install() (in module pimlico.core.dependencies.base), 84

install() (JavaJarsDependency method), 85

install() (Py4JSoftwareDependency method), 86

install() (PythonPackageOnPip method), 87

install() (SoftwareDependency method), 84

install\_core\_dependencies() (in module pimlico), 158

install\_dependencies() (in module pimlico.core.dependencies.base), 84

installable() (JavaDependency method), 85

installable() (JavaJarsDependency method), 85

installable() (Py4JSoftwareDependency method), 86

installable() (PythonPackageOnPip method), 87

installable() (PythonPackageSystemwideInstall method), 87

installable() (SoftwareDependency method), 83

installable() (SystemCommandDependency method), 84

installation\_instructions() (GraphvizDependency method), 107

installation\_instructions() (PythonPackageSystemwideInstall method), 87

installation\_instructions() (SoftwareDependency

method), 84  
 InstallationError, 84  
 InstallCmd (class in pimlico.cli.check), 77  
 instantiate\_additional\_datatype() (PimlicoDatatype method), 118, 142  
 instantiate\_output\_datatype() (BaseModuleInfo method), 97  
 instantiate\_output\_datatype() (InputModuleInfo method), 103  
 instantiate\_output\_datatype\_decorator() (in module pimlico.cli.debug.stepper), 75  
 int\_size (IntegerListDocumentType attribute), 130  
 IntegerListDocumentCorpus (class in pimlico.datatypes.ints), 130  
 IntegerListDocumentCorpusWriter (class in pimlico.datatypes.ints), 130  
 IntegerListDocumentType (class in pimlico.datatypes.ints), 130  
 IntegerListsDocumentCorpus (class in pimlico.datatypes.ints), 130  
 IntegerListsDocumentCorpusWriter (class in pimlico.datatypes.ints), 130  
 IntegerListsDocumentType (class in pimlico.datatypes.ints), 130  
 IntegerTableDocumentCorpus (class in pimlico.datatypes.table), 134  
 IntegerTableDocumentCorpusWriter (class in pimlico.datatypes.table), 134  
 IntegerTableDocumentType (class in pimlico.datatypes.table), 134  
 InternalModuleConnection (class in pimlico.core.modules.multistage), 105  
 invalid\_doc\_on\_error() (in module pimlico.core.modules.map), 93  
 invalid\_docs\_on\_error() (in module pimlico.core.modules.map), 93  
 invalid\_document\_or\_text() (InvalidDocument static method), 121, 144  
 InvalidDocument (class in pimlico.datatypes), 143  
 InvalidDocument (class in pimlico.datatypes.base), 120  
 InvalidDocumentFormatter (class in pimlico.cli.browser.formatter), 74  
 is\_filter() (pimlico.core.modules.base.BaseModuleInfo class method), 98  
 is\_identifier() (in module pimlico.utils.core), 152  
 is\_input() (pimlico.core.modules.base.BaseModuleInfo class method), 99  
 is\_locked() (BaseModuleInfo method), 100  
 is\_locked() (MultistageModuleInfo method), 104  
 is\_multiple\_input() (BaseModuleInfo method), 97  
 iterable\_input\_reader\_factory() (in module pimlico.core.modules.inputs), 104  
 IterableCorpus (class in pimlico.datatypes), 142  
 IterableCorpus (class in pimlico.datatypes.base), 120

IterableCorpusWriter (class in pimlico.datatypes), 143  
 IterableCorpusWriter (class in pimlico.datatypes.base), 120

## J

jars (Py4JSoftwareDependency attribute), 86  
 JavaDependency (class in pimlico.core.dependencies.java), 85  
 JavaJarsDependency (class in pimlico.core.dependencies.java), 85  
 JavaProcessError, 88  
 json\_string() (in module pimlico.core.modules.options), 106  
 JsonDocumentCorpus (class in pimlico.datatypes), 146  
 JsonDocumentCorpus (class in pimlico.datatypes.jsondoc), 131  
 JsonDocumentCorpusWriter (class in pimlico.datatypes), 146  
 JsonDocumentCorpusWriter (class in pimlico.datatypes.jsondoc), 131

## K

KerasModel (class in pimlico.datatypes.keras), 131  
 KerasModelBuilderClass (class in pimlico.datatypes.keras), 131  
 KerasModelBuilderClassWriter (class in pimlico.datatypes.keras), 131  
 KerasModelWriter (class in pimlico.datatypes.keras), 131  
 keypress() (InputDialog method), 74  
 KeyValueTypeListCorpus (class in pimlico.datatypes), 145  
 KeyValueTypeListCorpus (class in pimlico.datatypes.features), 124  
 KeyValueTypeListCorpusWriter (class in pimlico.datatypes), 145  
 KeyValueTypeListCorpusWriter (class in pimlico.datatypes.features), 124  
 KeyValueTypeListDocumentType (class in pimlico.datatypes), 145  
 KeyValueTypeListDocumentType (class in pimlico.datatypes.features), 124

## L

label (NumericResult attribute), 133  
 launch\_gateway() (in module pimlico.core.external.java), 88  
 launch\_shell() (in module pimlico.cli.shell.runner), 77  
 length (StreamCommunicationPacket attribute), 152  
 limited\_shuffle() (in module pimlico.utils.probability), 156  
 limited\_shuffle\_numpy() (in module pimlico.utils.probability), 156  
 list\_archive\_iter() (TarredCorpus method), 135, 147  
 load() (InvalidDocument static method), 120, 144  
 load() (PipelineConfig static method), 108

- `load_build_params()` (KerasModelBuilderClass method), [132](#)
  - `load_datatype()` (in module `pimlico.datatypes`), [144](#)
  - `load_datatype()` (in module `pimlico.datatypes.base`), [121](#)
  - `load_executor()` (BaseModuleInfo method), [95](#)
  - `load_formatter()` (in module `pimlico.cli.browser.formatter`), [74](#)
  - `load_local_config()` (PipelineConfig static method), [108](#)
  - `load_model()` (KerasModel method), [131](#)
  - `load_model()` (KerasModelBuilderClass method), [132](#)
  - `load_model()` (SklearnModel method), [134](#)
  - `load_model()` (Word2VecModel method), [136](#)
  - `load_module_executor()` (in module `pimlico.core.modules.base`), [101](#)
  - `load_module_info()` (in module `pimlico.core.modules.base`), [101](#)
  - `load_pipeline()` (TestPipeline static method), [150](#)
  - `LoadCmd` (class in `pimlico.cli.loaddump`), [78](#)
  - `lock()` (BaseModuleInfo method), [100](#)
  - `lock_path` (BaseModuleInfo attribute), [100](#)
  - `LongStoreCmd` (class in `pimlico.cli.main`), [79](#)
- ## M
- `main_module` (BaseModuleInfo attribute), [95](#)
  - `make_py4j_errors_safe()` (in module `pimlico.core.external.java`), [88](#)
  - `make_table()` (in module `pimlico.utils.docs.rest`), [151](#)
  - `Mention` (class in `pimlico.datatypes.coref.corenlp`), [111](#)
  - `Mention` (class in `pimlico.datatypes.coref.opennlp`), [112](#)
  - `MessageDialog` (class in `pimlico.cli.browser.tool`), [74](#)
  - `MessagePopupLauncher` (class in `pimlico.cli.browser.tool`), [75](#)
  - `metadata` (PimlicoDatatype attribute), [117](#), [142](#)
  - `metadata_filename` (BaseModuleInfo attribute), [96](#)
  - `MetadataCmd` (class in `pimlico.cli.shell.commands`), [76](#)
  - `missing_data()` (BaseModuleInfo method), [98](#)
  - `missing_module_data()` (BaseModuleInfo method), [98](#)
  - `model` (Word2VecModel attribute), [136](#)
  - `module_dependencies` (PipelineConfig attribute), [107](#)
  - `module_dependents` (PipelineConfig attribute), [107](#)
  - `module_executable` (BaseModuleInfo attribute), [95](#)
  - `module_executable` (InputModuleInfo attribute), [103](#)
  - `module_executable` (MultistageModuleInfo attribute), [104](#)
  - `module_executor_override` (BaseModuleInfo attribute), [95](#)
  - `module_inputs` (BaseModuleInfo attribute), [95](#)
  - `module_number_to_name()` (in module `pimlico.cli.util`), [83](#)
  - `module_numbers_to_names()` (in module `pimlico.cli.util`), [83](#)
  - `module_optional_outputs` (BaseModuleInfo attribute), [95](#)
  - `module_options` (BaseModuleInfo attribute), [95](#)
  - `module_outputs` (BaseModuleInfo attribute), [95](#)
  - `module_outputs` (DocumentMapModuleInfo attribute), [92](#)
  - `module_package_name()` (`pimlico.core.modules.base.BaseModuleInfo` class method), [100](#)
  - `module_readable_name` (BaseModuleInfo attribute), [95](#)
  - `module_status()` (in module `pimlico.cli.status`), [82](#)
  - `module_status_color()` (in module `pimlico.cli.status`), [82](#)
  - `module_type_name` (BaseModuleInfo attribute), [95](#)
  - `module_type_name` (InputModuleInfo attribute), [103](#)
  - `ModuleAlreadyCompletedError`, [103](#)
  - `ModuleConnection` (class in `pimlico.core.modules.multistage`), [105](#)
  - `ModuleExecutionError`, [102](#)
  - `ModuleExecutorLoadError`, [101](#)
  - `ModuleInfoLoadError`, [101](#)
  - `ModuleInputConnection` (class in `pimlico.core.modules.multistage`), [105](#)
  - `ModuleNotReadyError`, [103](#)
  - `ModuleOptionParseError`, [106](#)
  - `ModuleOutputConnection` (class in `pimlico.core.modules.multistage`), [106](#)
  - `modules` (PipelineConfig attribute), [107](#)
  - `ModuleStage` (class in `pimlico.core.modules.multistage`), [105](#)
  - `ModuleTypeError`, [101](#)
  - `multiline_tablate()` (in module `pimlico.utils.format`), [155](#)
  - `MultipleInputs` (class in `pimlico.datatypes`), [144](#)
  - `MultipleInputs` (class in `pimlico.datatypes.base`), [119](#)
  - `multiprocessing_executor_factory()` (in module `pimlico.core.modules.map.multiproc`), [90](#)
  - `MultiprocessingMapModuleExecutor` (class in `pimlico.core.modules.map.multiproc`), [90](#)
  - `MultiprocessingMapPool` (class in `pimlico.core.modules.map.multiproc`), [89](#)
  - `MultiprocessingMapProcess` (class in `pimlico.core.modules.map.multiproc`), [89](#)
  - `multistage_module()` (in module `pimlico.core.modules.multistage`), [105](#)
  - `MultistageModuleInfo` (class in `pimlico.core.modules.multistage`), [104](#)
  - `MultistageModulePreparationError`, [106](#)
  - `multiwith()` (in module `pimlico.utils.core`), [152](#)
- ## N
- `named_file_collection_union()` (in module `pimlico.datatypes.files`), [126](#)
  - `NamedFile()` (in module `pimlico.datatypes.files`), [128](#)
  - `NamedFileCollection` (class in `pimlico.datatypes.files`), [126](#)
  - `NamedFileCollectionWriter` (class in `pimlico.datatypes.files`), [126](#)
  - `NamedFileWriter` (class in `pimlico.datatypes.files`), [128](#)
  - `new_client()` (Py4JInterface method), [87](#)

[new\\_filename\(\)](#) (in module `pimlico.utils.filesystem`), 154  
[NewModuleCmd](#) (class in `pimlico.cli.newmodule`), 80  
[next\\_document\(\)](#) (`CorpusState` method), 74  
[no\\_retry\\_gateway\(\)](#) (in module `pimlico.core.external.java`), 88  
[non\\_filter\\_datatype](#) (`DocumentMapOutputTypeWrapper` attribute), 89  
[NonPTBTagError](#), 156  
[normalize\\_cell\(\)](#) (in module `pimlico.utils.docs.rest`), 151  
[normed\\_vectors](#) (`Embeddings` attribute), 123  
[notify\\_no\\_more\\_inputs\(\)](#) (`DocumentMapProcessMixin` method), 94  
[notify\\_no\\_more\\_inputs\(\)](#) (`DocumentProcessorPool` method), 94  
[notify\\_no\\_more\\_inputs\(\)](#) (`MultiprocessingMapPool` method), 90  
[notify\\_no\\_more\\_inputs\(\)](#) (`MultiprocessingMapProcess` method), 89  
[notify\\_no\\_more\\_inputs\(\)](#) (`ThreadingMapThread` method), 91  
[NumericResult](#) (class in `pimlico.datatypes.results`), 133  
[NumericResultWriter](#) (class in `pimlico.datatypes.results`), 133  
[NumpyArray](#) (class in `pimlico.datatypes`), 140  
[NumpyArray](#) (class in `pimlico.datatypes.arrays`), 116  
[NumpyArrayWriter](#) (class in `pimlico.datatypes`), 140  
[NumpyArrayWriter](#) (class in `pimlico.datatypes.arrays`), 116

## O

[OpenNLPCorefCorpus](#) (in module `pimlico.datatypes`), 138  
[OpenNLPCorefCorpus](#) (in module `pimlico.datatypes.coref`), 112  
[OpenNLPCorefCorpusWriter](#) (in module `pimlico.datatypes`), 138  
[OpenNLPCorefCorpusWriter](#) (in module `pimlico.datatypes.coref`), 112  
[opt\\_type\\_help\(\)](#) (in module `pimlico.core.modules.options`), 106  
[option\\_message\(\)](#) (in module `pimlico.cli.debug.stepper`), 75  
[output\\_datatype\\_text\(\)](#) (in module `pimlico.utils.docs.modulegen`), 151  
[output\\_name](#) (`DocumentMapOutputTypeWrapper` attribute), 89  
[output\\_names](#) (`BaseModuleInfo` attribute), 96  
[output\\_p4j\\_error\\_info\(\)](#) (in module `pimlico.core.external.java`), 88  
[output\\_stack\\_trace\(\)](#) (in module `pimlico.cli.debug`), 76  
[OutputCmd](#) (class in `pimlico.cli.locations`), 79  
[OutputConsumer](#) (class in `pimlico.core.external.java`), 88  
[OutputQueue](#) (class in `pimlico.utils.pipes`), 156

## P

[path](#) (`StringList` attribute), 122  
[path\\_name\\_to\\_doc\\_name\(\)](#) (`UnnamedFileCollection` method), 127  
[path\\_relative\\_to\\_config\(\)](#) (`PipelineConfig` method), 108  
[pdf\\_path](#) (`PlotOutput` attribute), 132  
[pimlico](#) (module), 158  
[pimlico.cli](#) (module), 83  
[pimlico.cli.browser](#) (module), 75  
[pimlico.cli.browser.formatter](#) (module), 73  
[pimlico.cli.browser.tool](#) (module), 74  
[pimlico.cli.check](#) (module), 77  
[pimlico.cli.clean](#) (module), 78  
[pimlico.cli.debug](#) (module), 75  
[pimlico.cli.debug.stepper](#) (module), 75  
[pimlico.cli.loaddump](#) (module), 78  
[pimlico.cli.locations](#) (module), 79  
[pimlico.cli.main](#) (module), 79  
[pimlico.cli.newmodule](#) (module), 80  
[pimlico.cli.pyshell](#) (module), 81  
[pimlico.cli.reset](#) (module), 81  
[pimlico.cli.run](#) (module), 81  
[pimlico.cli.shell](#) (module), 77  
[pimlico.cli.shell.base](#) (module), 76  
[pimlico.cli.shell.commands](#) (module), 76  
[pimlico.cli.shell.runner](#) (module), 77  
[pimlico.cli.status](#) (module), 82  
[pimlico.cli.subcommands](#) (module), 82  
[pimlico.cli.testemail](#) (module), 82  
[pimlico.cli.util](#) (module), 83  
[pimlico.core](#) (module), 111  
[pimlico.core.config](#) (module), 107  
[pimlico.core.dependencies](#) (module), 87  
[pimlico.core.dependencies.base](#) (module), 83  
[pimlico.core.dependencies.core](#) (module), 85  
[pimlico.core.dependencies.java](#) (module), 85  
[pimlico.core.dependencies.python](#) (module), 86  
[pimlico.core.dependencies.versions](#) (module), 87  
[pimlico.core.external](#) (module), 88  
[pimlico.core.external.java](#) (module), 87  
[pimlico.core.logs](#) (module), 110  
[pimlico.core.modules](#) (module), 106  
[pimlico.core.modules.base](#) (module), 95  
[pimlico.core.modules.execute](#) (module), 102  
[pimlico.core.modules.inputs](#) (module), 103  
[pimlico.core.modules.map](#) (module), 92  
[pimlico.core.modules.map.filter](#) (module), 89  
[pimlico.core.modules.map.multiproc](#) (module), 89  
[pimlico.core.modules.map.singleproc](#) (module), 91  
[pimlico.core.modules.map.threaded](#) (module), 91  
[pimlico.core.modules.multistage](#) (module), 104  
[pimlico.core.modules.options](#) (module), 106  
[pimlico.core.paths](#) (module), 111  
[pimlico.core.visualize](#) (module), 107



pimlico.core.visualize.deps (module), 107  
pimlico.core.visualize.status (module), 107  
pimlico.datatypes (module), 138  
pimlico.datatypes.arrays (module), 116  
pimlico.datatypes.base (module), 116  
pimlico.datatypes.caevo (module), 121  
pimlico.datatypes.core (module), 121  
pimlico.datatypes.coref (module), 112  
pimlico.datatypes.coref.corenlp (module), 111  
pimlico.datatypes.coref.opennlp (module), 111  
pimlico.datatypes.dictionary (module), 122  
pimlico.datatypes.documents (module), 122  
pimlico.datatypes.embeddings (module), 123  
pimlico.datatypes.features (module), 124  
pimlico.datatypes.files (module), 126  
pimlico.datatypes.floats (module), 129  
pimlico.datatypes.formatters (module), 113  
pimlico.datatypes.formatters.features (module), 113  
pimlico.datatypes.formatters.tokenized (module), 113  
pimlico.datatypes.ints (module), 130  
pimlico.datatypes.jsondoc (module), 131  
pimlico.datatypes.keras (module), 131  
pimlico.datatypes.parse (module), 114  
pimlico.datatypes.parse.candc (module), 113  
pimlico.datatypes.parse.dependency (module), 113  
pimlico.datatypes.plotting (module), 132  
pimlico.datatypes.r (module), 132  
pimlico.datatypes.results (module), 133  
pimlico.datatypes.sklearn (module), 133  
pimlico.datatypes.spans (module), 134  
pimlico.datatypes.table (module), 134  
pimlico.datatypes.tar (module), 135  
pimlico.datatypes.tokenized (module), 136  
pimlico.datatypes.word2vec (module), 136  
pimlico.datatypes.word\_annotations (module), 137  
pimlico.datatypes.xml (module), 138  
pimlico.modules (module), 36  
pimlico.modules.caevo (module), 36  
pimlico.modules.caevo.output (module), 37  
pimlico.modules.candc (module), 38  
pimlico.modules.corenlp (module), 38  
pimlico.modules.corpora (module), 39  
pimlico.modules.corpora.concat (module), 39  
pimlico.modules.corpora.corpus\_stats (module), 40  
pimlico.modules.corpora.format (module), 40  
pimlico.modules.corpora.list\_filter (module), 41  
pimlico.modules.corpora.split (module), 41  
pimlico.modules.corpora.subset (module), 42  
pimlico.modules.corpora.tar (module), 43  
pimlico.modules.corpora.tar\_filter (module), 44  
pimlico.modules.corpora.vocab\_builder (module), 44  
pimlico.modules.corpora.vocab\_mapper (module), 45  
pimlico.modules.embeddings (module), 45  
pimlico.modules.embeddings.dependencies (module), 46  
pimlico.modules.embeddings.word2vec (module), 46  
pimlico.modules.features (module), 47  
pimlico.modules.features.term\_feature\_compiler (module), 47  
pimlico.modules.features.term\_feature\_matrix\_builder (module), 48  
pimlico.modules.features.vocab\_builder (module), 48  
pimlico.modules.features.vocab\_mapper (module), 49  
pimlico.modules.input (module), 49  
pimlico.modules.input.embeddings (module), 50  
pimlico.modules.input.embeddings.fasttext (module), 50  
pimlico.modules.input.embeddings.fasttext\_gensim (module), 50  
pimlico.modules.input.text (module), 51  
pimlico.modules.input.text.raw\_text\_files (module), 51  
pimlico.modules.malt (module), 52  
pimlico.modules.malt.conll\_parser\_input (module), 52  
pimlico.modules.malt.parse (module), 52  
pimlico.modules.opennlp (module), 53  
pimlico.modules.opennlp.coreference (module), 53  
pimlico.modules.opennlp.coreference\_pipeline (module), 54  
pimlico.modules.opennlp.ner (module), 55  
pimlico.modules.opennlp.parse (module), 56  
pimlico.modules.opennlp.pos (module), 56  
pimlico.modules.opennlp.tokenize (module), 57  
pimlico.modules.r (module), 57  
pimlico.modules.r.script (module), 58  
pimlico.modules.regex (module), 58  
pimlico.modules.regex.annotated\_text (module), 58  
pimlico.modules.sklearn (module), 59  
pimlico.modules.sklearn.matrix\_factorization (module), 59  
pimlico.modules.utility (module), 60  
pimlico.modules.utility.alias (module), 60  
pimlico.modules.utility.copy\_file (module), 61  
pimlico.modules.visualization (module), 61  
pimlico.modules.visualization.bar\_chart (module), 62  
pimlico.modules.visualization.embeddings\_plot (module), 62  
pimlico.test (module), 151  
pimlico.test.pipeline (module), 150  
pimlico.utils (module), 158  
pimlico.utils.communicate (module), 152  
pimlico.utils.core (module), 152  
pimlico.utils.docs (module), 152  
pimlico.utils.docs.commandgen (module), 151  
pimlico.utils.docs.modulegen (module), 151  
pimlico.utils.docs.rest (module), 151  
pimlico.utils.email (module), 153  
pimlico.utils.filesystem (module), 154  
pimlico.utils.format (module), 155  
pimlico.utils.linguistic (module), 155  
pimlico.utils.logging (module), 155

- pimlico.utils.network (module), 155
- pimlico.utils.pipes (module), 155
- pimlico.utils.pos (module), 156
- pimlico.utils.probability (module), 156
- pimlico.utils.progress (module), 157
- pimlico.utils.strings (module), 158
- pimlico.utils.system (module), 158
- pimlico.utils.timeout (module), 158
- pimlico.utils.web (module), 158
- PimlicoCLISubcommand (class in pimlico.cli.subcommands), 82
- PimlicoDatatype (class in pimlico.datatypes), 140
- PimlicoDatatype (class in pimlico.datatypes.base), 116
- PimlicoDatatypeWriter (class in pimlico.datatypes), 142
- PimlicoDatatypeWriter (class in pimlico.datatypes.base), 120
- PimlicoJavaLibrary (class in pimlico.core.dependencies.java), 85
- PimlicoPythonShellContext (class in pimlico.cli.pyshell), 81
- PipelineCheckError, 109
- PipelineConfig (class in pimlico.core.config), 107
- PipelineConfigParseError, 109
- PipelineStructureError, 109
- plot() (PlotOutput method), 132
- PlotOutput (class in pimlico.datatypes.plotting), 132
- PlotOutputWriter (class in pimlico.datatypes.plotting), 132
- POOL\_TYPE (MultiprocessingMapModuleExecutor attribute), 90
- POOL\_TYPE (ThreadingMapModuleExecutor attribute), 92
- pos\_tag\_to\_ptb() (in module pimlico.utils.pos), 156
- pos\_tags\_to\_ptb() (in module pimlico.utils.pos), 156
- postloop() (DataShell method), 76
- postprocess() (DocumentMapModuleExecutor method), 93
- postprocess() (MultiprocessingMapModuleExecutor method), 90
- postprocess() (ThreadingMapModuleExecutor method), 92
- preloop() (DataShell method), 76
- prepare\_data() (PimlicoDatatype method), 118, 142
- prepare\_data() (RawTextDirectory method), 128
- prepare\_data() (XmlDocumentIterator method), 138, 149
- preprocess() (DocumentMapModuleExecutor method), 93
- preprocess\_config\_file() (in module pimlico.core.config), 110
- print\_dependency\_leaf\_problems() (in module pimlico.core.config), 110
- print\_execution\_error() (in module pimlico.cli.util), 83
- print\_missing\_dependencies() (in module pimlico.core.config), 110
- problems() (JavaDependency method), 85
- problems() (PythonPackageDependency method), 86
- problems() (SoftwareDependency method), 83
- problems() (SystemCommandDependency method), 84
- process\_document() (CorefDocumentType method), 111, 112
- process\_document() (DocumentMapProcessMixin method), 94
- process\_document() (FeatureListScoreDocumentType method), 125, 146
- process\_document() (FloatListDocumentType method), 129
- process\_document() (FloatListsDocumentType method), 129
- process\_document() (IntegerListDocumentType method), 130
- process\_document() (IntegerListsDocumentType method), 130
- process\_document() (IntegerTableDocumentType method), 134
- process\_document() (KeyValueListDocumentType method), 124, 145
- process\_document() (RawDocumentType method), 123, 150
- process\_document() (RawTextDocumentType method), 123, 150
- process\_document() (SentenceSpansDocumentType method), 134
- process\_document() (TarredCorpus method), 135, 147
- process\_document() (TermFeatureListDocumentType method), 124, 145
- process\_document() (TreeStringsDocumentType method), 114, 139
- process\_document() (WordAnnotationsDocumentType method), 137, 149
- process\_document\_data\_with\_datatype() (IterableCorpus method), 120, 143
- process\_documents() (DocumentMapProcessMixin method), 94
- process\_module\_options() (in module pimlico.core.modules.options), 106
- process\_module\_options() (pimlico.core.modules.base.BaseModuleInfo class method), 96
- PROCESS\_TYPE (MultiprocessingMapPool attribute), 90
- ProcessOutput (class in pimlico.core.modules.map), 93
- ProgressBarIter (class in pimlico.utils.progress), 158
- prompt (DataShell attribute), 76
- provide\_further\_outputs() (BaseModuleInfo method), 97
- Py4JInterface (class in pimlico.core.external.java), 87
- Py4JSafeJavaError, 88
- Py4JSoftwareDependency (class in pimlico.core.dependencies.java), 86

PythonCmd (class in pimlico.cli.shell.commands), 77  
 PythonPackageDependency (class in pimlico.core.dependencies.python), 86  
 PythonPackageOnPip (class in pimlico.core.dependencies.python), 87  
 PythonPackageSystemwideInstall (class in pimlico.core.dependencies.python), 86  
 PythonShellCmd (class in pimlico.cli.pyshell), 81

## Q

qget() (in module pimlico.utils.pipes), 155

## R

RAW\_INPUT (DocumentBrowserFormatter attribute), 73  
 RawDocumentType (class in pimlico.datatypes), 149  
 RawDocumentType (class in pimlico.datatypes.documents), 123  
 RawTextDirectory (class in pimlico.datatypes.files), 128  
 RawTextDocumentType (class in pimlico.datatypes), 150  
 RawTextDocumentType (class in pimlico.datatypes.documents), 123  
 RawTextFiles (class in pimlico.datatypes.files), 128  
 read() (StreamCommunicationPacket static method), 152  
 read\_annotation\_fields() (WordAnnotationCorpus method), 137, 148  
 read\_data() (SingleTextDocument method), 121  
 read\_floats() (FloatListDocumentType method), 129  
 read\_ints() (IntegerListDocumentType method), 130  
 read\_rows() (FloatListsDocumentType method), 129  
 read\_rows() (IntegerListsDocumentType method), 130  
 read\_rows() (IntegerTableDocumentType method), 134  
 ReaderOutputType (class in pimlico.core.modules.inputs), 103  
 recursive\_deps() (in module pimlico.core.dependencies.base), 85  
 remove\_duplicates() (in module pimlico.utils.core), 152  
 remove\_temporary\_redirects() (OutputConsumer method), 88  
 require\_tasks() (PimlicoDatatypeWriter method), 120, 142  
 requires\_data\_preparation (PimlicoDatatype attribute), 117, 142  
 requires\_data\_preparation (RawTextDirectory attribute), 128  
 requires\_data\_preparation (XmlDocumentIterator attribute), 138, 149  
 reset\_all\_modules() (PipelineConfig method), 108  
 reset\_execution() (BaseModuleInfo method), 100  
 reset\_execution() (MultistageModuleInfo method), 104  
 ResetCmd (class in pimlico.cli.reset), 81  
 result (NumericResult attribute), 133  
 retrieve\_processing\_status() (DocumentMapModuleExecutor method), 93  
 retry\_open() (in module pimlico.utils.filesystem), 154

row\_size (IntegerTableDocumentType attribute), 134  
 RTabSeparatedValuesFile (class in pimlico.datatypes.r), 132  
 RTabSeparatedValuesFileWriter (class in pimlico.datatypes.r), 133  
 run() (InputQueueFeeder method), 94  
 run() (MultiprocessingMapProcess method), 89  
 run() (OutputConsumer method), 88  
 run() (ThreadingMapThread method), 91  
 run\_command() (BrowseCmd method), 80  
 run\_command() (CleanCmd method), 78  
 run\_command() (DepsCmd method), 77  
 run\_command() (DumpCmd method), 78  
 run\_command() (EmailCmd method), 83  
 run\_command() (InputsCmd method), 79  
 run\_command() (InstallCmd method), 77  
 run\_command() (LoadCmd method), 79  
 run\_command() (LongStoreCmd method), 80  
 run\_command() (NewModuleCmd method), 80  
 run\_command() (OutputCmd method), 79  
 run\_command() (PimlicoCLISubcommand method), 82  
 run\_command() (PythonShellCmd method), 81  
 run\_command() (ResetCmd method), 81  
 run\_command() (RunCmd method), 81  
 run\_command() (ShellCLICmd method), 77  
 run\_command() (StatusCmd method), 82  
 run\_command() (UnlockCmd method), 80  
 run\_command() (VariantsCmd method), 79  
 run\_command() (VisualizeCmd method), 80  
 run\_test\_pipeline() (in module pimlico.test.pipeline), 150  
 run\_test\_suite() (in module pimlico.test.pipeline), 150  
 RunCmd (class in pimlico.cli.run), 81

## S

SafeProgressBar (class in pimlico.utils.progress), 157  
 satisfies\_typecheck() (in module pimlico.core.modules.base), 101  
 save\_popup\_launcher() (in module pimlico.cli.browser.tool), 75  
 ScipySparseMatrix (class in pimlico.datatypes), 140  
 ScipySparseMatrix (class in pimlico.datatypes.arrays), 116  
 ScipySparseMatrixWriter (class in pimlico.datatypes), 140  
 ScipySparseMatrixWriter (class in pimlico.datatypes.arrays), 116  
 script\_path (PlotOutput attribute), 132  
 send\_final\_report\_email() (in module pimlico.core.modules.execute), 102  
 send\_module\_report\_email() (in module pimlico.core.modules.execute), 102  
 send\_pimlico\_email() (in module pimlico.utils.email), 154  
 send\_text\_email() (in module pimlico.utils.email), 154



- sentence\_boundary\_re (WordAnnotationsDocumentType attribute), 137, 149
- SentenceSpansCorpus (class in pimlico.datatypes.spans), 134
- SentenceSpansCorpusWriter (class in pimlico.datatypes.spans), 134
- SentenceSpansDocumentType (class in pimlico.datatypes.spans), 134
- sequential\_document\_sample() (in module pimlico.utils.probability), 156
- sequential\_sample() (in module pimlico.utils.probability), 157
- set\_array() (NumpyArrayWriter method), 116, 140
- set\_matrix() (ScipySparseMatrixWriter method), 116, 140
- set\_metadata\_value() (BaseModuleInfo method), 96
- set\_metadata\_values() (BaseModuleInfo method), 96
- set\_proc\_title() (in module pimlico.utils.system), 158
- set\_up() (DocumentMapProcessMixin method), 94
- shell\_commands (IterableCorpus attribute), 120, 143
- shell\_commands (PimlicoDatatype attribute), 117, 142
- shell\_commands (Word2VecModel attribute), 136
- ShellCLICmd (class in pimlico.cli.shell.runner), 77
- ShellCommand (class in pimlico.cli.shell.base), 76
- ShellContextError, 81
- ShellError, 76
- shutdown() (DocumentProcessorPool method), 94
- shutdown() (InputQueueFeeder method), 94
- shutdown() (MultiprocessingMapPool method), 90
- shutdown() (ThreadingMapPool method), 91
- shutdown() (ThreadingMapThread method), 91
- signals (InputDialog attribute), 74
- similarities() (in module pimlico.utils.strings), 158
- SimpleWordAnnotationCorpusWriter (class in pimlico.datatypes), 148
- SimpleWordAnnotationCorpusWriter (class in pimlico.datatypes.word\_annotations), 137
- single\_process\_executor\_factory() (in module pimlico.core.modules.map.singleproc), 91
- SINGLE\_PROCESS\_TYPE (MultiprocessingMapPool attribute), 90
- SingleTextDocument (class in pimlico.datatypes.core), 121
- SingleTextDocumentWriter (class in pimlico.datatypes.core), 121
- SingleThreadMapModuleExecutor (class in pimlico.core.modules.map.singleproc), 91
- skip() (CorpusState method), 74
- skip\_invalid() (in module pimlico.core.modules.map), 93
- skip\_invalids() (in module pimlico.core.modules.map), 93
- skip\_popup\_launcher() (in module pimlico.cli.browser.tool), 75
- SklearnModel (class in pimlico.datatypes.sklearn), 133
- SklearnModelWriter (class in pimlico.datatypes.sklearn), 133
- slice\_progress() (in module pimlico.utils.progress), 157
- SoftwareDependency (class in pimlico.core.dependencies.base), 83
- SoftwareVersion (class in pimlico.core.dependencies.versions), 87
- sorted\_by\_similarity() (in module pimlico.utils.strings), 158
- split\_seq() (in module pimlico.utils.core), 153
- split\_seq\_after() (in module pimlico.utils.core), 153
- stages (MultistageModuleInfo attribute), 104
- StanfordDependencyParseCorpus (class in pimlico.datatypes), 139
- StanfordDependencyParseCorpus (class in pimlico.datatypes.parse), 115
- StanfordDependencyParseCorpus (class in pimlico.datatypes.parse.dependency), 113
- StanfordDependencyParseCorpusWriter (class in pimlico.datatypes), 139
- StanfordDependencyParseCorpusWriter (class in pimlico.datatypes.parse), 115
- StanfordDependencyParseCorpusWriter (class in pimlico.datatypes.parse.dependency), 114
- start() (Py4JInterface method), 87
- start\_java\_process() (in module pimlico.core.external.java), 87
- start\_worker() (MultiprocessingMapPool method), 90
- start\_worker() (ThreadingMapPool method), 91
- status (BaseModuleInfo attribute), 96
- status (MultistageModuleInfo attribute), 104
- status\_colored() (in module pimlico.cli.status), 82
- StatusCmd (class in pimlico.cli.status), 82
- step (PipelineConfig attribute), 109
- Stepper (class in pimlico.cli.debug.stepper), 75
- stop() (Py4JInterface method), 88
- StopProcessing, 103
- str\_to\_bool() (in module pimlico.core.modules.options), 106
- StreamCommunicationError, 152
- StreamCommunicationPacket (class in pimlico.utils.communicate), 152
- StringList (class in pimlico.datatypes.core), 121
- StringListWriter (class in pimlico.datatypes.core), 122
- strip\_common\_indent() (in module pimlico.utils.docs.commandgen), 151
- strip\_punctuation() (in module pimlico.utils.linguistic), 155
- subordinate\_additional\_name() (PimlicoDatatypeWriter method), 120, 142
- subsample() (in module pimlico.utils.probability), 157
- supplied\_additional (PimlicoDatatype attribute), 117, 142
- supplied\_additional (PlotOutput attribute), 132

SystemCommandDependency (class in pimlico.core.dependencies.base), 84

## T

table\_div() (in module pimlico.utils.docs.rest), 151  
 TarredCorpus (class in pimlico.datatypes), 147  
 TarredCorpus (class in pimlico.datatypes.tar), 135  
 TarredCorpusIterationError, 136, 148  
 TarredCorpusWriter (class in pimlico.datatypes), 147  
 TarredCorpusWriter (class in pimlico.datatypes.tar), 135  
 task\_complete() (PimlicoDatatypeWriter method), 120, 142  
 tear\_down() (DocumentMapProcessMixin method), 94  
 term\_dictionary (IndexedTermFeatureListCorpus attribute), 125, 145  
 TermFeatureListCorpus (class in pimlico.datatypes), 145  
 TermFeatureListCorpus (class in pimlico.datatypes.features), 124  
 TermFeatureListCorpusWriter (class in pimlico.datatypes), 145  
 TermFeatureListCorpusWriter (class in pimlico.datatypes.features), 124  
 TermFeatureListDocumentType (class in pimlico.datatypes), 145  
 TermFeatureListDocumentType (class in pimlico.datatypes.features), 124  
 terminate\_process() (in module pimlico.utils.communicate), 152  
 test\_all\_modules() (TestPipeline method), 150  
 test\_input\_module() (TestPipeline method), 150  
 test\_module\_execution() (TestPipeline method), 150  
 TestPipeline (class in pimlico.test.pipeline), 150  
 TestPipelineRunError, 151  
 THREAD\_TYPE (ThreadingMapPool attribute), 91  
 threading\_executor\_factory() (in module pimlico.core.modules.map.threaded), 92  
 ThreadingMapModuleExecutor (class in pimlico.core.modules.map.threaded), 91  
 ThreadingMapPool (class in pimlico.core.modules.map.threaded), 91  
 ThreadingMapThread (class in pimlico.core.modules.map.threaded), 91  
 timeout() (in module pimlico.utils.timeout), 158  
 timeout\_process() (in module pimlico.utils.communicate), 152  
 title\_box() (in module pimlico.utils.format), 155  
 to\_json\_dict() (Entity method), 112  
 to\_json\_dict() (Mention method), 111, 112  
 to\_keyed\_vectors() (Embeddings method), 123  
 TokenizedCorpus (class in pimlico.datatypes), 148  
 TokenizedCorpus (class in pimlico.datatypes.tokenized), 136  
 TokenizedCorpusWriter (class in pimlico.datatypes), 148

TokenizedCorpusWriter (class in pimlico.datatypes.tokenized), 136  
 TokenizedDocumentFormatter (class in pimlico.datatypes.formatters.tokenized), 113  
 TreeStringsDocumentType (class in pimlico.datatypes), 139  
 TreeStringsDocumentType (class in pimlico.datatypes.parse), 114  
 trim\_docstring() (in module pimlico.utils.docs), 152  
 truncate() (in module pimlico.utils.strings), 158  
 type\_checking\_name() (DynamicInputDatatypeRequirement method), 119, 143  
 type\_checking\_name() (pimlico.datatypes.base.IterableCorpus class method), 120  
 type\_checking\_name() (pimlico.datatypes.base.PimlicoDatatype class method), 119  
 type\_checking\_name() (pimlico.datatypes.IterableCorpus class method), 143  
 type\_checking\_name() (pimlico.datatypes.PimlicoDatatype class method), 142  
 typecheck\_formatter() (in module pimlico.cli.browser.formatter), 74  
 typecheck\_inputs() (BaseModuleInfo method), 99  
 typecheck\_inputs() (MultistageModuleInfo method), 104  
 TypeCheckError, 101

## U

unlock() (BaseModuleInfo method), 100  
 UnlockCmd (class in pimlico.cli.main), 80  
 UnnamedFileCollection (class in pimlico.datatypes.files), 127  
 UnnamedFileCollectionWriter (class in pimlico.datatypes.files), 127  
 unpacker (IntegerListDocumentType attribute), 130  
 unpacker (IntegerListsDocumentType attribute), 130  
 unpacker (IntegerTableDocumentType attribute), 134  
 update() (SafeProgressBar method), 157  
 update\_processing\_status() (DocumentMapModuleExecutor method), 93

## V

VariantsCmd (class in pimlico.cli.main), 79  
 vector\_size (Embeddings attribute), 123  
 vectors (Embeddings attribute), 123  
 VisualizeCmd (class in pimlico.cli.main), 80  
 Vocab (class in pimlico.datatypes.embeddings), 123  
 vocab (Embeddings attribute), 123

## W

walk() (RawTextDirectory method), 128

[weights\\_filename](#) (KerasModelBuilderClass attribute), [132](#)  
[Word2VecModel](#) (class in `pimlico.datatypes.word2vec`), [136](#)  
[Word2VecModelWriter](#) (class in `pimlico.datatypes.word2vec`), [136](#)  
[word\\_boundary](#) (WordAnnotationsDocumentType attribute), [137](#), [149](#)  
[word\\_re](#) (WordAnnotationsDocumentType attribute), [137](#), [149](#)  
[word\\_vec\(\)](#) (Embeddings method), [123](#)  
[word\\_vecs\(\)](#) (Embeddings method), [123](#)  
[WordAnnotationCorpus](#) (class in `pimlico.datatypes`), [148](#)  
[WordAnnotationCorpus](#) (class in `pimlico.datatypes.word_annotations`), [137](#)  
[WordAnnotationCorpusWithRequiredFields](#) (class in `pimlico.datatypes`), [149](#)  
[WordAnnotationCorpusWithRequiredFields](#) (class in `pimlico.datatypes.word_annotations`), [137](#)  
[WordAnnotationCorpusWriter](#) (class in `pimlico.datatypes`), [148](#)  
[WordAnnotationCorpusWriter](#) (class in `pimlico.datatypes.word_annotations`), [137](#)  
[WordAnnotationsDocumentType](#) (class in `pimlico.datatypes`), [149](#)  
[WordAnnotationsDocumentType](#) (class in `pimlico.datatypes.word_annotations`), [137](#)  
[WorkerShutdownError](#), [94](#)  
[WorkerStartupError](#), [94](#)  
[wrap\\_module\\_info\\_as\\_filter\(\)](#) (in module `pimlico.core.modules.map.filter`), [89](#)  
[wrap\\_tarred\\_corpus\(\)](#) (in module `pimlico.cli.debug.stepper`), [75](#)  
[wrapped\\_module\\_info](#) (DocumentMapOutputTypeWrapper attribute), [89](#)  
[write\\_architecture\(\)](#) (KerasModelWriter method), [131](#)  
[write\\_data\(\)](#) (NamedFileWriter method), [128](#)  
[write\\_dictionaries\(\)](#) (IndexedTermFeatureListCorpusWriter method), [125](#), [146](#)  
[write\\_file\(\)](#) (NamedFileCollectionWriter method), [126](#)  
[write\\_file\(\)](#) (UnnamedFileCollectionWriter method), [128](#)  
[write\\_keyed\\_vectors\(\)](#) (EmbeddingsWriter method), [124](#)  
[write\\_keyed\\_vectors\(\)](#) (Word2VecModelWriter method), [137](#)  
[write\\_metadata\(\)](#) (PimlicoDatatypeWriter method), [120](#), [142](#)  
[write\\_model\(\)](#) (KerasModelWriter method), [131](#)  
[write\\_model\(\)](#) (SklearnModelWriter method), [133](#)  
[write\\_row\(\)](#) (RTabSeparatedValuesFileWriter method), [133](#)  
[write\\_vectors\(\)](#) (EmbeddingsWriter method), [124](#)  
[write\\_vocab\\_list\(\)](#) (EmbeddingsWriter method), [124](#)  
[write\\_weights\(\)](#) (KerasModelBuilderClassWriter method), [131](#)  
[write\\_weights\(\)](#) (KerasModelWriter method), [131](#)  
[write\\_word2vec\\_model\(\)](#) (Word2VecModelWriter method), [137](#)  
[write\\_word\\_counts\(\)](#) (EmbeddingsWriter method), [124](#)

## X

[XmlDocumentIterator](#) (class in `pimlico.datatypes`), [149](#)  
[XmlDocumentIterator](#) (class in `pimlico.datatypes.xml`), [138](#)