
nctoolkit

Robert Wilson

Oct 22, 2021

QUICK OVERVIEW

1	Installation	3
2	Introduction to nctoolkit	5
3	News	9
4	Datasets	13
5	Cheat sheet	17
6	Plotting	19
7	Importing and exporting data	21
8	Temporal statistics	23
9	Subsetting data	27
10	Manipulating variables	29
11	Interpolation	33
12	Ensemble methods	37
13	Parallel processing	41
14	Examples	43
15	Random Data Hacks	47
16	Global settings	49
17	API Reference	51
18	Package info	103
	Python Module Index	105
	Index	107

nctoolkit is a comprehensive Python package for analyzing netCDF data on Linux and macOS.

Core abilities include:

- Cropping to geographic regions
- Interactive plotting of data
- Subsetting to specific time periods
- Calculating time averages
- Calculating spatial averages
- Calculating rolling averages
- Calculating climatologies
- Creating new variables using arithmetic operations
- Calculating anomalies
- Horizontally and vertically remapping data
- Calculating the correlations between variables
- Calculating vertical averages for the likes of oceanic data
- Calculating ensemble averages
- Calculating phenological metrics

INSTALLATION

1.1 How to install nctoolkit

nctoolkit is available from the [Python Packaging Index](#). To install nctoolkit using pip:

```
$ pip install numpy
$ pip install nctoolkit
```

If you already have numpy installed, ignore the first line. This is only included as it will make installing some dependencies smoother. nctoolkit partly relies on cartopy for plotting. This has some additional dependencies, so you may need to follow their guide [here](#) to ensure cartopy is installed fully. If you install nctoolkit using conda, you will not need to worry about that.

If you install nctoolkit from pypi, you will need to install the system dependencies listed below.

nctoolkit can also be installed using conda, as follows:

```
$ conda install -c conda-forge nctoolkit
```

Note that recent releases are not available for Python 3.8 on macOS on conda. This issue is being investigated at the minute, and will hopefully be resolved shortly. In the meantime, if you are using macOS and Python 3.8, it is best to install using pip.

At present this can be slow due to the time taken to resolve dependency versions. If you run into problems just use pip.

To install the development version from GitHub:

```
$ pip install git+https://github.com/r4ecology/nctoolkit.git
```

1.2 Python dependencies

- Python (3.6 or later)
- [numpy](#) (1.14 or later)
- [pandas](#) (0.24 or later)
- [xarray](#) (0.14 or later)
- [netCDF4](#) (1.53 or later)
- [ncplot](#)

1.3 System dependencies

There are two main system dependencies: [Climate Data Operators](#), and [NCO](#). The easiest way to install them is using conda:

```
$ conda install -c conda-forge cdo
$ conda install -c conda-forge nco
```

CDO is necessary for the package to work. NCO is an optional dependency and does not have to be installed.

If you want to install CDO from source, you can use one of the bash scripts available [here](#).

INTRODUCTION TO NCTOOLKIT

nctoolkit is a multi-purpose tool for analyzing and post-processing netCDF files. It is designed to carry out almost all analysis and post-processing chains, and to do so easily and efficiently. It is designed explicitly with climate change and oceanographic work in mind. Under the hood, it uses [Climate Data Operators](#) (CDO), but it operates as a stand-alone package with no knowledge of CDO being required to use it.

Let's look at what it can do using a historical global dataset of sea surface temperature, which you can find [here](#).

The preferred way to import nctoolkit is:

```
[1]: import nctoolkit as nc
```

```
nctoolkit is using Climate Data Operators version 1.9.10
```

2.1 It let's you quickly visualize data

nctoolkit offers plotting functionality that will let you automatically plot data from almost any type of netCDF file. It's as simple as the following, which calculates mean historical sea surface temperature and then plots it:

```
[2]: ds = nc.open_data("sst.mon.mean.nc")  
ds.plot()
```

Data type cannot be displayed: application/javascript, application/vnd.holoviews_load.v0+json

Data type cannot be displayed: application/javascript, application/vnd.holoviews_load.v0+json

Data type cannot be displayed: application/javascript, application/vnd.holoviews_load.v0+json

Data type cannot be displayed: application/javascript, application/vnd.holoviews_load.v0+json

```
[2]: :DynamicMap    [time]  
      :Overlay  
      .Image.I      :Image    [lon,lat]    (sst)  
      .Coastline.I  :Feature   [Longitude, Latitude]
```

2.2 It let's you calculate spatial averages

Calculating the spatial mean

```
[3]: ds = nc.open_data("sst.mon.mean.nc")
      ds.spatial_mean()
      ds.plot()
```

```
[3]: :DynamicMap    [variable]
      :Curve       [time]    (value)
```

2.3 It let's you do mathematical operations

nctoolkit offers an ‘assign’ method for performing mathematical operations on variables. This works in a way that will be familiar to users of Pandas. The method is illustrated below in a processing chain that works out how much warmer each part of the ocean is than the global mean.

```
[4]: ds = nc.open_data("sst.mon.mean.nc")
      ds.tmean()
      ds.assign(delta = lambda x: x.sst - spatial_mean(x.sst), drop = True)
      ds.plot()
```

```
[4]: :Overlay
      .Image.I      :Image    [lon,lat]    (delta)
      .Coastline.I :Feature   [Longitude, Latitude]
```

2.4 It let's you crop data

We can crop to a specific region using the crop method. To get a region covering most of Europe, we could do this:

```
[5]: ds = nc.open_data("sst.mon.mean.nc")
      ds.crop(lon = [-13, 38], lat = [30, 67])
      ds.plot()
```

```
[5]: :DynamicMap    [time]
      :Overlay
      .Image.I      :Image    [lon,lat]    (sst)
      .Coastline.I :Feature   [Longitude, Latitude]
```

2.5 It let's you regrid data

nctoolkit has built-in methods for regridding data to user-specified grids. One of the most useful is `to_latlon`. This let's you regrid to a regular latlon grid. You just need to specify the extent of the new grid, the resolution and the regridding method.

```
[6]: ds = nc.open_data("sst.mon.mean.nc")
ds.to_latlon(lon = [-13, 38], lat = [30, 67], res = 0.5, method = "nn")
ds.plot()

[6]: :DynamicMap    [time]
      :Overlay
      .Image.I      :Image    [lon,lat]    (sst)
      .Coastline.I  :Feature   [Longitude, Latitude]
```

2.6 It let's you calculate temporal averages

nctoolkit features a suite of methods, beginning with the letter t, that let you calculate temporal statistics. For example, if we wanted to calculate how much sea surface temperature varies each year, we could do this:

```
[7]: ds = nc.open_data("sst.mon.mean.nc")
ds.trange("year")
ds.tmean()
ds.plot()

[7]: :Overlay
      .Image.I      :Image    [lon,lat]    (sst)
      .Coastline.I  :Feature   [Longitude, Latitude]
```

2.7 It let's you calculate anomalies

In an example above we calculated the global mean sea surface temperature every month since 1850. But calculate the anomaly might be more interesting. The code below will calculate the change in global annual mean sea surface temperature since 1850-1969. The window argument let's you calculate it on a rolling basis.

```
[8]: ds = nc.open_data("sst.mon.mean.nc")
ds.spatial_mean()
ds.annual_anomaly(baseline = [1850, 1869], window= 20)
ds.plot("sst")

[8]: :DynamicMap    [variable]
      :Curve       [time]    (value)
```

2.8 It let's you calculate zonal averages

It is easy to calculate zonal averages using nctoolkit. In the example below change in temperature since 1850-1869 in each latitude band is calculated:

```
[9]: ds = nc.open_data("sst.mon.mean.nc")
      ds.annual_anomaly(baseline = [1850, 1869], window= 20)
      ds.zonal_mean()
      ds.plot()
```

```
[9]: :QuadMesh    [time,lat]    (sst)
```

2.9 Getting started with nctoolkit

To get started with nctoolkit it is best to start [here](#), and to consider getting the [cheatsheet](#).

3.1 Release of v0.3.8

Version 0.3.8 was released in October 2021. This is a minor release, focusing on under-the-hood improvements and introducing better handling of files with varying vertical layers.

A method, `vertical_integration` for calculating vertically integrated totals for netCDF data of the likes of oceanic data, where the vertical levels vary spatially, were introduced. `vertical_mean` has been improved and can now calculate vertical mean in cases where the cell thickness varies in space.

`merge_time` is deprecated, and its functionality will be incorporated into `merge`. So, following this release ensemble merging should use `merge`.

`open_url` is now able to handle multiple urls. Previously it could only handle one.

Some under-the-hood improvements have been made to `assign` to ensure that truth statements do not occasionally throw an error.

3.2 Release of v0.3.7

Version 0.3.7 was released in August 2021. This is a minor release.

New mathematical methods for simple operations on variables were added: `abs`, `power`, `square`, `sqrt`, `exp`, `log` and `log10`. These methods match numpy names.

Bug fixes: `assign` previously did not work with `log10`. Now fixed.

`compare_all` was deleted after a period of deprecation.

3.3 Release of v0.3.6

Version 0.3.6 was released in July 2021. This was a minor release.

New methods `ensemble_var` and `ensemble_stdev` were introduced for calculating variance and standard deviation across ensembles. The method `tvvariance` will be deprecated and is now renamed `tvar` for naming consistency.

3.4 Release of v0.3.5

Version 0.3.5 was released in May 2021.

This is a minor release focusing on some under-the-hood improvements in performance and a couple of new methods. It drops support for CDO version 1.9.3, as this is becoming too time-consuming to continue given the increasingly low reward.

A couple of new methods have been added. `distribute` enables files to be split up spatially into equally sized `m` by `n` rectangles. `collect` is the reverse of `distribute`. It will collect distributed data into one file.

In prior releases `assign` calls could not be split over multiple lines. This is now fixed.

There was a bug in previous releases where `regrid` did not work with multi-file datasets. This was due to the enabling of parallel processing with nctoolkit. The issue is now fixed.

The deprecated methods `mutate` and `assign` have now been removed. Variable creation should use `assign`.

3.5 Release of v0.3.4

Version 0.3.3 was released in April 2021.

This was a minor release focusing on performance improvements, removal of deprecated methods and introduction of one new method.

A new method `fill_na` has been introduced that allows missing values to be filled with the distanced weighted average.

The methods `remove_variables` and `cell_areas` have been removed and are replaced permanently by `drop` and `cell_area`.

3.6 Release of v0.3.2

Version 0.3.2 was released in March 2021. This was a quick release to fix a bug causing `to_nc` to not save output in the base directory.

3.7 Release of v0.3.1

Version 0.3.1 was released in March 2021. This is a minor release that includes new methods, under-the-hood improvements and the removal of deprecated methods.

New methods are introduced for identifying the first time step will specific numerical thresholds are first exceeded or fallen below etc: `first_above`, `first_below`, `last_above` and `last_below`. The thresholds are either single numbers or can come from a gridded dataset for grid-cell specific thresholds.

Methods to compare a dataset with another dataset or netCDF file have been added: `gt` and `lt`, which stand for 'greater than' and 'less than'.

Users are be able to recycle the weights calculated when interpolating data. This can enable much faster interpolation of multiple files with the same grid.

The temporal methods replaced by `tmean` etc. have now been removed from the package. So `monthly_mean` etc. can no longer be used.

3.8 Release of v0.3.0

Version 0.3.0 was released in February 2021. This will be a major release introducing major improvements to the package.

A new method `assign` is now available for generating new variables. This replaces the `mutate` and `transmute`, which were place-holder functions in the early releases of nctoolkit until a proper method for creating variables was put in place. `assign` operates in the same way as the `assign` method in Pandas. Users can generate new variables using lambda functions.

A major-change in this release is that evaluation is now lazy by default. The previous default of non-lazy evaluation was designed to make life slightly easier for new users of the package, but it is probably overly annoying for users to have to set evaluation to lazy each time they use the package.

This release features a subtle shift in how datasets work, so that they have consistent list-like properties. Previously, the files in a dataset given by the ``current`` attribute could be both a str or a list, depending on whether there was one or more files in the dataset. This now always gives a list. As a result datasets in nctoolkit have list-like properties, with ``append`` and `remove` methods available for adding and removing files. `remove` is a new method in this release. As before datasets are iterable.

This release will also allow users to run nctoolkit in parallel. Previous releases allowed files in multi-file datasets to be processed in parallel. However, it was not possible to create processing chains and process files in parallel. This is now possible in version thanks to under-the-hood changes in nctoolkit's code base.

Users are now able to add a configuration file, which means global settings do not need to be set in every session or in every script.

DATASETS

nctoolkit works with what it calls datasets. Each dataset is made up of or more netCDF files.

4.1 Opening datasets

There are 3 ways to create a dataset: `open_data`, `open_url` or `open_thredds`.

If the data you want to analyze that is available on your computer use `open_data`. This will accept either a path to a single file or a list of files. It will also accept wildcards. So if you wanted to open all of the files in a folder called data as a dataset, you could do the following:

```
ds = nc.open_data("data/*.nc")
```

If you want to use data that can be downloaded from a url, just use `open_url`. This will download the netCDF files to a temporary folder, and it can then be analyzed.

If you want to analyze data that is available from a thredds server or opendap, then user `open_thredds`. The file paths should end with `.nc`.

4.2 Visualization of datasets

You can visualize the contents of a dataset using the `plot` method. Below, we will plot temperature for January and the North Atlantic:

```
ds = nc.open_thredds("https://psl.noaa.gov/thredds/dodsC/Datasets/COBE/data.mon.ltm.1981-  
↪2010.nc")  
ds.plot()
```

Please note there may be some issues due to bugs in nctoolkit's dependencies that cause problems plotting some data types. If data does not plot, raise an issue [here](#).

4.3 Modifying datasets and lazy evaluation

nctoolkit works by performing operations and then saving the results as either a temporary file or in a file specified by the user. We can illustrate this with the following code. This will select the first time step from a file available over thredds and will plot the results.

```
ds = nc.open_thredds("https://psl.noaa.gov/thredds/dodsC/Datasets/COBE/data.mon.ltm.1981-
→2010.nc")
ds.select(time = 0)
ds.plot()
```

You will notice, once this is done, that the file associated with the dataset is now a temporary file.

```
ds.current
```

This will happen each time nctoolkit carries out an operation. This is potentially an invitation to slow-running code. You do not want to be constantly reading and writing data. Ideally, you want a processing chain which minimizes IO. nctoolkit enables this by allowing method chaining, thanks to the method chaining of its computational back-end CDO.

Let's look at this chain of code:

```
ds = nc.open_thredds("https://psl.noaa.gov/thredds/dodsC/Datasets/COBE/data.mon.ltm.1981-
→2010.nc")
ds.assign(sst = lambda x: x.sst + 273.15)
ds.select(months = 1)
ds.crop(lon = [-80, 20], lat = [30, 70])
ds.spatial_mean()
```

What is potentially wrong with this? It carries out four operations, so we absolutely do not want to create temporary file in each step. So instead of evaluating the operations line by line nctoolkit only evaluates them either when you tell it to or it has to. So in the code example above we have told nctoolkit what to do to that dataset, but have not told it to actually do any of it.

The quickest way to evaluate everything using run. The code above would become:

```
ds = nc.open_thredds("https://psl.noaa.gov/thredds/dodsC/Datasets/COBE/data.mon.ltm.1981-
→2010.nc")
ds.assign(sst = lambda x: x.sst + 273.15)
ds.select(months = 1)
ds.crop(lon = [-80, 20], lat = [30, 70])
ds.spatial_mean()
ds.run()
```

Evaluation is, to use the technical term, lazy within nctoolkit. It only evaluates things until it needs to or is forced to.

This allows us to create efficient processing chain where we read the input file and write to the output file with no intermediate file writing. If, in the example above, we wanted to save the output file, we could do this:

```
ds = nc.open_thredds("https://psl.noaa.gov/thredds/dodsC/Datasets/COBE/data.mon.ltm.1981-
→2010.nc")
ds.select(months = 1)
ds.crop(lon = [-80, 20], lat = [30, 70])
ds.spatial_mean()
ds.to_nc("foo.nc")
```

4.4 List-like behaviour of datasets

If you want to view the files within a dataset view the `current` attribute.

This is a list that gives the file(s) within the dataset. To make processing these files easier nctoolkit features a number of methods similar to lists.

First, datasets are iterable. So, you can loop through each element of a dataset as follows:

You can find out how many files are in a dataset, using `len`:

You can add a new file to a dataset using `append`:

This method also let you add the files from another dataset.

Similarly, you can remove files from a dataset using `remove`:

In line with typical list behaviours, you can also create empty datasets as follows:

This is particularly useful if you need to create an ensemble based on multiple files that need significant processing before being added to the dataset.

4.5 Dataset attributes

We can find out key information about a dataset using its attributes.

If we want to know the variables available in a dataset called `ds`, we would do:

```
ds.variables
```

If we want to know the vertical levels available in the dataset, we use the following.

```
ds.levels
```

If we want to know the files in a dataset, we would do this. nctoolkit works by generating temporary files, so if you have carried out any operations, this will show a list of temporary files.

```
ds.current
```

If we want to find out what times are in the dataset we do this:

```
ds.times
```

If we want to find out what months are in the dataset:

```
ds.months
```

If we want to find out what years are in the dataset:

```
ds.years
```

We can also access the history of operations carried out on the dataset. This will show the operations carried out by nctoolkit's computational back-end CDO:

```
ds.history
```


CHEAT SHEET

A cheat sheet providing a quick 2-page overview of nctoolkit is available [here](#).

PLOTTING

nctoolkit provides automatic plotting of netCDF data in a similar way to the command line tool ncview.

If you have a dataset, simply use the `plot` method to create an interactive plot that matches the data type.

We can illustrate this using a sea surface temperature dataset available [here](#).

Let's start by calculating mean sea surface temperature for the year 2000 and plotting it:

```
[1]: import nctoolkit as nc
ff = "sst.mon.mean.nc"
ds = nc.open_data(ff)
ds.select(year = 2000)
ds.tmean()
ds.plot()
```

nctoolkit is using Climate Data Operators version 1.9.10

Data type cannot be displayed: application/javascript, application/vnd.holoviews_load.v0+json

Data type cannot be displayed: application/javascript, application/vnd.holoviews_load.v0+json

Data type cannot be displayed: application/javascript, application/vnd.holoviews_load.v0+json

Data type cannot be displayed: application/javascript, application/vnd.holoviews_load.v0+json

```
[1]: :Overlay
      .Image.I      :Image    [lon,lat]    (sst)
      .Coastline.I  :Feature   [Longitude, Latitude]
```

We might be interested in the zonal mean. nctoolkit can automatically plot this easily:

```
[2]: ff = "sst.mon.mean.nc"
ds = nc.open_data(ff)
ds.select(year = 2000)
ds.tmean()
```

(continues on next page)

(continued from previous page)

```
ds.zonal_mean()
ds.plot()
```

```
[2]: :DynamicMap    [variable]
      :Curve      [lat]    (value)
```

nctoolkit can also easily handle heat maps. So, we can easily plot the change in zonal mean over time:

```
[3]: ff = "sst.mon.mean.nc"
     ds = nc.open_data(ff)
     ds.zonal_mean()
     ds.annual_anomaly(baseline = [1850, 1869], window = 20)
     ds.plot()
```

```
[3]: :QuadMesh    [time,lat]    (sst)
```

In a similar vein, it can automatically handle time series. Below we plot a time series of global mean sea surface temperature since 1850:

```
[4]: ff = "sst.mon.mean.nc"
     ds = nc.open_data(ff)
     ds.spatial_mean()
     ds.plot()
```

```
[4]: :DynamicMap    [variable]
      :Curve      [time]    (value)
```

6.1 Internal: ncplot

Plotting is carried out using the ncplot package. If you come across any errors, please raise an issue [here](#).

This is a package that aims to deliver easy use. Colour scales for heat map default to a diverging blue-to-red palette when there are positives and negatives and a viridis palette otherwise.

IMPORTING AND EXPORTING DATA

nctoolkit can work with data available on local file systems, urls and over thredds and OPeNDAP.

7.1 Opening single files and ensembles

If you want to import a single netCDF file as a dataset, do the following:

```
import nctoolkit as nc
ds = nc.open_data(infile)
```

The *open_data* function can also import multiple files. This can be done in two ways. If we have a list of files we can do the following:

```
import nctoolkit as nc
ds = nc.open_data(file_list)
```

Alternatively, *open_data* is capable of handling wildcards. So if we have a folder called data, we can import all files in it as follows:

```
import nctoolkit as nc
ds = nc.open_data("data/*.nc")
```

7.2 Opening files from urls/ftp

If we want to work with a file that is available at a url or ftp, we can use the *open_url* function. This will start by downloading the file to a temporary folder, so that it can be analysed.

```
import nctoolkit as nc
ds = nc.open_url(www.foo.nc)
```

7.3 Opening data available over thredds servers or OPeNDAP

If you want to work with data that is available over a thredds server or OPeNDAP, you can use the `open_thredds` method. This will require that the url ends with “.nc”.

```
import nctoolkit as nc
ds = nc.open_thredds(www.foo.nc)
```

7.4 Exporting datasets

nctoolkit has a number of built in methods for exporting data to netCDF, pandas dataframes and xarray datasets.

7.5 Save as a netCDF

The method `to_nc` lets users export a dataset to a netCDF file. If you want this to be a zipped netCDF file use the `zip` method before to `to_nc`. An example of usage is as follows:

```
ds = nc.open_data(infile)
ds.tmean()
ds.zip()
ds.to_nc(outfile)
```

7.6 Convert to xarray Dataset

The method `to_xarray` lets users export a dataset to an xarray dataset. An example of usage is as follows:

```
ds = nc.open_data(infile)
ds.tmean()
xr_ds = ds.to_xarray()
```

7.7 Convert to pandas dataframe

The method `to_dataframe` lets users export a dataset to a pandas dataframe.

```
ds = nc.open_data(infile)
ds.tmean()
df = ds.to_dataframe()
```

TEMPORAL STATISTICS

nctoolkit has a number of built-in methods for calculating temporal statistics, all of which are prefixed with `t`: `tmean`, `tmin`, `tmax`, `trange`, `tpercentile`, `tmedian`, `tvariance`, `tstdev` and `tcumsum`.

These methods allow you to quickly calculate temporal statistics over specified time periods using the `over` argument. By default the methods calculate the value over all time steps available. For example the following will calculate the temporal mean:

```
import nctoolkit as nc
ds = nc.open_data("sst.mon.mean.nc")
ds.tmean()
```

However, you may want to calculate, for example, an annual average. To do this we use `over`. This is a list which tells the function which time periods to average over. For example, the following will calculate an annual average:

```
ds.tmean(["year"])
```

If you are only averaging over one time period, as above, you can simply use a character string:

```
ds.tmean("year")
```

The possible options for `over` are “day”, “month”, “year”, and “season”. In this case “day” stands for day of year, not day of month.

In the example below we are calculating the maximum value in each month of each year in the dataset.

```
ds.tmax(["month", "year"])
```

8.1 Calculating rolling averages

nctoolkit has a range of methods to calculate rolling averages: `rolling_mean`, `rolling_min`, `rolling_max`, `rolling_range` and `rolling_sum`. These methods let you calculate rolling statistics over a specified time window. For example, if you had daily data and you wanted to calculate a rolling weekly mean value, you could do the following:

```
ds.rolling_mean(7)
```

If you wanted to calculate a rolling weekly sum, this would do:

```
ds.rolling_sum(7)
```

8.2 Calculating anomalies

nctoolkit has two methods for calculating anomalies: `annual_anomaly` and `monthly_anomaly`. Both methods require you to specify a baseline period to calculate the anomaly against. They require that you specify a baseline period showing the minimum and maximum years of the climatological period to compare against.

So, if you wanted to calculate the annual anomaly compared with a baseline period of 1950-1969, you would do this:

```
ds.annual_anomaly(baseline = [1950, 1969])
```

By default, the annual anomaly is calculated as the absolute difference between the annual mean in a year and the mean across the baseline period. However, in some cases this is not suitable. Instead you might want the relative change. In that case, you would do the following:

```
ds.annual_anomaly(baseline = [1950, 1969], metric = "relative")
```

You can also smooth out the anomalies, so that they are calculated on a rolling basis. The following will calculate the anomaly using a rolling window of 10 years.

```
ds.annual_anomaly(baseline = [1950, 1969], window = 10)
```

Monthly anomalies are calculated in the same way:

```
ds.monthly_anomaly(baseline = [1950, 1969])
```

Here the anomaly is the difference between the value in each month compared with the mean in that month during the baseline period.

8.3 Calculating climatologies

This means we can easily calculate climatologies. For example the following will calculate a seasonal climatology:

```
ds.tmean("season")
```

These methods allow partial matches for the arguments, which means you do not need to remember the precise argument each time. For example, the following will also calculate a seasonal climatology:

```
ds.tmean("Seas")
```

Calculating a climatological monthly mean would require the following:

```
ds.tmean("month")
```

and daily would be the following:

```
ds.tmean("day")
```

8.4 Calculating climatologies

This means we can easily calculate climatologies. For example the following will calculate a seasonal climatology:

```
ds.tmean("season")
```

8.5 Cumulative sums

We can calculate the cumulative sum as follows:

```
ds.tcumsum()
```

Please note that this can only calculate over all time periods, and does not accept an `over` argument.

SUBSETTING DATA

nctoolkit has many built in methods for subsetting data. The main method is `select`. This let's you select specific variables, years, months, seasons and timesteps.

9.1 Selecting variables

If you want to select specific variables, you would do the following:

```
ds.select(variables = ["var1", "var2"])
```

If you only want to select one variable, you can do this:

```
ds.select(variables = "var1")
```

9.2 Selecting years

If you want to select specific years, you can do the following:

```
ds.select(years = [2000, 2001])
```

Again, if you want a single year the following will work:

```
ds.select(years = 2000)
```

The `select` method allows partial matches for its arguments. So if we want to select the year 2000, the following will work:

```
ds.select(year = 2000)
```

In this case we can also select a range. So the following will work:

```
ds.select(years = range(2000, 2010))
```

9.3 Selecting months

You can select months in the same way as years. The following examples will all do the same thing:

```
ds.select(months = [1,2,3,4])
ds.select(months = range(1,5))
ds.select(mon = [1,2,3,4])
```

9.4 Selecting seasons

You can easily select seasons. For example if you wanted to select winter, you would do the following:

```
ds.select(season = "DJF")
```

9.5 Selecting timesteps

You can select specific timesteps from a dataset in a similar manner. For example if you wanted to select the first two timesteps in a dataset the following two methods will work:

```
ds.select(time = [0,1])
ds.select(time = range(0,2))
```

9.6 Geographic subsetting

If you want to select a geographic subregion of a dataset, you can use `crop`. This method will select all data within a specific longitude/latitude box. You just need to supply the minimum longitude and latitude required. In the example below, a dataset is cropped with longitudes between -80 and 90 and latitudes between 50 and 80:

```
ds.crop(lon = [-80, 90], lat = [50, 80])
```

MANIPULATING VARIABLES

10.1 Creating new variables

Variable creation in nctoolkit can be done using the `assign` method, which works in a similar way to the method available in pandas.

The `assign` method works using lambda functions. Let's say we have a dataset with a variable 'var' and we simply want to add 10 to it and call the new variable 'new'. We would do the following:

```
ds.assign(new = lambda x: x.var + 10)
```

If you are unfamiliar with lambda functions, note that the `x` after lambda signifies that `x` represents the dataset in whatever comes after ':', which is the actual equation to evaluate. The `x.var` term is `var` from the dataset.

By default `assign` keeps the original variables in the dataset. However, we may only want the new variable or variables. In that case you can use the `drop` argument:

```
ds.assign(new = lambda x: x.var+ 10, drop = True)
```

This results in only one variable.

Note that the `assign` method uses kwargs for the lambda functions, so `drop` can be positioned anywhere. So the following will do the same thing

```
ds.assign(new = lambda x: x.var+ 10, drop = True)  
ds.assign(drop = True, new = lambda x: x.var+ 10)
```

At present, `assign` requires that it is written on a single line. So avoid doing something like the following:

```
ds.assign(new = lambda x: x.var+ 10,  
drop = True)
```

The `assign` method will evaluate the lambda functions sent to it for each dataset grid cell for each time step. So every part of the lambda function must evaluate to a number. So the following will work:

```
k = 273.15  
ds.assign(drop = True, sst_k = lambda x: x.sst + k)
```

However, if you set `k` to a string or anything other than a number it will throw an error. For example, this will throw an error:

```
k = "273.15"  
ds.assign(drop = True, sst_k = lambda x: x.sst + k)
```

10.2 Applying mathematical functions to dataset variables

As part of your lambda function you can use a number of standard mathematical functions. These all have the same names as those in numpy: `abs`, `floor`, `ceil`, `sqrt`, `exp`, `log10`, `sin`, `cos`, `tan`, `arcsin`, `arccos` and `arctan`.

For example if you wanted to calculate the ceiling of a variable you could do the following:

```
ds.assign(new = lambda x: ceil(x.old))
```

An example of using logs would be the following:

```
ds.assign(new = lambda x: log10(x.old+1))
```

10.3 Using spatial statistics

The `assign` method carries out its calculations in each time step, and you can access spatial statistics for each time step when generating new variables. A series of functions are available that have the same names as nctoolkit methods for spatial statistics: `spatial_mean`, `spatial_max`, `spatial_min`, `spatial_sum`, `vertical_mean`, `vertical_max`, `vertical_min`, `vertical_sum`, `zonal_mean`, `zonal_max`, `zonal_min` and `zonal_sum`.

An example of the usefulness of these functions would be if you were working with global temperature data and you wanted to map regions that are warmer than average. You could do this by working out the difference between temperature in one location and the global mean:

```
ds.assign(temp_comp = lambda x: x.temperature - spatial_mean(x.temperature), drop = True)
```

You can also do comparisons. In the above case, we instead might simply want to identify regions that are hotter than the global average. In that case we can simply do this:

```
ds.assign(temp_comp = lambda x: x.temperature > spatial_mean(x.temperature), drop = True)
```

Let's say we wanted to map regions which are 3 degrees hotter than average. We could that as follows:

```
ds.assign(temp_comp = lambda x: x.temperature > spatial_mean(x.temperature + 3), drop = True)
```

or like this:

```
ds.assign(temp_comp = lambda x: x.temperature > (spatial_mean(x.temperature)+3), drop = True)
```

Logical operators work in the standard Python way. So if we had a dataset with a variable called 'var' and we wanted to find cells with values between 1 and 10, we could do this:

```
ds.assign(one2ten = lambda x: x.var > 1 & x.var < 10)
```

You can process multiple variables at once using `assign`. Variables will be created in the order given, and variables created by the first lambda function can be used by the next one, and so on. The simple example below shows how this works. First we create a `var1`, which is temperature plus 1. Then `var2`, which is `var1` plus 1. Finally, we calculate the difference between `var1` and `var2`, and this should be 1 everywhere:

```
ds.assign(var1 = lambda x: x.var + 1, var2 = lambda x: x.var1 + 1, diff = lambda x: x.var2 - x.var1)
```

10.4 Functions that work with nctoolkit variables

The following functions can be used on nctoolkit variables as part of lambda functions.

Function	Description	Example
abs	Absolute value	abs(x.sst)
ceiling	Ceiling of variable	ceiling(x.sst -1)
cell_area	Area of grid-cell (m2)	cell_area(x.var)
cos	Trigonometric cosine of variable	cos(x.var)
day	Day of the month of the variable	day(x.var)
exp	Exponential of variable	exp(x.sst)
floor	Floor of variable	floor(x.sst + 8.2)
hour	Hour of the day of the variable	hour(x.var)
isnan	Is variable a missing value/NA?	isnan(x.var)
latitude	Latitude of the grid cell	latitude(x.var)
level	Vertical level of variable.	level(x.var)
log	Natural log of variable	log10(x.sst + 1)
log10	Base log10 of variable	log10(x.sst + 1)
longitude	Longitude of the grid cell	longitude(x.var)
month	Month of the variable	month(x.var)
sin	Trigonometric sine of variable	sin(x.var)
spatial_max	Spatial max of variable at time-step	spatial_max(x.var)
spatial_mean	Spatial mean of variable at time-step	spatial_mean(x.var)
spatial_min	Spatial min of variable at time-step	spatial_min(x.var)
spatial_sum	Spatial sum of variable at time-step	spatial_sum(x.var)
sqrt	Square root of variable	sqrt(x.sst + 273.15)
tan	Trigonometric tangent of variable	tan(x.var)
timestep	Time step of variable. Using Python indexing.	timestep(x.var)
year	Year of the variable	year(x.var)
zonal_max	Zonal max of variable at time-step	zonal_max(x.var)
zonal_mean	Zonal mean of variable at time-step	zonal_mean(x.var)
zonal_min	Zonal min of variable at time-step	zonal_min(x.var)
zonal_sum	Zonal sum of variable at time-step	zonal_sum(x.var)

10.5 Simple mathematical operations on variables

If you want to do simple operations like adding or subtracting numbers from the variables in datasets you can use the add, subtract, divide and multiply methods. For example if you wanted to add 10 to every variable in a dataset, you would do the following:

```
ds.add(10)
```

If you wanted to multiply everything by 10, you would do this:

```
ds.multiply(10)
```

These methods will also let you use other datasets or netCDF files. So, you could add the values in a dataset data2 to a dataset called data1 as follows:

```
ds1.add(ds2)
```

Please note that this will require that the datasets are structured in a way that the operation makes sense. So each dimension in the datasets will either have to be identical, with the exception of when one dataset has a single value for a dimension. So for example if ds2 above has data covering only 1 timestep, but ds1 has multiple timesteps the data from that single time step will be added to all timesteps in ds1. But if the time steps match, then the data from the first time step in ds2 will be added to the data in the first time step in ds1, and the same will happen with the following time steps.

10.6 Simple numerical comparisons

If you want to do something as simple as working out whether the values of the variables in a dataset are greater than zero, you can use the `compare` method. This method accepts a simple comparison formula, which follows Python conventions. For example, if you wanted to figure out if the values in a dataset were greater than zero, you would do the following:

```
ds.compare(">0")
```

If you wanted to know if they were equal to zero you would do this:

```
ds.compare("==0")
```

INTERPOLATION

nctoolkit features built in methods for horizontal and vertical interpolation.

11.1 Interpolating to a set of coordinates

If you want to regrid a dataset to a specified set of coordinates you can `regrid` and a pandas dataframe. The first column of the dataframe should be the longitudes and the second should be latitudes. The example below regrids a sea-surface temperature dataset to a single location with longitude -30 and latitude 50.

```
import nctoolkit as nc
import pandas as pd
ds = nc.open_thredds("https://psl.noaa.gov/thredds/dodsC/Datasets/COBE2/sst.mon.mean.nc")
ds.select(timestep = range(0, 12))
coords = pd.DataFrame({"lon": [-30], "lat": [50]})
ds.regrid(coords)
```

11.1.1 Interpolating to a regular latlon grid

If you want to interpolate to a regular latlon grid, you can use `to_latlon`. `lon` and `lat` specify the minimum and maximum longitudes and latitudes, while `res`, a 2 variable list specifies the resolution. For example, if we wanted to regrid the globe to 0.5 degree north-south by 1 degree east-west resolution, we could do the following:

```
ds = nc.open_thredds("https://psl.noaa.gov/thredds/dodsC/Datasets/COBE2/sst.mon.mean.nc")
ds.select(timestep = 0)
ds.to_latlon(lon = [-79.5, 79.5], lat = [0.75, 89.75], res = [1, 0.5])
```

11.1.2 Interpolating to another dataset's grid

If we are working with two datasets and want to put them on a common grid, we can interpolate one onto the other's grid. We can illustrate this with a dataset of global sea surface temperature. Let's start by regridding the first timestep in this dataset to a regular latlon grid covering the North Atlantic.

```
ds1 = nc.open_thredds("https://psl.noaa.gov/thredds/dodsC/Datasets/COBE2/sst.mon.mean.nc")
ds1.select(timestep = 0)
ds1.to_latlon(lon = [-79.5, 79.5], lat = [-0.75, 89.75], res = [1, 0.5])
```

We can then use this new dataset as the target grid in `regrid`. So

```
ds2 = nc.open_thredds("https://psl.noaa.gov/thredds/dodsC/Datasets/COBE2/sst.mon.mean.nc")
ds2.select(timestep = 0)
ds2.regrid(ds1)
```

This method will also work using netCDF files. So, if you wanted you can also use a path to a netCDF file as the target grid.

11.1.3 How to reuse the weights for regridding

Under the hood nctoolkit regrids data by first generating a weights file. There are situations where you will want to be able to re-use these weights. For example, if you are post-processing a large number of files one after the other. To make this easier nctoolkit let's you recycle the regridding info. This let's you interpolate using either `regrid` or `to_latlon`, but keep the regridding data for future use by `regrid`.

The example below illustrates this. First, we regrid a global dataset to a regular latlon grid covering the North Atlantic, setting the `recycle` argument to `True`.

```
ds = nc.open_thredds("https://psl.noaa.gov/thredds/dodsC/Datasets/COBE2/sst.mon.mean.nc")
ds.select(timestep = 0)
ds.to_latlon(lon = [-79.5, 79.5], lat = [-0.75, 89.75], res = [1, 0.5], recycle = True)
```

We can then use the grid from data for regridding:

```
ds1 = nc.open_thredds("https://psl.noaa.gov/thredds/dodsC/Datasets/COBE2/sst.mon.mean.nc")
ds1.select(timestep = 0)
ds1.regrid(ds)
```

This, of course, requires that the grids in the datasets are consistent. If you want to access the weights and grid files generated, you can do the following:

These files are deleted either when data is deleted or when the Python session is existed.

11.1.4 Resampling

If you want to make data more coarse spatially, just use the `resample_grid` method. This will, for example, let you select every 2nd grid cell in a north-south and east-west direction. This is illustrated in the example below, where a dataset which has spatial resolution of 1 by 1 degrees is coarsened, so that only every 10th cell is selected in a north-south and east-west. In other words it is now a 10 degrees by 10 degrees dataset.

```
ds = nc.open_thredds("https://psl.noaa.gov/thredds/dodsC/Datasets/COBE2/sst.mon.mean.nc")
ds.select(timestep = 0)
ds.resample_grid(10)
```

11.1.5 Vertical interpolation

We can carry out vertical interpolation using the `vertical_interp` method. This is particularly useful for oceanic data. This is illustrated below by interpolating ocean temperatures from NOAA's World Ocean Atlas for January to a depth of 500 metres. The `vertical_interp` method requires a `levels` argument, which is sea-depth in this case.

```
ds = nc.open_thredds("https://data.nodc.noaa.gov/thredds/dodsC/ncei/woa/temperature/A5B7/  
↳1.00/woa18_A5B7_t01_01.nc")  
ds.select(variables="t_an")  
ds.vertical_interp(levels= [500])
```


ENSEMBLE METHODS

12.1 Merging files with different variables

This notebook will outline some general methods for doing comparisons of multiple files. We will work with two different sea surface temperature data sets from NOAA and the Met Office Hadley Centre.

```
[1]: import nctoolkit as nc
import pandas as pd
import xarray as xr
import numpy as np
```

nctoolkit is using CD0 version 1.9.8

Let's start by downloading the files using wget. Uncomment the code below to do this (note: you will need to extract the HadISST dataset):

```
[2]: # ! wget ftp://ftp.cdc.noaa.gov/Datasets/COBE2/sst.mon.mean.nc
# ! wget https://www.metoffice.gov.uk/hadobs/hadisst/data/HadISST_sst.nc.gz
```

The first step is to get the data. We will start by creating two separate datasets for each file.

```
[3]: sst_noaa = nc.open_data("sst.mon.mean.nc")
sst_hadley = nc.open_data("HadISST_sst.nc")
```

We can see that both variables have sea surface temperature labelled as sst. So we will need to change that.

```
[4]: sst_noaa.variables
```

```
[4]: ['sst']
```

```
[5]: sst_hadley.variables
```

```
[5]: ['sst', 'time_bnds']
```

```
[6]: sst_noaa.rename({"sst": "noaa"})
sst_hadley.rename({"sst": "hadley"})
```

The data sets also cover different time periods, and only have overlapping between 1870 and 2018. so we will need to select those years

```
[7]: sst_noaa.select(years = range(1870, 2019))
sst_hadley.select(years = range(1870, 2019))
```

We also have a problem in that there are two horizontal grids in the Hadley Centre file. We can solve this by selecting the sst variable only

```
[8]: sst_hadley.select(variables = "hadley")
```

At this point, the datasets have the same number of time steps and months covered. However, the grids are still a bit different. So we want to unify them by regridding one dataset on to the other's grid. This can be done using regrid, or any grid of your choosing.

```
[9]: sst_noaa.regrid(grid = sst_hadley)
```

We now have two separate datasets. Let's create a new dataset that has both of them, and then merge them. When doing this we need to make sure nas are treated properly. In this case Hadley Centre values not being NAs as they should be, so we need to fix that. The merge method also requires a strict matching criteria for the dates in the merging files. In this case the Hadley Centre and NOAA data sets both give monthly means, but use a different day of the month. So we will set match to ["year", "month"] this will ensure there are no mis-matches

```
[10]: all_sst = nc.merge(sst_noaa, sst_hadley, match = ["year", "month"])
all_sst.set_missing([-9000, - 900])
```

Let's work out what the global mean SST was over the time period. Note that this will not be totally accurate as there are some missing values here and there that might bias things.

```
[11]: all_sst.spatial_mean()
all_sst.tmean("year")
all_sst.rolling_mean(10)
```

```
[12]: all_sst.plot("noaa")
```

Data type cannot be displayed: application/javascript, application/vnd.holoviews_load.v0+json

Data type cannot be displayed: application/javascript, application/vnd.holoviews_load.v0+json

Data type cannot be displayed: application/javascript, application/vnd.holoviews_load.v0+json

Data type cannot be displayed: application/javascript, application/vnd.holoviews_load.v0+json

Data type cannot be displayed: application/javascript, application/vnd.holoviews_load.v0+json

Data type cannot be displayed: application/javascript, application/vnd.holoviews_load.v0+json

```
[12]: :DynamicMap    [variable]
      :Curve      [time]    (value)
```

We can also work out the difference between the two. Here we wil work out the monthly bias per cell. Then calculate the mean global difference per year, and then calculate a rolling 10 year mean.

```
[13]: all_sst = nc.open_data([sst_noaa.current, sst_hadley.current])
all_sst.merge(match = ["year", "month"])
all_sst.transmute({"bias": "hadley-noaa"})
all_sst.set_missing([-9000, - 900])
all_sst.spatial_mean()
all_sst.tmean("year")
all_sst.rolling_mean(10)
all_sst.plot("bias")
```

```
[13]: :DynamicMap    [variable]
      :Curve      [time]    (value)
```

You can see that there is a notable difference at the start of the time series.

12.2 Merging files with different times

TBC

12.3 Ensemble averaging

TBC

PARALLEL PROCESSING

nctoolkit is written to enable rapid processing and analysis of netCDF files, and this includes the ability to process in parallel. Two methods of parallel processing are available. First is the ability to carry out operations on multi-file datasets in parallel. Second is the ability to define a processing chain in nctoolkit, and then use the multiprocessing package to process files in parallel using that chain.

13.1 Parallel processing of multi-file datasets

If you have a multi-file dataset, processing the files within it in parallel is easy. All you need to is the following:

```
nc.options(cores = 6)
```

This will tell nctoolkit to process the files in multi-file datasets in parallel and to use 6 cores when doing so. You can, of course, set the number of cores as high as you want. The only thing nctoolkit will do is limit it to the number of cores on your machine.

13.2 Parallel processing using multiprocessing

A common task is taking a bunch of files in a folder, doing things to them, and then saving a modified version of each file in a new folder. We want to be able to parallelize that, and we can use the multiprocessing package in the usual way.

But first, we need to change the global settings:

```
import nctoolkit as nc
nc.options(parallel = True)
```

This tells nctoolkit that we are about to do something in parallel. This is critical because of the internal workings of nctoolkit. Behind the scenes nctoolkit is constantly creating and deleting temporary files. It manages this process by creating a safe-list, i.e. a list of files in use that should not be deleted. But if you are running in parallel, you are adding to this list in parallel, and this can cause problems. Telling nctoolkit it will be run in parallel tells it to switch to using a type of list that can be safely added to in parallel.

We can use multiprocessing to do the following: take all of the files in folder foo, do a bunch of things to them, then save the results in a new folder:

We start with a function giving a processing chain. There are obviously different ways of doing this, but I like to use a function that takes the input file and output file:

```
def process_chain(infile, outfile):
    ds = nc.open_data(ff)
    ds.assign(tos = lambda x: x.sst + 273.15)
    ds.tmean()
    ds.to_nc(outfile)
```

We now want to loop through all of the files in a folder, apply the function to them and then save the results in a new folder called new:

```
ensemble = nc.create_ensemble("../data/ensemble")
import multiprocessing
pool = multiprocessing.Pool(3)
for ff in ensemble:
    pool.apply_async(process_chain, [ff, ff.replace("ensemble", "new")])
pool.close()
pool.join()
```

The number 3 in this case signifies that 3 cores are to be used.

Please note that if you are working interactively or in a Jupyter notebook, it is best to reset parallel as follows once you have stopped any parallel processing:

```
nc.options(parallel = False)
```

This is because of the effects of manually terminating commands on multiprocessing lists, which nctoolkit uses when in parallel mode.

EXAMPLES

This tutorial runs through a number of example work flows.

14.1 Global sea surface temperature since 1850

This example analyzes a global sea surface temperature dataset, covering the years since 1850. The data is available from the National Oceanic and Atmospheric Administration (NOAA) [here](#).

We are looking at global sea surface temperature since 1850, so an obvious question is how much the oceans have warmed over this time period. We can use nctoolkit's `spatial_mean` method to calculate this:

Once the file is downloaded, we should set it to `ff`:

```
[1]: import nctoolkit as nc
ff = "sst.mon.mean.nc"
```

nctoolkit is using CD0 version 1.9.9

```
[2]: ds = nc.open_data(ff)
ds.spatial_mean()
ds.plot()
```

Data type cannot be displayed: application/javascript, application/vnd.holoviews_load.v0+json

Data type cannot be displayed: application/javascript, application/vnd.holoviews_load.v0+json

Data type cannot be displayed: application/javascript, application/vnd.holoviews_load.v0+json

Data type cannot be displayed: application/javascript, application/vnd.holoviews_load.v0+json

```
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
```

(continues on next page)

(continued from previous page)

```
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
```

```
[2]: :DynamicMap    [variable]
      :Curve      [time]    (value)
```

We can see a clear temperature rise of about 1 degree Celcius. But this is monthly data, so a bit noisy. We can smooth it out by taking an annual mean. In this case we send “year” to `tmean` to tell it to calculate the mean for each year:

```
[3]: ds = nc.open_data(ff)
      ds.tmean(["year"])
      ds.spatial_mean()
      ds.plot("sst")
```

```
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
```

```
[3]: :DynamicMap    [variable]
      :Curve      [time]    (value)
```

That is getting better. But again, we possibly want a rolling average. We can use the `rolling_mean` method to calculate the mean over every 20-year period:

```
[4]: ds = nc.open_data(ff)
      ds.tmean(["year"])
      ds.spatial_mean()
      ds.rolling_mean(20)
      ds.plot("sst")
```

```
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
```

```
[4]: :DynamicMap    [variable]
      :Curve      [time]    (value)
```

We'll finish things off by tweaking this so that we can work out how much temperature has increased since the first 20 years in the time period. For this we can use the `annual_anomaly` method.

```
[5]: ds = nc.open_data(ff)
ds.annual_anomaly(baseline = [1850, 1869], window = 20)
ds.spatial_mean()
ds.plot("sst")
```

WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
 WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
 WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
 WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
 WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
 WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
 WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
 WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
 WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace
 WARNING:param.ParameterizedMetaclass: Use method 'params' via param namespace

```
[5]: :DynamicMap    [variable]
      :Curve      [time]    (value)
```

14.2 More to come....

RANDOM DATA HACKS

nctoolkit features a number of useful methods to tweak data.

15.1 Shifting time

Sometimes the times in datasets are not quite what we want, and we need some way to adjust time. An example of this is when you are missing a year of data, so want to copy data from the prior year and use it. But first you would need to shift the times in that year forward by a year. You can do this with the `shift` method. This lets you shift time forward by a specified number of hours, days, months or years. You just need to supply hours, days, months or years as an argument. So, if you wanted to shift time backward by one year, you would do the following:

```
ds.shift(years = -1)
```

If you wanted to shift time forward by 12 hours, this would do it:

```
ds.shift(hours = 12)
```

Note: this method allows partial matches to the arguments, so you could use hour, day, month or year just as easily.

15.2 Adding cell areas to a dataset

You can add grid cell areas to a dataset as follows:

```
ds.cell_area()
```

By default, this will add the cell area (in square metres) to the dataset. If you want the dataset to only include cell areas you need to set the `join` argument to `False`:

```
ds.cell_area(join = False)
```

Of course, this method will only if it is possible to calculate the areas the grid cells.

15.3 Changing the format of the netCDF files in a dataset

Sometimes you will want to change the format of the files in a dataset. You can do this using the `format` method. This lets you set the format, with the following options:

- `netCDF` = “nc1”
- `netCDF` version 2 (64-bit offset) = “nc2”/”nc”
- `netCDF4` (HDF5) = “nc4”
- `netCDF4-classic` = “nc4c”
- `netCDF` version 5 (64-bit data) = “nc5”

So, if you want to set the format to `netCDF4`, you would do the following:

```
ds.format("nc4")
```

15.4 Getting rid of dimensions with only one value

Sometimes you will have a dataset that has a dimension with only one value, and you might want to get rid of that dimension. For example, you might only have one timestep and keeping it may have no value. Getting rid of that dimension can be done using the `reduce_dims` method. It works as follows:

```
ds.reduce_dims()
```

GLOBAL SETTINGS

nctoolkit let's you set global settings using options.

The most important and recommended to update is to set evaluation to lazy. This can be done as follows:

```
nc.options(lazy = True)
```

This means that commands will only be evaluated when either request them to be or they need to be.

For example, in the code below the 3 specified commands will only be calculated after it is told to `run`. This cuts down on IO, and can result in significant improvements in run time. At present lazy defaults to `False`, but this may change in a future release of nctoolkit.

```
nc.options(lazy = True)
data.tmean()
data.crop(lat = [0,90])
data.spatial_mean()
data.run()
```

If you are working with ensembles, you may want to change the number of cores used for processing multiple files. For example, you can process multiple files in parallel using 6 cores as follows. By default `cores = 1`. Most methods can run in parallel when working with multi-file datasets.

```
nc.options(cores = 6)
```

By default nctoolkit uses the OS's temporary directories when it needs to create temporary files. In most cases this is optimal. Most of the time reading and writing to temporary folders is faster. However, in some cases this may not be a good idea because you may not have enough space in the temporary folder. In this case you can change the directory used for saving temporary files as follows:

```
nc.options(temp_dir = "/foo")
```

16.1 Setting global settings using a configuration file

You may want to set some global settings either permanently or on a project level. You can do this by setting up a configuration file. This should be a plain text file called `.nctoolkitrc` or `nctoolkitrc`. It should be placed in one of two locations: your working directory or your home directory. When nctoolkit is imported, it will look first in your working directory and then in your home directory for a file called `.nctoolkitrc` or `nctoolkitrc`. It will then use the first it finds to change the global settings from the defaults.

The structure of this file is straightforward. For example, if you wanted to set evaluation to lazy and the number of cores used for processing multi-file datasets, you would the following in your configuration file:

lazy : True

cores : 6

The files roughly follow Python dictionary syntax, with the setting and value separate by `:`. Note that unless the setting is specified in the file, the defaults will be used. If you do not provide a configuration file, nctoolkit will use the default settings.

API REFERENCE

17.1 Session options

<code>options(**kwargs)</code>	Define session options.
---------------------------------	-------------------------

17.1.1 nctoolkit.options

`nctoolkit.options(**kwargs)`

Define session options. Set the options in the session. Available options are `thread_safe` and `lazy`. Set `thread_safe = True` if `hdf5` was built to be thread safe. Set `lazy = True` if you want methods to evaluate lazy by default. Set `cores = n`, if you want `nctoolkit` to process the individual files in multi-file datasets in parallel. Note this only applies to multi-file datasets and will not improve performance with single files. Set `temp_dir = "/foo"` if you want to change the temporary directory used by `nctoolkit` to save temporary files.

Parameters `**kwargs` – Define options using key, value pairs.

Examples

If you wanted to process the files in multi-file datasets in parallel with 6 cores, do the following:

```
>>> import nctoolkit as nc
>>> nc.options(cores = 6)
```

If you want to set evaluation to always be lazy do the following:

```
>>> nc.options(lazy = True)
```

If you want `nctoolkit` to store temporary files in a specific directory, do this:

```
>>> nc.options(temp_dir = "/foo")
```

17.2 Opening/copying data

<code>open_data([x, checks])</code>	Read netCDF data as a DataSet object
<code>open_url([x, ftp_details, wait, file_stop])</code>	Read netCDF data from a url as a DataSet object
<code>open_thredds([x, wait, checks])</code>	Read thredds data as a DataSet object
<code>DataSet.copy(self)</code>	Make a deep copy of an DataSet object

17.2.1 nctoolkit.open_data

`nctoolkit.open_data(x=[], checks=False, **kwargs)`

Read netCDF data as a DataSet object

Parameters

- **x** (*str or list*) – A string or list of netCDF files or a single url. The function will check the files exist. If x is not a list, but an iterable it will be converted to a list. If a `*.nc` style wildcard is supplied, `open_data` will use all files available. By default an empty dataset is created, ie. using `open_data()` will create an empty dataset that can then be expanded using `append`.
- **checks** (*boolean*) – Do you want basic checks to ensure cdo can read files?
- ****kwargs** (*kwargs*) – Optional arguments for internal use by `open_thredds` and `open_url`.

Returns open_data

Return type `nctoolkit.DataSet`

Examples

If you want to open a single file as a dataset, do the following:

```
>>> import nctoolkit as nc
>>> ds = nc.open_data("example.nc")
```

If you want to open a list of files as a multi-file dataset, you would do something like this:

```
>>> import nctoolkit as nc
>>> ds = nc.open_data(["file1.nc", "file2.nc", "file3.nc"])
```

If you wanted to open all files in a directory “data” as a multi-file dataset, you can use a wildcard:

```
>>> import nctoolkit as nc
>>> ds = nc.open_data("data/*.nc")
```

17.2.2 nctoolkit.open_url

`nctoolkit.open_url(x=None, ftp_details=None, wait=None, file_stop=None)`

Read netCDF data from a url as a DataSet object

Parameters

- **x** (*str*) – A string with a url. Prior to processing data will be downloaded to a temp folder.
- **ftp_details** (*dict*) – A dictionary giving the user name and password combination for ftp downloads: {"user":user, "password":pass}
- **wait** (*int*) – Time to wait, in seconds, for data to download. A minimum of 3 attempts will be made to download the data.
- **file_stop** (*int*) – Time limit, in minutes, for individual attempts at downloading data. This is useful to get around download freezes.

Returns `open_url`

Return type `nctoolkit.DataSet`

Examples

If you want to open a file available over a url do the following:

```
>>> import nctoolkit as nc
>>> ds = nc.open_url("http://foo.nc")
```

This will download the file as a temporary folder for use in the dataset.

17.2.3 nctoolkit.open_thredds

`nctoolkit.open_thredds(x=None, wait=None, checks=False)`

Read thredds data as a DataSet object

Parameters

- **x** (*str or list*) – A string or list of thredds urls, which must end with .nc.
- **checks** (*boolean*) – Do you want to check if data is available over thredds?
- **wait** (*int*) – Time to wait for thredds server to be checked. Limitless if not supplied.

Returns `open_thredds`

Return type `nctoolkit.DataSet`

Examples

If you want to open a file available over thredds or opendap, do the following:

```
>>> import nctoolkit as nc
>>> ds = nc.open_thredds("http://foo.nc")
```

17.2.4 nctoolkit.DataSet.copy

`DataSet.copy(self)`

Make a deep copy of an DataSet object

17.3 Merging or analyzing multiple datasets

<code>merge(*datasets[, match])</code>	Merge datasets
<code>cor_time([x, y])</code>	Calculate the temporal correlation coefficient between two datasets This will calculate the temporal correlation coefficient, for each time step, between two datasets.
<code>cor_space([x, y])</code>	Calculate the spatial correlation coefficient between two datasets This will calculate the spatial correlation coefficient, for each time step, between two datasets.

17.3.1 nctoolkit.merge

`nctoolkit.merge(*datasets, match=['day', 'year', 'month'])`

Merge datasets

Parameters

- **datasets** (*kwargs*) – Datasets to merge.
- **match** (*list*) – Temporal matching criteria. This is a list which must be made up of a subset of day, year, month. This checks that the datasets have compatible times. For example, if you want to ensure the datasets have the same years, then use `match = ["year"]`.

17.3.2 nctoolkit.cor_time

`nctoolkit.cor_time(x=None, y=None)`

Calculate the temporal correlation coefficient between two datasets This will calculate the temporal correlation coefficient, for each time step, between two datasets. The datasets must either have the same variables or only have one variable.

Parameters

- **x** (*dataset*) – First dataset to use
- **y** (*dataset*) – Second dataset to use

17.3.3 nctoolkit.cor_space

`nctoolkit.cor_space(x=None, y=None)`

Calculate the spatial correlation coefficient between two datasets This will calculate the spatial correlation coefficient, for each time step, between two datasets. The datasets must either have the same variables or only have one variable.

Parameters

- **x** (*dataset*) – First dataset to use
- **y** (*dataset*) – Second dataset to use

17.4 Adding and removing files to a dataset

append

remove

17.4.1 nctoolkit.append

Functions

<code>append(self[, x])</code>	Add new file(s) to a dataset.
<code>remove(self[, x])</code>	Remove file(s) from a dataset

17.4.2 nctoolkit.remove

Functions

<code>nc_remove(ff[, deep])</code>	Method for removing netCDF files.
------------------------------------	-----------------------------------

17.5 Accessing attributes

<code>DataSet.variables</code>	List variables contained in a dataset
<code>DataSet.years</code>	List years contained in a dataset
<code>DataSet.months</code>	List months contained in a dataset
<code>DataSet.times</code>	List times contained in a dataset
<code>DataSet.levels</code>	List levels contained in a dataset
<code>DataSet.size</code>	The size of an object This will print the number of files, total size, and smallest and largest files in an DataSet object.
<code>DataSet.current</code>	The current file or files in the DataSet object
<code>DataSet.history</code>	The history of operations on the DataSet
<code>DataSet.start</code>	The starting file or files of the DataSet object

17.5.1 nctoolkit.DataSet.variables

property `DataSet.variables`
List variables contained in a dataset

17.5.2 nctoolkit.DataSet.years

property `DataSet.years`
List years contained in a dataset

17.5.3 nctoolkit.DataSet.months

property `DataSet.months`
List months contained in a dataset

17.5.4 nctoolkit.DataSet.times

property `DataSet.times`
List times contained in a dataset

17.5.5 nctoolkit.DataSet.levels

property `DataSet.levels`
List levels contained in a dataset

17.5.6 nctoolkit.DataSet.size

property `DataSet.size`
The size of an object This will print the number of files, total size, and smallest and largest files in an DataSet object.

17.5.7 nctoolkit.DataSet.current

property `DataSet.current`
The current file or files in the DataSet object

17.5.8 nctoolkit.DataSet.history

property `DataSet.history`
The history of operations on the DataSet

17.5.9 nctoolkit.DataSet.start

property `DataSet.start`

The starting file or files of the DataSet object

17.6 Plotting

`DataSet.plot(self[, vars])`

17.6.1 nctoolkit.DataSet.plot

`DataSet.plot(self, vars=None)`

17.7 Variable modification

<code>DataSet.assign(self[, drop])</code>	Create new variables Existing columns that are re-assigned will be overwritten. :param drop: Set to True if you want existing variables to be removed once the new ones have been created. Defaults to False.
<code>DataSet.rename(self, newnames)</code>	Rename variables in a dataset
<code>DataSet.set_missing(self[, value])</code>	Set the missing value for a single number or a range
<code>DataSet.sum_all(self[, drop])</code>	Calculate the sum of all variables for each time step

17.7.1 nctoolkit.DataSet.assign

`DataSet.assign(self, drop=False, **kwargs)`

Create new variables Existing columns that are re-assigned will be overwritten. :param drop: Set to True if you want existing variables to be removed once the new ones have been created.

Defaults to False.

should evaluate to a numeric. New variables are calculated for each grid cell and time step.

Parameters ****kwargs** (*dict of {str: callable}*) – New variable names are keywords. All terms in the equation given by the lamda function should evaluate to a numeric. New variables are calculated for each grid cell and time step.

Notes

Operations are carried out in the order give. So if a new variable is created in the first argument, it can then be used in following arguments.

17.7.2 nctoolkit.DataSet.rename

`DataSet.rename(self, newnames)`

Rename variables in a dataset

Parameters **newnames** (*dict*) – Dictionary with key-value pairs being original and new variable names

Examples

If you want to rename a variable x to y, do the following:

```
>>> ds.rename({"x": "y"})
```

17.7.3 nctoolkit.DataSet.set_missing

`DataSet.set_missing(self, value=None)`

Set the missing value for a single number or a range

Parameters **value** (*2 variable list or int/float*) – If int/float is provided, the missing value will be set to that. If a list is provided, values between the two values (inclusive) of the list are set to missing.

17.7.4 nctoolkit.DataSet.sum_all

`DataSet.sum_all(self, drop=True)`

Calculate the sum of all variables for each time step

Parameters **drop** (*boolean*) – Do you want to keep variables?

17.8 netCDF file attribute modification

<code>DataSet.set_longnames(self[, name_dict])</code>	Set the long names of variables
<code>DataSet.set_units(self[, unit_dict])</code>	Set the units for variables

17.8.1 nctoolkit.DataSet.set_longnames

`DataSet.set_longnames(self, name_dict=None)`

Set the long names of variables

Parameters `name_dict` (*dict*) – Dictionary with key, value pairs representing the variable names and their long names

17.8.2 nctoolkit.DataSet.set_units

`DataSet.set_units(self, unit_dict=None)`

Set the units for variables

Parameters `unit_dict` (*dict*) – A dictionary where the key-value pairs are the variables and new units respectively.

17.9 Vertical/level methods

<code>DataSet.surface(self)</code>	Extract the top/surface level from a dataset This extracts the first vertical level from each file in a dataset.
<code>DataSet.bottom(self)</code>	Extract the bottom level from a dataset This extracts the bottom level from each netCDF file.
<code>DataSet.vertical_interp(self[, levels])</code>	Vertically interpolate a dataset based on given vertical levels This is calculated for each time step and grid cell
<code>DataSet.vertical_mean(self[, thickness, ...])</code>	Calculate the depth-averaged mean for each variable This is calculated for each time step and grid cell
<code>DataSet.vertical_min(self)</code>	Calculate the vertical minimum of variable values This is calculated for each time step and grid cell
<code>DataSet.vertical_max(self)</code>	Calculate the vertical maximum of variable values This is calculated for each time step and grid cell
<code>DataSet.vertical_range(self)</code>	Calculate the vertical range of variable values This is calculated for each time step and grid cell
<code>DataSet.vertical_sum(self)</code>	Calculate the vertical sum of variable values This is calculated for each time step and grid cell
<code>DataSet.vertical_integration(self[, ...])</code>	Calculate the vertically integrated sum over the water column This calculates the sum of the variable multiplied by the cell thickness
<code>DataSet.vertical_cumsum(self)</code>	Calculate the vertical sum of variable values This is calculated for each time step and grid cell
<code>DataSet.invert_levels(self)</code>	Invert the levels of 3D variables This is calculated for each time step and grid cell
<code>DataSet.bottom_mask(self)</code>	Create a mask identifying the deepest cell without missing values.

17.9.1 nctoolkit.DataSet.surface

`DataSet.surface(self)`

Extract the top/surface level from a dataset This extracts the first vertical level from each file in a dataset.

Examples

If you wanted to extract the top vertical level of a dataset, do the following:

```
>>> ds.surface()
```

This method is most useful for things like oceanic data, where this method will extract the sea surface.

17.9.2 nctoolkit.DataSet.bottom

`DataSet.bottom(self)`

Extract the bottom level from a dataset This extracts the bottom level from each netCDF file. Please note that for ensembles, it uses the first file to derive the index of the bottom level. Use `bottom_mask` for files when the bottom cell in netCDF files do not represent the actual bottom.

Examples

If you wanted to extract the bottom vertical level of a dataset, do the following:

```
>>> ds.bottom()
```

This method is most useful for things like oceanic model data, where the bottom cell corresponds to the bottom of the ocean.

17.9.3 nctoolkit.DataSet.vertical_interp

`DataSet.vertical_interp(self, levels=None)`

Vertically interpolate a dataset based on given vertical levels This is calculated for each time step and grid cell

Parameters `levels` (*list, int or str*) – list of vertical levels, for example depths for an ocean model, to vertically interpolate to. These must be floats or ints.

Examples

If you wanted to vertically interpolate a dataset to 5 and 10 metres, you would do the following:

```
>>> ds.vertical_interp([5, 10])
```

This method is most useful for things like oceanic data, where you need to interpolate to certain depth levels. It will require that vertical levels are the same in every grid cell.

17.9.4 nctoolkit.DataSet.vertical_mean

`DataSet.vertical_mean(self, thickness=None, depth_range=None)`

Calculate the depth-averaged mean for each variable This is calculated for each time step and grid cell

thickness: str or Dataset Only use when vertical levels vary in space One of: a variable, in the dataset, which contains the variable thicknesses; a .nc file which contains the thicknesses; or a Dataset that contains the thicknesses. Note: the .nc file or Dataset must only contain one variable.

depth_range: list Only use when vertical levels vary in space Set a depth range if desired. Should be of the form [min_depth, max_depth].

Examples

If you wanted to vertical mean of every variable in a dataset, you would do this:

```
>>> ds.vertical_mean()
```

This method will calculate the vertical mean weighted by the thickness of each cell. Note that if cell thickness cannot be derived it will just average the values in each vertical cell.

17.9.5 nctoolkit.DataSet.vertical_min

`DataSet.vertical_min(self)`

Calculate the vertical minimum of variable values This is calculated for each time step and grid cell

Examples

If you wanted to vertical minimum of every variable in a dataset, you would do this:

```
>>> ds.vertical_min()
```

17.9.6 nctoolkit.DataSet.vertical_max

`DataSet.vertical_max(self)`

Calculate the vertical maximum of variable values This is calculated for each time step and grid cell

Examples

If you wanted to vertical maximum of every variable in a dataset, you would do this:

```
>>> ds.vertical_max()
```

17.9.7 nctoolkit.DataSet.vertical_range

`DataSet.vertical_range(self)`

Calculate the vertical range of variable values This is calculated for each time step and grid cell

Examples

If you wanted to range of values across all vertical levels of every variable in a dataset, you would do this:

```
>>> ds.vertical_range()
```

17.9.8 nctoolkit.DataSet.vertical_sum

`DataSet.vertical_sum(self)`

Calculate the vertical sum of variable values This is calculated for each time step and grid cell

Examples

If you wanted to sum of values across all vertical levels of every variable in a dataset, you would do this:

```
>>> ds.vertical_sum()
```

17.9.9 nctoolkit.DataSet.vertical_integration

`DataSet.vertical_integration(self, thickness=None, depth_range=None)`

Calculate the vertically integrated sum over the water column This calculates the sum of the variable multiplied by the cell thickness

Parameters

- **thickness** (*str or Dataset*) – One of: a variable, in the dataset, which contains the variable thicknesses; a .nc file which contains the thicknesses; or a Dataset that contains the thicknesses. Note: the .nc file or Dataset must only contain one variable.
- **depth_range** (*list*) – Set a depth range if desired. Should be of the form [min_depth, max_depth].

Examples

If you wanted to sum of values across all vertical levels of every variable in a dataset, you would do this:

```
>>> ds.vertical_sum()
```

17.9.10 nctoolkit.DataSet.vertical_cumsum

`DataSet.vertical_cumsum(self)`

Calculate the vertical sum of variable values This is calculated for each time step and grid cell

Examples

If you wanted to calculate the cumulative sum of values across all vertical levels of every variable in a dataset, you would do this:

```
>>> ds.vertical_sum()
```

The cumulative sum will be calculated from the first to the last vertical level. For example, in oceanic data it would start at the sea surface.

17.9.11 nctoolkit.DataSet.invert_levels

`DataSet.invert_levels(self)`

Invert the levels of 3D variables This is calculated for each time step and grid cell

Examples

If you wanted to invert the vertical levels, you would do this:

```
>>> ds.invert_levels()
```

17.9.12 nctoolkit.DataSet.bottom_mask

`DataSet.bottom_mask(self)`

Create a mask identifying the deepest cell without missing values. This converts a dataset to a mask identifying which cell represents the bottom, for example the seabed. 1 identifies the deepest cell with non-missing values. Everything else is 0, or missing. At present this method only uses the first available variable from netCDF files, so it may not be suitable for all data

17.10 Rolling methods

<code>DataSet.rolling_mean(self[, window])</code>	Calculate a rolling mean based on a window
<code>DataSet.rolling_min(self[, window])</code>	Calculate a rolling minimum based on a window
<code>DataSet.rolling_max(self[, window])</code>	Calculate a rolling maximum based on a window
<code>DataSet.rolling_sum(self[, window])</code>	Calculate a rolling sum based on a window
<code>DataSet.rolling_range(self[, window])</code>	Calculate a rolling range based on a window

17.10.1 nctoolkit.DataSet.rolling_mean

`DataSet.rolling_mean(self, window=None)`

Calculate a rolling mean based on a window

Parameters = `int (window)` – The size of the window for the calculation of the rolling mean

Examples

If you wanted to calculate a rolling mean with the mean calculated over every 10 time steps, do the following:

```
>>> ds.rolling_mean(10)
```

17.10.2 nctoolkit.DataSet.rolling_min

`DataSet.rolling_min(self, window=None)`

Calculate a rolling minimum based on a window

Parameters = `int (window)` – The size of the window for the calculation of the rolling minimum

Examples

If you wanted to calculate a rolling minimum with the minimum calculated over every 10 time steps, do the following:

```
>>> ds.rolling_min(10)
```

17.10.3 nctoolkit.DataSet.rolling_max

`DataSet.rolling_max(self, window=None)`

Calculate a rolling maximum based on a window

Parameters = `int (window)` – The size of the window for the calculation of the rolling maximum

Examples

If you wanted to calculate a rolling maximum with the maximum calculated over every 10 time steps, do the following:

```
>>> ds.rolling_max(10)
```

17.10.4 nctoolkit.DataSet.rolling_sum

`DataSet.rolling_sum(self, window=None)`

Calculate a rolling sum based on a window

Parameters = `int (window)` – The size of the window for the calculation of the rolling sum

Examples

If you wanted to calculate a rolling sum with the sum calculated over every 10 time steps, do the following:

```
>>> ds.rolling_sum(10)
```

17.10.5 nctoolkit.DataSet.rolling_range

`DataSet.rolling_range(self, window=None)`

Calculate a rolling range based on a window

Parameters = `int (window)` – The size of the window for the calculation of the rolling range

Examples

If you wanted to calculate a rolling range with the range calculated over every 10 time steps, do the following:

```
>>> ds.rolling_range(10)
```

17.11 Evaluation setting

`DataSet.run(self)`

Run all stored commands in a dataset

17.11.1 nctoolkit.DataSet.run

`DataSet.run(self)`

Run all stored commands in a dataset

Examples

If evaluation is lazy and you need to evaluate commands on a dataset, do the following:

```
>>> ds.run()
```

17.12 Cleaning functions

17.13 Ensemble creation

<code>create_ensemble([path, recursive])</code>	Generate an ensemble
---	----------------------

17.13.1 nctoolkit.create_ensemble

`nctoolkit.create_ensemble(path="", recursive=True)`

Generate an ensemble

Parameters

- **path** (*str*) – The directory to search for netCDF files
- **recursive** (*boolean*) – True/False depending on whether you want to search the path recursively. Defaults to True.

Returns A list of files

Return type list

Examples

If you wanted to recursively find all netCDF files available in a directory “data”, you would do this:

```
>>> import nctoolkit as nc
>>> nc.create_ensemble("data")
```

If you wanted to find the files in that directory and ignore subdirectories, you would instead do this:

```
>>> nc.create_ensemble("data", recursive = False)
```

17.14 Arithmetic methods

<i>DataSet.abs(self)</i>	Method to get the absolute value of variables
<i>DataSet.add(self[, x, var])</i>	Add to a dataset This will add a constant, another dataset or a netCDF file to the dataset. :param x: An int, float, single file dataset or netCDF file to add to the dataset. If a dataset or netCDF file is supplied, this must have only one variable, unless var is provided. The grids must be the same. :type x: int, float, DataSet or netCDF file :param var: A variable in the x to use for the operation :type var: str.
<i>DataSet.assign(self[, drop])</i>	Create new variables Existing columns that are re-assigned will be overwritten. :param drop: Set to True if you want existing variables to be removed once the new ones have been created. Defaults to False.
<i>DataSet.exp(self)</i>	Method to get the exponential of variables
<i>DataSet.log(self)</i>	Method to get the natural log of variables
<i>DataSet.log10(self)</i>	Method to get the base 10 log of variables
<i>DataSet.multiply(self[, x, var])</i>	Multiply a dataset This will multiply a dataset by a constant, another dataset or a netCDF file. :param x: An int, float, single file dataset or netCDF file to multiply the dataset by. If multiplying by a dataset or single file there must only be a single variable in it, unless var is supplied. The grids must be the same. :type x: int, float, DataSet or netCDF file :param var: A variable in the x to multiply the dataset by :type var: str.
<i>DataSet.power(self[, x])</i>	Powers of variables in dataset :param x: An int or float to take the variables to the power of :type x: int, float
<i>DataSet.sqrt(self)</i>	Method to get the square root of variables
<i>DataSet.square(self)</i>	Method to get the square of variables
<i>DataSet.subtract(self[, x, var])</i>	Subtract from a dataset This will subtract a constant, another dataset or a netCDF file from the dataset. :param x: An int, float, single file dataset or netCDF file to subtract from the dataset. If a dataset or netCDF is supplied this must only have one variable, unless var is provided. The grids must be the same. :type x: int, float, DataSet or netCDF file :param var: A variable in the x to use for the operation :type var: str.
<i>DataSet.divide(self[, x, var])</i>	Divide the data This will divide the dataset by a constant, another dataset or a netCDF file. :param x: An int, float, single file dataset or netCDF file to divide the dataset by. If a dataset or netCDF file is supplied, this must have only one variable, unless var is provided. The grids must be the same. :type x: int, float, DataSet or netCDF file :param var: A variable in the x to use for the operation :type var: str.

17.14.1 nctoolkit.DataSet.abs

`DataSet.abs(self)`

Method to get the absolute value of variables

Examples

If you wanted to get the absolute value of each variable, you just need do this:

```
>>> ds.abs()
```

17.14.2 nctoolkit.DataSet.add

`DataSet.add(self, x=None, var=None)`

Add to a dataset This will add a constant, another dataset or a netCDF file to the dataset. :param x: An int, float, single file dataset or netCDF file to add to the dataset.

If a dataset or netCDF file is supplied, this must have only one variable, unless var is provided. The grids must be the same.

Parameters `var (str)` – A variable in the x to use for the operation

Examples

If you wanted to add 10 to all variables in a dataset, you would do the following:

```
>>> ds.add(10)
```

To add the values in a dataset ds2 from a dataset ds1, you would do the following:

```
>>> ds1.add(ds2)
```

Grids in the datasets must match. Addition will occur in matching timesteps in ds1 and ds2. If there is only 1 timestep in ds2, then the data from that timestep will be added to the data in all ds1 time steps.

Adding the data from another netCDF file will work in the same way:

```
>>> ds1.add("example.nc")
```

17.14.3 nctoolkit.DataSet.exp

`DataSet.exp(self)`

Method to get the exponential of variables

Examples

If you wanted to calculate the exponential of a variable, you just need to do this:

```
>>> ds.exp(0.5)
```

17.14.4 nctoolkit.DataSet.log

`DataSet.log(self)`

Method to get the natural log of variables

Examples

If you wanted to calculate the natural log of each variable, you just need to do this:

```
>>> ds.log()
```

17.14.5 nctoolkit.DataSet.log10

`DataSet.log10(self)`

Method to get the base 10 log of variables

Examples

If you wanted to calculate the base 10 log of each variable, you just need to do this:

```
>>> ds.log10()
```

17.14.6 nctoolkit.DataSet.multiply

`DataSet.multiply(self, x=None, var=None)`

Multiply a dataset This will multiply a dataset by a constant, another dataset or a netCDF file. :param x: An int, float, single file dataset or netCDF file to multiply the dataset by.

If multiplying by a dataset or single file there must only be a single variable in it, unless var is supplied. The grids must be the same.

Parameters **var** (*str*) – A variable in the x to multiply the dataset by

Examples

If you wanted to multiply variables in a dataset by 10, you would do the following:

```
>>> ds.multiply(10)
```

To multiply the values in a dataset by the values of variables in dataset ds2, you would do the following:

```
>>> ds1.multiply(ds2)
```

Grids in the datasets must match. Multiplication will occur in matching timesteps in ds1 and ds2. If there is only 1 timestep in ds2, then the data from that timestep in ds2 will multiply the data in all timesteps in ds1.

Multiplying a dataset by the data from another netCDF file will work in the same way:

```
>>> ds.multiply("example.nc")
```

17.14.7 nctoolkit.DataSet.power

`DataSet.power(self, x=None)`

Powers of variables in dataset :param x: An int or float to take the variables to the power of :type x: int, float

Examples

If you wanted to take each variable to the power of 0.5 you would do this:

```
>>> ds.power(0.5)
```

17.14.8 nctoolkit.DataSet.sqrt

`DataSet.sqrt(self)`

Method to get the square root of variables

Examples

If you wanted to calculate the square root of each variable, you just need to do this:

```
>>> ds.sqrt()
```

17.14.9 nctoolkit.DataSet.square

`DataSet.square(self)`

Method to get the square of variables

Examples

If you wanted to calculate the square of each variable, you just need to do this:

```
>>> ds.power()
```

17.14.10 nctoolkit.DataSet.subtract

DataSet.subtract(*self*, *x=None*, *var=None*)

Subtract from a dataset This will subtract a constant, another dataset or a netCDF file from the dataset. :param x: An int, float, single file dataset or netCDF file to subtract from the dataset.

If a dataset or netCDF is supplied this must only have one variable, unless var is provided. The grids must be the same.

Parameters **var** (*str*) – A variable in the x to use for the operation

Examples

If you wanted to subtract 10 from all variables in a dataset, you would do the following:

```
>>> ds.subtract(10)
```

To subtract the values in a dataset ds2 from those in a dataset ds1, you would do the following:

```
>>> ds1.subtract(ds2)
```

Grids in the datasets must match. Division will occur in matching timesteps in ds1 and ds2 if there are matching timesteps. If there is only 1 timestep in ds2, then the data from that timestep in ds2 will be subtracted from the data in all timesteps in ds1.

Subtracting of the data from another netCDF file will work in the same way:

```
>>> ds1.subtract("example.nc")
```

17.14.11 nctoolkit.DataSet.divide

DataSet.divide(*self*, *x=None*, *var=None*)

Divide the data This will divide the dataset by a constant, another dataset or a netCDF file. :param x: An int, float, single file dataset or netCDF file to divide the dataset by.

If a dataset or netCDF file is supplied, this must have only one variable, unless var is provided. The grids must be the same.

Parameters **var** (*str*) – A variable in the x to use for the operation

Examples

If you wanted to divide all variables in a dataset by 20, you would do the following:

```
>>> ds.divide(10)
```

To divide values in a dataset by those in the dataset ds2 from a dataset ds1, you would do the following:

```
>>> ds1.divide(ds2)
```

Grids in the datasets must match. Division will occur in matching timesteps in ds1 and ds2. If there is only 1 timestep in ds2, then the data from that timestep in ds2 will divided the data in all ds1 time steps.

Adding the data from another netCDF file will work in the same way:

```
>>> ds.divide("example.nc")
```

17.15 Ensemble statistics

<code>DataSet.ensemble_mean(self[, nco, ignore_time])</code>	Calculate an ensemble mean
<code>DataSet.ensemble_min(self[, nco, ignore_time])</code>	Calculate an ensemble min
<code>DataSet.ensemble_max(self[, nco, ignore_time])</code>	Calculate an ensemble maximum
<code>DataSet.ensemble_percentile(self[, p])</code>	Calculate an ensemble percentile This will calculate the percentiles for each time step in the files.
<code>DataSet.ensemble_range(self)</code>	Calculate an ensemble range The range is calculated for each time step; for example, if each file in the ensemble has 12 months of data the statistic will be calculated for each month.
<code>DataSet.ensemble_stdev(self)</code>	Calculate an ensemble standard deviation
<code>DataSet.ensemble_sum(self)</code>	Calculate an ensemble sum The sum is calculated for each time step; for example, if each file in the ensemble has 12 months of data the statistic will be calculated for each month.
<code>DataSet.ensemble_var(self)</code>	Calculate an ensemble variance

17.15.1 nctoolkit.DataSet.ensemble_mean

`DataSet.ensemble_mean(self, nco=False, ignore_time=False)`

Calculate an ensemble mean

Parameters

- **nco** (*boolean*) – Do you want to use NCO for the calculation? Default is False, i.e. CDO is used. Modify default if run time is an issue.
- **ignore_time** (*boolean*) – If True the mean is calculated over all time steps. If False, the ensemble mean is calculated for each time steps; for example, if the ensemble is made up of monthly files the mean for each month will be calculated.

17.15.2 nctoolkit.DataSet.ensemble_min

`DataSet.ensemble_min(self, nco=False, ignore_time=False)`

Calculate an ensemble min

Parameters

- **nco** (*boolean*) – Do you want to use NCO for the calculation? Default is False, i.e. CDO is used. Modify default if run time is an issue.
- **ignore_time** (*boolean*) – If True the min is calculated over all time steps. If False, the ensemble min is calculated for each time steps; for example, if the ensemble is made up of monthly files the min for each month will be calculated.

17.15.3 nctoolkit.DataSet.ensemble_max

`DataSet.ensemble_max(self, nco=False, ignore_time=False)`

Calculate an ensemble maximum

Parameters

- **nco** (*boolean*) – Do you want to use NCO for the calculation? Default is False, i.e. CDO is used. Modify default if run time is an issue.
- **ignore_time** (*boolean*) – If True the max is calculated over all time steps. If False, the ensemble max is calculated for each time steps; for example, if the ensemble is made up of monthly files the max for each month will be calculated.

17.15.4 nctoolkit.DataSet.ensemble_percentile

`DataSet.ensemble_percentile(self, p=None)`

Calculate an ensemble percentile This will calculate the percentiles for each time step in the files. For example, if you had an ensemble of files where each file included 12 months of data, it would calculate the percentile for each month.

Parameters *p* (*float* or *int*) – percentile to calculate. $0 \leq p \leq 100$.

17.15.5 nctoolkit.DataSet.ensemble_range

`DataSet.ensemble_range(self)`

Calculate an ensemble range The range is calculated for each time step; for example, if each file in the ensemble has 12 months of data the statistic will be calculated for each month.

17.15.6 nctoolkit.DataSet.ensemble_stdev

`DataSet.ensemble_stdev(self)`

Calculate an ensemble standard deviation

The ensemble standard deviation is calculated for each time steps; for example, if the ensemble is made up of monthly files the standard deviation for each month will be calculated.

17.15.7 nctoolkit.DataSet.ensemble_sum

`DataSet.ensemble_sum(self)`

Calculate an ensemble sum The sum is calculated for each time step; for example, if each file in the ensemble has 12 months of data the statistic will be calculated for each month.

17.15.8 nctoolkit.DataSet.ensemble_var

`DataSet.ensemble_var(self)`

Calculate an ensemble variance

The ensemble variance is calculated for each time steps; for example, if the ensemble is made up of monthly files the standard deviation for each month will be calculated.

17.16 Subsetting operations

<code>DataSet.crop(self[, lon, lat, nco, nco_vars])</code>	Crop to a rectangular longitude and latitude box
<code>DataSet.select(self, **kwargs)</code>	A method for subsetting datasets to specific variables, years, longitudes etc.
<code>DataSet.drop(self[, vars])</code>	Remove variables This will remove stated variables from files in the dataset.

17.16.1 nctoolkit.DataSet.crop

`DataSet.crop(self, lon=[- 180, 180], lat=[- 90, 90], nco=False, nco_vars=None)`

Crop to a rectangular longitude and latitude box

Parameters

- **lon** (*list*) – The longitude range to select. This must be two variables, between -180 and 180 when `nco = False`.
- **lat** (*list*) – The latitude range to select. This must be two variables, between -90 and 90 when `nco = False`.
- **nco** (*boolean*) – Do you want this to use NCO for cropping? Defaults to `False`, and uses CDO. Set to `True` if you want to call NCO. NCO is typically better at handling very large horizontal grids.
- **nco_vars** (*str or list*) – If using NCO, the variables you want to select

Examples

If you wanted to crop a dataset to longitudes between -40 and 30 and latitudes between -10 and 40, you would do the following:

```
>>> ds.crop(lon = [-40, 30], lat = [-10, 40])
```

If you wanted to select only the northern hemisphere, the following will work:

```
>>> ds.crop(lat = [0, 90])
```

17.16.2 nctoolkit.DataSet.select

`DataSet.select(self, **kwargs)`

A method for subsetting datasets to specific variables, years, longitudes etc. Operations are applied in the order supplied.

Parameters ***kwargs** – Possible arguments: variables, years, months, seasons, timesteps, lon, lat

Note: this uses partial matches. So year, month, var etc. will also work

Each kwarg works as follows:

variables [str or list] A variable or list of variables to select

seasons [str] Seasons to select. One of “DJF”, “MAM”, “JJA”, “SON”.

months [list, range or int] Month(s) to select.

years [list, range or int] Years(s) to select. These should be integers

timesteps [list or int] time step(s) to select. For example, if you wanted the first time step set times=0.

Examples

If you want to select a single variable do the following:

```
>>> ds.select(variable = "var")
```

If you want to select a list of variables, do this:

```
>>> ds.select(variable = ["var1", "var2"])
```

If you want to select data for January, do the following:

```
>>> ds.select(month = 1)
```

If you want to select a range of months, do the following:

```
>>> ds.select(months = range(1, 7))
```

If you want to select a range of years, for example the 2010s, do the following:

```
>>> ds.select(years = range(2010, 2020))
```

If you want to select the first two timesteps in a dataset, do the following:

```
>>> ds.select(timesteps = [0,1])
```

17.16.3 nctoolkit.DataSet.drop

`DataSet.drop(self, vars=None)`

Remove variables This will remove stated variables from files in the dataset.

Parameters `vars` (*str or list*) – Variable or variables to be removed from the dataset. Variables that are listed but not in the dataset will be ignored

Examples

If you wanted to remove a single variable ‘var1’ from a dataset data, you would do the following:

```
>>> ds.drop('var')
```

If you wanted to remove a list of variables, you would do the following:

```
>>> ds.drop(['var1', 'var2', 'var2'])
```

17.17 Time-based methods

`DataSet.set_date(self[, year, month, day, ...])`

Set the date in a dataset You should only do this if you have to fix/change a dataset with a single, not multiple dates.

`DataSet.shift(self, **kwargs)`

Shift method.

17.17.1 nctoolkit.DataSet.set_date

`DataSet.set_date(self, year=None, month=None, day=None, base_year=1900)`

Set the date in a dataset You should only do this if you have to fix/change a dataset with a single, not multiple dates.

Parameters

- **year** (*int*) – The year
- **month** (*int*) – The month
- **day** (*int*) – The day
- **base_year** (*int*) – The base year for time creation in the netCDF. Defaults to 1900.

17.17.2 nctoolkit.DataSet.shift

`DataSet.shift(self, **kwargs)`

Shift method. A wrapper for shift_days, shift_hours Operations are applied in the order supplied.

Parameters ***kwargs** – hours maps to shift_hours days maps to shift_days months maps to shift_months years maps to shift_years

Note: this uses partial matches. So hour, day, month, year will also work.

Examples

If you wanted to shift all times back 1 hour, you would do the following:

```
>>> ds.shift(hours = -1)
```

If you wanted to shift all times forward 2 days, you would do the following:

```
>>> ds.shift(days = 2)
```

If you wanted to shift all times forward 6 months, you would do the following:

```
>>> ds.shift(months = 6)
```

If you wanted to shift all times forward 1 year, you would do the following:

```
>>> ds.shift(years = 1)
```

This method will allow partial matches in arguments. So the following will do the same thing:

```
>>> ds.shift(year = 2)
```

```
>>> ds.shift(years = 2)
```

17.18 Interpolation and resampling methods

<code>DataSet.regrid(self[, grid, method, recycle])</code>	Regrid a dataset to a target grid
<code>DataSet.to_latlon(self[, lon, lat, res, ...])</code>	Regrid a dataset to a regular latlon grid
<code>DataSet.resample_grid(self[, factor])</code>	Resample the horizontal grid of a dataset
<code>DataSet.time_interp(self[, start, end, ...])</code>	Temporally interpolate variables based on date range and time resolution
<code>DataSet.timestep_interp(self[, steps])</code>	Temporally interpolate a dataset to given number of time steps between existing time steps
<code>DataSet.fill_na(self[, n])</code>	Fill missing values with distance-weighted average.

17.18.1 nctoolkit.DataSet.regrid

`DataSet.regrid(self, grid=None, method='bil', recycle=False)`

Regrid a dataset to a target grid

Parameters

- **grid** (*nctoolkit.DataSet*, *pandas data frame* or *netCDF file*) – The grid to remap to
- **method** (*str*) – Remapping method. Defaults to “bil”. Methods available are: bilinear - “bil”; nearest neighbour - “nn” - “nearest neighbour” bicubic interpolation - “bic” Distance-weighted average - “dis” First order conservative remapping - “con” Second order conservative remapping - “con2” Large area fraction remapping - “laf”

17.18.2 nctoolkit.DataSet.to_latlon

`DataSet.to_latlon(self, lon=None, lat=None, res=None, method='bil', recycle=False)`

Regrid a dataset to a regular latlon grid

Parameters

- **lon** (*list*) – 2 element list giving minimum and maximum longitude of target grid
- **lat** (*list*) – 2 element list giving minimum and maximum latitude of target grid
- **res** (*float, int or list*) – If float or int given, this will be the horizontal and vertical resolution of the target grid. If 2 element list is given, the first element is the longitudinal resolution and the second is the latitudinal resolution.
- **method** (*str*) – Remapping method. Defaults to “bil”. Methods available are: bilinear - “bil”; nearest neighbour - “nn” - “nearest neighbour” bicubic interpolation - “bic” Distance-weighted average - “dis” First order conservative remapping - “con” Second order conservative remapping - “con2” Large area fraction remapping - “laf”
- **recycle** (*bool*) – Do you want the grid and weights to be available for recycling and use in regrid? Defaults to False

17.18.3 nctoolkit.DataSet.resample_grid

`DataSet.resample_grid(self, factor=None)`

Resample the horizontal grid of a dataset

Parameters **factor** (*int*) – The resampling factor. Must be a positive integer. No interpolation occurs. Example: factor of 2 will sample every other grid cell

Examples

If you wanted to select every other grid cell, you could do the following:

```
>>> ds.resample_grid(2)
```

17.18.4 nctoolkit.DataSet.time_interp

`DataSet.time_interp(self, start=None, end=None, resolution='monthly')`

Temporally interpolate variables based on date range and time resolution

Parameters

- **start** (*str*) – Start date for interpolation. Needs to be of the form YYYY/MM/DD or YYYY-MM-DD.
- **end** (*str*) – End date for interpolation. Needs to be of the form YYYY/MM/DD or YYYY-MM-DD. If end is not given interpolation will be to the final available time in the dataset.
- **resolution** (*str*) – Time steps used for interpolation. Needs to be “daily”, “weekly”, “monthly” or “yearly”. Defaults to monthly.

17.18.5 nctoolkit.DataSet.timestep_interp

`DataSet.timestep_interp(self, steps=None)`

Temporally interpolate a dataset to given number of time steps between existing time steps

Parameters `steps` (*int*) – Number of time steps to interpolate between existing time steps. For example, if you wanted to go from daily to hourly data you would set `steps=24`.

17.18.6 nctoolkit.DataSet.fill_na

`DataSet.fill_na(self, n=1)`

Fill missing values with distance-weighted average. :param n: Number of nearest neighbours to use. Defaults to 1. :type n: int

17.19 Masking methods

`DataSet.mask_box(self[, lon, lat])`

Mask a lon/lat box

17.19.1 nctoolkit.DataSet.mask_box

`DataSet.mask_box(self, lon=[- 180, 180], lat=[- 90, 90])`

Mask a lon/lat box

Parameters

- `lon` (*list*) – Longitude range to mask. Must be of the form: `[lon_min, lon_max]`
- `lat` (*list*) – Latitude range to mask. Must be of the form: `[lat_min, lat_max]`

17.20 Anomaly methods

`DataSet.annual_anomaly(self[, baseline, ...])`

Calculate annual anomalies for each variable based on a baseline period The anomaly is derived by first calculating the climatological annual mean for the given baseline period.

`DataSet.monthly_anomaly(self[, baseline])`

Calculate monthly anomalies based on a baseline period The anomaly is derived by first calculating the climatological monthly mean for the given baseline period.

17.20.1 nctoolkit.DataSet.annual_anomaly

`DataSet.annual_anomaly(self, baseline=None, metric='absolute', window=1)`

Calculate annual anomalies for each variable based on a baseline period. The anomaly is derived by first calculating the climatological annual mean for the given baseline period. Annual means are then calculated for each year and the anomaly is calculated compared with the baseline mean. This will be calculated on a per-file basis in a multi-file dataset.

Parameters

- **baseline** (*list*) – Baseline years. This needs to be the first and last year of the climatological period. Example: a baseline of [1980,1999] will result in anomalies against the 20 year climatology from 1980 to 1999.
- **metric** (*str*) – Set to “absolute” or “relative”, depending on whether you want the absolute or relative anomaly to be calculated.
- **window** (*int*) – A window for the anomaly. By default window = 1, i.e. the annual anomaly is calculated. If, for example, window = 20, the 20 year rolling means will be used to calculate the anomalies.

Examples

If you wanted to calculate an annual anomaly where values are compared with the mean for the years 1950-1969, you would do this:

```
>>> ds.annual_anomaly(baseline = [1950, 1969])
```

By default, this results in the absolute difference to be used. If you wanted the anomaly to be calculated relative to the baseline mean, you would do this:

```
>>> ds.annual_anomaly(baseline = [1950, 1969], metric = "relative")
```

You might want to smooth out the anomalies, so that you are looking at rolling averages. In that case you can supply a windows. So if you wanted to calculate the anomaly using a rolling average with a 10 year window, you would do this:

```
>>> ds.annual_anomaly(baseline = [1950, 1969], window = 10)
```

17.20.2 nctoolkit.DataSet.monthly_anomaly

`DataSet.monthly_anomaly(self, baseline=None)`

Calculate monthly anomalies based on a baseline period. The anomaly is derived by first calculating the climatological monthly mean for the given baseline period. Monthly means are then calculated for each year and the anomaly is calculated compared with the baseline mean. This is calculated separately for each file in a multi-file dataset.

Parameters **baseline** (*list*) – Baseline years. This needs to be the first and last year of the climatological period. Example: a baseline of [1985,2005] will result in anomalies against 20 year climatology from 1986 to 2005.

Examples

If you wanted to calculate a monthly anomaly where values are compared with the climatological monthly mean for the years 1950-1969, you would do this:

```
>>> ds.monthly_anomaly(baseline = [1950, 1969])
```

17.21 Statistical methods

<i>DataSet.tmean</i> (self[, over])	Calculate the temporal mean of all variables
<i>DataSet.tmin</i> (self[, over])	Calculate the temporal minimum of all variables
<i>DataSet.tmedian</i> (self[, over])	Calculate the temporal median of all variables :param over: Time periods to average over.
<i>DataSet.tpercentile</i> (self[, p, over])	Calculate the temporal percentile of all variables
<i>DataSet.tmax</i> (self[, over])	Calculate the temporal maximum of all variables
<i>DataSet.tsum</i> (self[, over])	Calculate the temporal sum of all variables
<i>DataSet.trange</i> (self[, over])	Calculate the temporal range of all variables
<i>DataSet.tvariance</i> (self[, over])	Calculate the temporal variance of all variables
<i>DataSet.tstdev</i> (self[, over])	Calculate the temporal standard deviation of all variables
<i>DataSet.tcumsum</i> (self)	Calculate the temporal cumulative sum of all variables
<i>DataSet.tvar</i> (self[, over])	Calculate the temporal variance of all variables
<i>DataSet.cor_space</i> (self[, var1, var2])	Calculate the correlation correct between two variables in space This is calculated for each time step.
<i>DataSet.cor_time</i> (self[, var1, var2])	Calculate the correlation correct in time between two variables The correlation is calculated for each grid cell, ignoring missing values.
<i>DataSet.spatial_mean</i> (self)	Calculate the area weighted spatial mean for all variables This is performed for each time step.
<i>DataSet.spatial_min</i> (self)	Calculate the spatial minimum for all variables This is performed for each time step.
<i>DataSet.spatial_max</i> (self)	Calculate the spatial maximum for all variables This is performed for each time step.
<i>DataSet.spatial_percentile</i> (self[, p])	Calculate the spatial sum for all variables This is performed for each time step.
<i>DataSet.spatial_range</i> (self)	Calculate the spatial range for all variables This is performed for each time step.
<i>DataSet.spatial_sum</i> (self[, by_area])	Calculate the spatial sum for all variables This is performed for each time step.
<i>DataSet.centre</i> (self[, by, by_area])	Calculate the latitudinal or longitudinal centre for each year/month combination in files. This applies to each file in an ensemble. by : str Set to 'latitude' if you want the latitudinal centre calculated. 'longitude' for longitudinal. by_area : bool If the variable is a value/m2 type variable, set to True, otherwise set to False.
<i>DataSet.zonal_mean</i> (self)	Calculate the zonal mean for each year/month combination in files.
<i>DataSet.zonal_min</i> (self)	Calculate the zonal minimum for each year/month combination in files.

continues on next page

Table 23 – continued from previous page

<code>DataSet.zonal_max(self)</code>	Calculate the zonal maximum for each year/month combination in files.
<code>DataSet.zonal_range(self)</code>	Calculate the zonal range for each year/month combination in files.
<code>DataSet.meridional_mean(self)</code>	Calculate the meridional mean for each year/month combination in files.
<code>DataSet.meridional_min(self)</code>	Calculate the meridional minimum for each year/month combination in files.
<code>DataSet.meridional_max(self)</code>	Calculate the meridional maximum for each year/month combination in files.
<code>DataSet.meridional_range(self)</code>	Calculate the meridional range for each year/month combination in files.

17.21.1 nctoolkit.DataSet.tmean

`DataSet.tmean(self, over='time')`

Calculate the temporal mean of all variables

Parameters `over` (*str* or *list*) – Time periods to average over. Options are ‘year’, ‘month’, ‘day’.

Examples

If you want to calculate mean over all time steps. Do the following:

```
>>> ds.tmean()
```

If you want to calculate the mean for each year in a dataset, do this:

```
>>> ds.tmean("year")
```

If you want to calculate the mean for each month in a dataset, do this:

```
>>> ds.tmean("month")
```

If you want to calculate the mean for each month in each year in a dataset, do this:

```
>>> ds.tmean(["year", "month"])
```

This method will also let you easily calculate climatologies. So, if you wanted to calculate a monthly climatological mean, you would do this:

```
>>> ds.tmean("month")
```

A daily climatological mean would be the following:

```
>>> ds.tmean("day")
```

17.21.2 nctoolkit.DataSet.tmin

`DataSet.tmin(self, over='time')`

Calculate the temporal minimum of all variables

Parameters `over` (*str* or *list*) – Time periods to average over. Options are ‘year’, ‘month’, ‘day’.

Examples

If you want to calculate minimum over all time steps. Do the following:

```
>>> ds.tmin()
```

If you want to calculate the minimum for each year in a dataset, do this:

```
>>> ds.tmin("year")
```

If you want to calculate the minimum for each month in a dataset, do this:

```
>>> ds.tmin("month")
```

If you want to calculate the minimum for each month in each year in a dataset, do this:

```
>>> ds.tmin(["year", "month"])
```

This method will also let you easily calculate climatologies. So, if you wanted to calculate a monthly climatological min, you would do this:

```
>>> ds.tmin("month")
```

A daily climatological minimum would be the following:

```
>>> ds.tmin("day")
```

17.21.3 nctoolkit.DataSet.tmedian

`DataSet.tmedian(self, over='time')`

Calculate the temporal median of all variables :param over: Time periods to average over. Options are ‘year’, ‘month’, ‘day’. :type over: str or list

Examples

If you want to calculate median over all time steps. Do the following:

```
>>> ds.tmedian()
```

If you want to calculate the median for each year in a dataset, do this:

```
>>> ds.tmedian("year")
```

If you want to calculate the median for each month in a dataset, do this:

```
>>> ds.tmedian("month")
```

If you want to calculate the median for each month in each year in a dataset, do this:

```
>>> ds.tmedian(["year", "month"])
```

This method will also let you easily calculate climatologies. So, if you wanted to calculate a monthly climatological median, you would do this:

```
>>> ds.tmedian("month")
```

A daily climatological median would be the following:

```
>>> ds.tmedian("day")
```

17.21.4 nctoolkit.DataSet.tpercentile

`DataSet.tpercentile(self, p=None, over='time')`

Calculate the temporal percentile of all variables

Parameters `p` (*float or int*) – Percentile to calculate

Examples

If you want to calculate the 20th percentile over all time steps. Do the following:

```
>>> ds.tpercentile(20)
```

If you want to calculate the 20th percentile for each year in a dataset, do this:

```
>>> ds.tpercentile(20)
```

17.21.5 nctoolkit.DataSet.tmax

`DataSet.tmax(self, over='time')`

Calculate the temporal maximum of all variables

Parameters `over` (*str or list*) – Time periods to average over. Options are 'year', 'month', 'day'.

Examples

If you want to calculate maximum over all time steps. Do the following:

```
>>> ds.tmax()
```

If you want to calculate the maximum for each year in a dataset, do this:

```
>>> ds.tmax("year")
```

If you want to calculate the maximum for each month in a dataset, do this:

```
>>> ds.tmax("month")
```

If you want to calculate the maximum for each month in each year in a dataset, do this:

```
>>> ds.tmax(["year", "month"])
```

This method will also let you easily calculate climatologies. So, if you wanted to calculate a monthly climatological max, you would do this:

```
>>> ds.tmax("month")
```

A daily climatological maximum would be the following:

```
>>> ds.tmax("day")
```

17.21.6 nctoolkit.DataSet.tsum

`DataSet.tsum(self, over='time')`

Calculate the temporal sum of all variables

17.21.7 nctoolkit.DataSet.trange

`DataSet.trange(self, over='time')`

Calculate the temporal range of all variables

Parameters `over` (*str* or *list*) – Time periods to average over. Options are ‘year’, ‘month’, ‘day’.

Examples

If you want to calculate range over all time steps. Do the following:

```
>>> ds.trange()
```

If you want to calculate the range for each year in a dataset, do this:

```
>>> ds.trange("year")
```

If you want to calculate the range for each month in a dataset, do this:

```
>>> ds.trange("month")
```

If you want to calculate the range for each month in each year in a dataset, do this:

```
>>> ds.trange(["year", "month"])
```

This method will also let you easily calculate climatologies. So, if you wanted to calculate a monthly climatological range, you would do this:

```
>>> ds.trange("month")
```

A daily climatological range would be the following:

```
>>> ds.trange( "day")
```

17.21.8 nctoolkit.DataSet.tvariance

`DataSet.tvariance(self, over='time')`

Calculate the temporal variance of all variables

Parameters `over` (*str* or *list*) – Time periods to average over. Options are ‘year’, ‘month’, ‘day’.

Examples

If you want to calculate variance over all time steps. Do the following:

```
>>> ds.tvariance()
```

If you want to calculate the variance for each year in a dataset, do this:

```
>>> ds.tvariance("year")
```

If you want to calculate the variance for each month in a dataset, do this:

```
>>> ds.tvariance("month")
```

If you want to calculate the variance for each month in each year in a dataset, do this:

```
>>> ds.tvariance(["year", "month"])
```

This method will also let you easily calculate climatologies. So, if you wanted to calculate a monthly climatological var, you would do this:

```
>>> ds.tvariance( "month")
```

A daily climatological variance would be the following:

```
>>> ds.tvariance( "day")
```

17.21.9 nctoolkit.DataSet.tstdev

`DataSet.tstdev(self, over='time')`

Calculate the temporal standard deviation of all variables

Parameters `over` (*str* or *list*) – Time periods to average over. Options are ‘year’, ‘month’, ‘day’.

Examples

If you want to calculate standard deviation over all time steps. Do the following:

```
>>> ds.tstdev()
```

If you want to calculate the standard deviation for each year in a dataset, do this:

```
>>> ds.tstdev("year")
```

If you want to calculate the standard deviation for each month in a dataset, do this:

```
>>> ds.tstdev("month")
```

If you want to calculate the standard deviation for each month in each year in a dataset, do this:

```
>>> ds.tstdev(["year", "month"])
```

This method will also let you easily calculate climatologies. So, if you wanted to calculate a monthly climatological var, you would do this:

```
>>> ds.tstdev("month")
```

A daily climatological standard deviation would be the following:

```
>>> ds.tstdev("day")
```

17.21.10 nctoolkit.DataSet.tcumsum

`DataSet.tcumsum(self)`

Calculate the temporal cumulative sum of all variables

Examples

If you want to calculate the cumulative sum for all variables over all timesteps, do this:

```
>>> ds.tcumsum()
```

17.21.11 nctoolkit.DataSet.tvar

`DataSet.tvar(self, over='time')`

Calculate the temporal variance of all variables

Parameters `over` (*str* or *list*) – Time periods to average over. Options are ‘year’, ‘month’, ‘day’.

Examples

If you want to calculate variance over all time steps. Do the following:

```
>>> ds.tvar()
```

If you want to calculate the variance for each year in a dataset, do this:

```
>>> ds.tvar("year")
```

If you want to calculate the variance for each month in a dataset, do this:

```
>>> ds.tvar("month")
```

If you want to calculate the variance for each month in each year in a dataset, do this:

```
>>> ds.tvar(["year", "month"])
```

This method will also let you easily calculate climatologies. So, if you wanted to calculate a monthly climatological var, you would do this:

```
>>> ds.tvar("month")
```

A daily climatological variance would be the following:

```
>>> ds.tvar("day")
```

17.21.12 nctoolkit.DataSet.cor_space

`DataSet.cor_space(self, var1=None, var2=None)`

Calculate the correlation correct between two variables in space This is calculated for each time step. The correlation coefficient is calculated using values in all grid cells, ignoring missing values.

Parameters

- **var1** (*str*) – The first variable
- **var2** (*str*) – The second variable

Examples

If you wanted to calculate the spatial correlation coefficient between variables x and y in a dataset, you would do this:

```
>>> ds.cor_space("x", "y")
```

The correlation coefficient will be calculated for each time step.

17.21.13 nctoolkit.DataSet.cor_time

`DataSet.cor_time(self, var1=None, var2=None)`

Calculate the correlation correct in time between two variables The correlation is calculated for each grid cell, ignoring missing values.

Parameters

- **var1** (*str*) – The first variable
- **var2** (*str*) – The second variable

Examples

If you wanted to calculate the temporal correlation coefficient between variables x and y in a dataset, you would do this:

```
>>> ds.cor_space("x", "y")
```

The correlation coefficient will be calculated for each grid cell. This method will indicate how temporally correlated variables are in different spatial regions.

17.21.14 nctoolkit.DataSet.spatial_mean

`DataSet.spatial_mean(self)`

Calculate the area weighted spatial mean for all variables This is performed for each time step.

Examples

If you want to calculate the spatial mean for a dataset, just do the following:

```
>>> ds.spatial_mean()
```

Note that this calculation will calculate the average using weights based on each cell's area. If cell areas cannot be calculated, it will take a straight average, and a warning will say this.

17.21.15 nctoolkit.DataSet.spatial_min

`DataSet.spatial_min(self)`

Calculate the spatial minimum for all variables This is performed for each time step.

Examples

If you want to calculate the spatial minimum for a dataset, just do the following:

```
>>> ds.spatial_min()
```

17.21.16 nctoolkit.DataSet.spatial_max

`DataSet.spatial_max(self)`

Calculate the spatial maximum for all variables This is performed for each time step.

Examples

If you want to calculate the spatial maximum for a dataset, just do the following:

```
>>> ds.spatial_max()
```

17.21.17 nctoolkit.DataSet.spatial_percentile

`DataSet.spatial_percentile(self, p=None)`

Calculate the spatial sum for all variables This is performed for each time step. :param p: Percentile to calculate. 0<=p<=100. :type p: int or float

Examples

If you want to calculate the median of each variable across space for a dataset, just do the following:

```
>>> ds.spatial_percentile(50)
```

17.21.18 nctoolkit.DataSet.spatial_range

`DataSet.spatial_range(self)`

Calculate the spatial range for all variables This is performed for each time step.

Examples

If you want to calculate the range of each variable across space for a dataset, just do the following:

```
>>> ds.spatial_max()
```

17.21.19 nctoolkit.DataSet.spatial_sum

`DataSet.spatial_sum(self, by_area=False)`

Calculate the spatial sum for all variables This is performed for each time step.

Parameters `by_area` (*boolean*) – Set to True if you want to multiply the values by the grid cell area before summing over space. Default is False.

Examples

If you want to calculate the spatial sum each variable across space for a dataset, just do the following:

```
>>> ds.spatial_sum()
```

By default, this method simply sums up each grid cell value. In some cases this is not suitable. For example, the values in each cell may concentrations or values per square metre etc. In this case multiplying each cell value by the cell area is more suitable. Do the following:

```
>>> ds.spatial_sum(by_area = True)
```

Each cell's value will be multiplied by the area of the cell (in square metres) prior to calculating the spatial sum.

17.21.20 nctoolkit.DataSet.centre

DataSet.centre(*self*, *by*='latitude', *by_area*=False)

Calculate the latitudinal or longitudinal centre for each year/month combination in files. This applies to each file in an ensemble. *by* : str

Set to 'latitude' if you want the latitudinal centre calculated. 'longitude' for longitudinal.

by_area [bool] If the variable is a value/m2 type variable, set to True, otherwise set to False.

17.21.21 nctoolkit.DataSet.zonal_mean

DataSet.zonal_mean(*self*)

Calculate the zonal mean for each year/month combination in files. This applies to each file in an ensemble.

Examples

If you want to calculate the zonal mean for a dataset, do the following:

```
>>> ds.zonal_mean()
```

17.21.22 nctoolkit.DataSet.zonal_min

DataSet.zonal_min(*self*)

Calculate the zonal minimum for each year/month combination in files. This applies to each file in an ensemble.

Examples

If you want to calculate the zonal minimum for a dataset, do the following:

```
>>> ds.zonal_min()
```

17.21.23 nctoolkit.DataSet.zonal_max

`DataSet.zonal_max(self)`

Calculate the zonal maximum for each year/month combination in files. This applies to each file in an ensemble.

Examples

If you want to calculate the zonal maximum for a dataset, do the following:

```
>>> ds.zonal_max()
```

17.21.24 nctoolkit.DataSet.zonal_range

`DataSet.zonal_range(self)`

Calculate the zonal range for each year/month combination in files. This applies to each file in an ensemble.

Examples

If you want to calculate the zonal range for a dataset, do the following:

```
>>> ds.zonal_range()
```

17.21.25 nctoolkit.DataSet.meridional_mean

`DataSet.meridional_mean(self)`

Calculate the meridional mean for each year/month combination in files. This applies to each file in an ensemble.

Examples

If you want to calculate the meridional mean for a dataset, do the following:

```
>>> ds.meridional_mean()
```

17.21.26 nctoolkit.DataSet.meridional_min

`DataSet.meridional_min(self)`

Calculate the meridional minimum for each year/month combination in files. This applies to each file in an ensemble.

Examples

If you want to calculate the meridional minimum for a dataset, do the following:

```
>>> ds.meridional_min()
```

17.21.27 nctoolkit.DataSet.meridional_max

DataSet.meridional_max(*self*)

Calculate the meridional maximum for each year/month combination in files. This applies to each file in an ensemble.

Examples

If you want to calculate the meridional maximum for a dataset, do the following:

```
>>> ds.meridional_max()
```

17.21.28 nctoolkit.DataSet.meridional_range

DataSet.meridional_range(*self*)

Calculate the meridional range for each year/month combination in files. This applies to each file in an ensemble.

Examples

If you want to calculate the meridional range for a dataset, do the following:

```
>>> ds.meridional_max()
```

17.22 Merging methods

DataSet.merge(*self*[, *join*, *match*])

Merge a multi-file ensemble into a single file 2 methods are available.

17.22.1 nctoolkit.DataSet.merge

DataSet.merge(*self*, *join*='variables', *match*=['year', 'month', 'day'])

Merge a multi-file ensemble into a single file 2 methods are available. 1) merging files with different variables, but the same time steps. 2) merging files with the same variables, with different times.

Parameters

- **join** (*str*) – This defines the type of merging to carry out. “variables”: this will merge by variable, so that an ensemble with different variables, but the same number of time steps is merged to a single file. “time”: this will merge files with the same variables, but different times to a single file, into a single file with ordered times. *join* defaults to “variables”, and uses partial matches, so “var” will give variable based merging.

- **match** (*list*, *str*) – Optional argument when join = ‘variables’. A list or str stating what must match in the netCDF files. Defaults to year/month/day. This list must be some combination of year/month/day. An error will be thrown if the elements of time in match do not match across all netCDF files. The only exception is if there is a single date file in the ensemble.

17.23 Splitting methods

<code>DataSet.split(self[, by])</code>	Split the dataset Each file in the ensemble will be separated into new files based on the splitting argument.
--	---

17.23.1 nctoolkit.DataSet.split

`DataSet.split(self, by=None)`

Split the dataset Each file in the ensemble will be separated into new files based on the splitting argument.

Parameters *by* (*str*) – Available by arguments are ‘year’, ‘month’, ‘yearmonth’, ‘season’, ‘day’. year will split files by year, month will split files by month, yearmonth will split files by year and month; season will split files by year, day will split files by day.

Examples

If you want to split each file into a dataset into a separate files for each year, do the following:

```
>>> ds.split("year")
```

If you wanted to split by month, do the following:

```
>>> ds.split("month")
```

17.24 Output and formatting methods

<code>DataSet.to_nc(self, out[, zip, overwrite])</code>	Save a dataset to a named file This will only work with single file datasets.
<code>DataSet.to_xarray(self[, decode_times, ...])</code>	Open a dataset as an xarray object
<code>DataSet.to_dataframe(self[, decode_times, ...])</code>	Open a dataset as a pandas data frame
<code>DataSet.zip(self)</code>	Zip the dataset This will compress the files within the dataset.
<code>DataSet.format(self[, ext])</code>	Zip the dataset This will compress the files within the dataset. This works lazily. :param ext: New format. Must be one of “nc”, “nc1”, “nc2”, “nc4” and “nc5”. netCDF = nc1 netCDF version 2 (64-bit offset) = nc2/netCDF4 (HDF5) = nc4 netCDF4-classic = nc4c netCDF version 5 (64-bit data) = nc5 :type ext: str.

17.24.1 nctoolkit.DataSet.to_nc

`DataSet.to_nc(self, out, zip=True, overwrite=False)`

Save a dataset to a named file This will only work with single file datasets.

Parameters

- **out** (*str*) – Output file name.
- **zip** (*boolean*) – True/False depending on whether you want to zip the file. Default is True.
- **overwrite** (*boolean*) – If out file exists, do you want to overwrite it? Default is False.

Examples

If you want to export a dataset to a netCDF file, do the following:

```
>>> ds.to_nc("out.nc")
```

By default this file will be zipped. If you do not want it zipped, do this:

```
>>> ds.to_nc("out.nc", zip = False)
```

By default this cannot overwrite files. If the output file exists, do the following:

```
>>> ds.to_nc("out.nc", overwrite = True)
```

17.24.2 nctoolkit.DataSet.to_xarray

`DataSet.to_xarray(self, decode_times=True, cdo_times=False)`

Open a dataset as an xarray object

Parameters

- **decode_times** (*boolean*) – Set to False if you do not want xarray to decode the times. Default is True. If xarray cannot decode times, CDO will be used.
- **cdo_times** (*boolean*) – Set to True if you do not want CDO to decode the times

Returns to_xarray

Return type xarray.Dataset

Examples

If you want to convert a dataset to an xarray dataset, do the following:

```
>>> ds.to_xarray()
```

This will return an xarray dataset.

If you do not want time to be decoded, do the following:

```
>>> ds.to_xarray(decode_times = False)
```

17.24.3 nctoolkit.DataSet.to_dataframe

`DataSet.to_dataframe(self, decode_times=True, cdo_times=False)`

Open a dataset as a pandas data frame

Parameters

- **decode_times** (*boolean*) – Set to False if you do not want xarray to decode the times prior to conversion to data frame. Default is True.
- **cdo_times** (*boolean*) – Set to True if you do not want CDO to decode the times

Returns to_dataframe

Return type pandas.DataFrame

17.24.4 nctoolkit.DataSet.zip

`DataSet.zip(self)`

Zip the dataset This will compress the files within the dataset. This works lazily.

Examples

If you want to zip the files in a dataset, do the following:

```
>>> ds.zip()
```

This will occur lazily, so will only occur after everything has been evaluated.

17.24.5 nctoolkit.DataSet.format

`DataSet.format(self, ext=None)`

Zip the dataset This will compress the files within the dataset. This works lazily. :param ext: New format. Must be one of “nc”, “nc1”, “nc2”, “nc4” and “nc5” .

netCDF = nc1 netCDF version 2 (64-bit offset) = nc2/nc netCDF4 (HDF5) = nc4 netCDF4-classi = nc4c netCDF version 5 (64-bit data) = nc5

17.25 Miscellaneous methods

<code>DataSet.distribute(self[, m, n])</code>	Split the dataset into multiple evenly sized horizontal and vertical new files
<code>DataSet.collect(self)</code>	Collect a dataset that has been split using distribute
<code>DataSet.cell_area(self[, join])</code>	Calculate the area of grid cells.

continues on next page

Table 27 – continued from previous page

<code>DataSet.first_above(self[, x])</code>	Identify the time step when a value is first above a threshold This will do the comparison with either a number, a Dataset or a netCDF file. :param x: An int, float, single file dataset or netCDF file to use for the threshold(s). If comparing with a dataset or single file there must only be a single variable in it. The grids must be the same. :type x: int, float, DataSet or netCDF file.
<code>DataSet.first_below(self[, x])</code>	Identify the time step when a value is first below a threshold This will do the comparison with either a number, a Dataset or a netCDF file. :param x: An int, float, single file dataset or netCDF file to use for the threshold(s). If comparing with a dataset or single file there must only be a single variable in it. The grids must be the same. :type x: int, float, DataSet or netCDF file.
<code>DataSet.last_above(self[, x])</code>	Identify the final time step when a value is above a threshold This will do the comparison with either a number, a Dataset or a netCDF file. :param x: An int, float, single file dataset or netCDF file to use for the threshold(s). If comparing with a dataset or single file there must only be a single variable in it. The grids must be the same. :type x: int, float, DataSet or netCDF file.
<code>DataSet.last_below(self[, x])</code>	Identify the final time step when a value is below a threshold This will do the comparison with either a number, a Dataset or a netCDF file. :param x: An int, float, single file dataset or netCDF file to use for the threshold(s). If comparing with a dataset or single file there must only be a single variable in it. The grids must be the same. :type x: int, float, DataSet or netCDF file.
<code>DataSet.cdo_command(self[, command])</code>	Apply a cdo command
<code>DataSet.nco_command(self[, command, ensemble])</code>	Apply an nco command
<code>DataSet.compare(self[, expression])</code>	Compare all variables to a constant
<code>DataSet.gt(self, x)</code>	Method to calculate if variable in dataset is greater than that in another file or dataset This currently only works with single file datasets
<code>DataSet.lt(self, x)</code>	Method to calculate if variable in dataset is less than that in another file or dataset This currently only works with single file datasets
<code>DataSet.reduce_dims(self)</code>	Reduce dimensions of data This will remove any dimensions with only one value.
<code>DataSet.reduce_grid(self[, mask])</code>	Reduce the dataset to non-zero locations in a mask :param mask: single variable dataset or path to .nc file. The mask must have an identical grid to the dataset. :type mask: str or dataset.

17.25.1 nctoolkit.DataSet.distribute

`DataSet.distribute(self, m=1, n=1)`

Split the dataset into multiple evenly sized horizontal and vertical new files

Parameters

- **m** (*int*) – Number of rows
- **n** (*int*) – Number of columns

17.25.2 nctoolkit.DataSet.collect

`DataSet.collect(self)`

Collect a dataset that has been split using distribute

17.25.3 nctoolkit.DataSet.cell_area

`DataSet.cell_area(self, join=True)`

Calculate the area of grid cells. Area of grid cells is given in square meters.

Parameters **join** (*boolean*) – Set to False if you only want the cell areas to be in the output. join=True adds the areas as a variable to the dataset. Defaults to True.

Examples

If you wanted to add the cell_areas as a new variable in a dataset, you would do the following:

```
>>> ds.cell_area()
```

If you wanted to replace a dataset with the cell areas of that dataset, you would do the following:

```
>>> ds.cell_area(join = False)
```

17.25.4 nctoolkit.DataSet.first_above

`DataSet.first_above(self, x=None)`

Identify the time step when a value is first above a threshold This will do the comparison with either a number, a Dataset or a netCDF file. :param x: An int, float, single file dataset or netCDF file to use for the threshold(s).

If comparing with a dataset or single file there must only be a single variable in it. The grids must be the same.

Examples

If you wanted to calculate the first time step where the value in a grid cell goes above 10, you would do the following

```
>>> ds.first_above(10)
```

If you wanted to calculate the first time step where the value in a grid cell goes above that in another dataset, the following will work. Note that both datasets must have the same grid, and can only have single variables. The second dataset can, of course, only have one timestep.

```
>>> ds.first_above(ds1)
```

17.25.5 nctoolkit.DataSet.first_below

DataSet.first_below(*self*, *x=None*)

Identify the time step when a value is first below a threshold This will do the comparison with either a number, a Dataset or a netCDF file. :param x: An int, float, single file dataset or netCDF file to use for the threshold(s).

If comparing with a dataset or single file there must only be a single variable in it. The grids must be the same.

Examples

If you wanted to calculate the first time step where the value in a grid cell goes below 10, you would do the following

```
>>> ds.first_below(10)
```

If you wanted to calculate the first time step where the value in a grid cell goes above that in another dataset, the following will work. Note that both datasets must have the same grid, and can only have single variables. The second dataset can, of course, only have one timestep.

```
>>> ds.first_below(ds1)
```

17.25.6 nctoolkit.DataSet.last_above

DataSet.last_above(*self*, *x=None*)

Identify the final time step when a value is above a threshold This will do the comparison with either a number, a Dataset or a netCDF file. :param x: An int, float, single file dataset or netCDF file to use for the threshold(s).

If comparing with a dataset or single file there must only be a single variable in it. The grids must be the same.

Examples

If you wanted to calculate the last time step where the value in a grid cell is above 10, you would do the following

```
>>> ds.first_above(10)
```

If you wanted to calculate the last time step where the value in a grid cell goes above that in another dataset, the following will work. Note that both datasets must have the same grid, and can only have single variables. The second dataset can, of course, only have one timestep.

```
>>> ds.first_above(ds1)
```

17.25.7 nctoolkit.DataSet.cdo_command

`DataSet.cdo_command(self, command=None)`

Apply a cdo command

Parameters **command** (*string*) – cdo command to call. This command must be such that “cdo {command} infile outfile” will run.

17.25.8 nctoolkit.DataSet.nco_command

`DataSet.nco_command(self, command=None, ensemble=False)`

Apply an nco command

Parameters

- **command** (*string*) – nco command to call. This must be of a form such that “nco {command} infile outfile” will run.
- **ensemble** (*boolean*) – Set to True if you want the command to take all of the files as input. This is useful for ensemble methods.

17.25.9 nctoolkit.DataSet.compare

`DataSet.compare(self, expression=None)`

Compare all variables to a constant

Parameters **expression** (*str*) – This a regular comparison such as “<0”, “>0”, “==0”

Examples

If you wanted to identify grid cells with positive values you would do the following:

```
>>> ds.compare(">0")
```

This will be calculated for each time step.

If you wanted to identify grid cells with negative values, you would do this

```
>>> ds.compare("<0")
```

17.25.10 nctoolkit.DataSet.gt

`DataSet.gt(self, x)`

Method to calculate if variable in dataset is greater than that in another file or dataset This currently only works with single file datasets

Parameters *x* (*str or single file dataset*) – File path or nctoolkit dataset

17.25.11 nctoolkit.DataSet.lt

`DataSet.lt(self, x)`

Method to calculate if variable in dataset is less than that in another file or dataset This currently only works with single file datasets

Parameters *x* (*str or single file dataset*) – File path or nctoolkit dataset

17.25.12 nctoolkit.DataSet.reduce_dims

`DataSet.reduce_dims(self)`

Reduce dimensions of data This will remove any dimensions with only one value. For example, if only selecting one vertical level, the vertical dimension will be removed.

Examples

If you want to remove any dimensions that have only one value, do the following:

```
>>> ds.reduce_dims("out.nc")
```

Note that this will work lazily. This method is most useful when you want to simplify datasets before exporting them to something like a pandas dataframe.

17.25.13 nctoolkit.DataSet.reduce_grid

`DataSet.reduce_grid(self, mask=None)`

Reduce the dataset to non-zero locations in a mask :param mask: single variable dataset or path to .nc file.

The mask must have an identical grid to the dataset.

17.26 Ecological methods

`DataSet.phenology(self[, var, metric, p])`

Calculate phenologies from a dataset Each file in an ensemble must only cover a single year, and ideally have all days.

17.26.1 nctoolkit.DataSet.phenology

`DataSet.phenology(self, var=None, metric=None, p=None)`

Calculate phenologies from a dataset Each file in an ensemble must only cover a single year, and ideally have all days. The method assumes datasets have daily resolution.

Parameters

- **var** (*str*) – Variable to analyze.
- **metric** (*str*) – Must be peak, middle, start or end. Peak is defined as the day of the maximum value. Middle is the day when the cumulative total of the variable first exceeds the cumulative total for the entire year. Start or end is defined as the first day when the cumulative total exceeds a percentile *p* of the maximum cumulative total.
- **p** (*str*) – Percentile to use for start or end.

PACKAGE INFO

This package was created by Robert Wilson at Plymouth Marine Laboratory (PML).

18.1 Acknowledgements

The current codebase of nctoolkit was developed using funding from the NERC Climate Linked Atlantic Sector Science programme ([NE/R015953/1](#)) and a combination of UK Research and Innovation (UKRI) and European Research Council (ERC) funded research projects.

18.2 Bugs and issues

If you identify bugs or issues with the package please raise an issue at PML's Marine Systems Modelling group's GitHub page [here](#) or contact nctoolkit's creator at rwil@pml.ac.uk.

18.3 Contributions welcome

The package is new, with new features being added each month. There remain a large number of features that could be added, especially for dealing with atmospheric data. If packages users are interested in contributing or suggesting new features they are welcome to raise and issue at the package's GitHub page or contact me.

PYTHON MODULE INDEX

n

`nctoolkit.append`, [55](#)

`nctoolkit.remove`, [55](#)

A

`abs()` (*nctoolkit.DataSet* method), 68
`add()` (*nctoolkit.DataSet* method), 68
`annual_anomaly()` (*nctoolkit.DataSet* method), 80
`assign()` (*nctoolkit.DataSet* method), 57

B

`bottom()` (*nctoolkit.DataSet* method), 60
`bottom_mask()` (*nctoolkit.DataSet* method), 63

C

`cdo_command()` (*nctoolkit.DataSet* method), 100
`cell_area()` (*nctoolkit.DataSet* method), 98
`centre()` (*nctoolkit.DataSet* method), 91
`collect()` (*nctoolkit.DataSet* method), 98
`compare()` (*nctoolkit.DataSet* method), 100
`copy()` (*nctoolkit.DataSet* method), 54
`cor_space()` (in module *nctoolkit*), 54
`cor_space()` (*nctoolkit.DataSet* method), 88
`cor_time()` (in module *nctoolkit*), 54
`cor_time()` (*nctoolkit.DataSet* method), 89
`create_ensemble()` (in module *nctoolkit*), 66
`crop()` (*nctoolkit.DataSet* method), 74
`current` (*nctoolkit.DataSet* property), 56

D

`distribute()` (*nctoolkit.DataSet* method), 98
`divide()` (*nctoolkit.DataSet* method), 71
`drop()` (*nctoolkit.DataSet* method), 76

E

`ensemble_max()` (*nctoolkit.DataSet* method), 73
`ensemble_mean()` (*nctoolkit.DataSet* method), 72
`ensemble_min()` (*nctoolkit.DataSet* method), 73
`ensemble_percentile()` (*nctoolkit.DataSet* method), 73
`ensemble_range()` (*nctoolkit.DataSet* method), 73
`ensemble_stddev()` (*nctoolkit.DataSet* method), 73
`ensemble_sum()` (*nctoolkit.DataSet* method), 74
`ensemble_var()` (*nctoolkit.DataSet* method), 74
`exp()` (*nctoolkit.DataSet* method), 68

F

`fill_na()` (*nctoolkit.DataSet* method), 79
`first_above()` (*nctoolkit.DataSet* method), 98
`first_below()` (*nctoolkit.DataSet* method), 99
`format()` (*nctoolkit.DataSet* method), 96

G

`gt()` (*nctoolkit.DataSet* method), 101

H

`history` (*nctoolkit.DataSet* property), 56

I

`invert_levels()` (*nctoolkit.DataSet* method), 63

L

`last_above()` (*nctoolkit.DataSet* method), 99
`levels` (*nctoolkit.DataSet* property), 56
`log()` (*nctoolkit.DataSet* method), 69
`log10()` (*nctoolkit.DataSet* method), 69
`lt()` (*nctoolkit.DataSet* method), 101

M

`mask_box()` (*nctoolkit.DataSet* method), 79
`merge()` (in module *nctoolkit*), 54
`merge()` (*nctoolkit.DataSet* method), 93
`meridional_max()` (*nctoolkit.DataSet* method), 93
`meridional_mean()` (*nctoolkit.DataSet* method), 92
`meridional_min()` (*nctoolkit.DataSet* method), 92
`meridional_range()` (*nctoolkit.DataSet* method), 93
`module`
 nctoolkit.append, 55
 nctoolkit.remove, 55
`monthly_anomaly()` (*nctoolkit.DataSet* method), 80
`months` (*nctoolkit.DataSet* property), 56
`multiply()` (*nctoolkit.DataSet* method), 69

N

`nco_command()` (*nctoolkit.DataSet* method), 100
`nctoolkit.append`
 module, 55

nctoolkit.remove
module, 55

O

open_data() (in module nctoolkit), 52
open_thredds() (in module nctoolkit), 53
open_url() (in module nctoolkit), 53
options() (in module nctoolkit), 51

P

phenology() (nctoolkit.DataSet method), 102
plot() (nctoolkit.DataSet method), 57
power() (nctoolkit.DataSet method), 70

R

reduce_dims() (nctoolkit.DataSet method), 101
reduce_grid() (nctoolkit.DataSet method), 101
regrid() (nctoolkit.DataSet method), 77
rename() (nctoolkit.DataSet method), 58
resample_grid() (nctoolkit.DataSet method), 78
rolling_max() (nctoolkit.DataSet method), 64
rolling_mean() (nctoolkit.DataSet method), 64
rolling_min() (nctoolkit.DataSet method), 64
rolling_range() (nctoolkit.DataSet method), 65
rolling_sum() (nctoolkit.DataSet method), 65
run() (nctoolkit.DataSet method), 65

S

select() (nctoolkit.DataSet method), 75
set_date() (nctoolkit.DataSet method), 76
set_longnames() (nctoolkit.DataSet method), 59
set_missing() (nctoolkit.DataSet method), 58
set_units() (nctoolkit.DataSet method), 59
shift() (nctoolkit.DataSet method), 76
size (nctoolkit.DataSet property), 56
spatial_max() (nctoolkit.DataSet method), 90
spatial_mean() (nctoolkit.DataSet method), 89
spatial_min() (nctoolkit.DataSet method), 89
spatial_percentile() (nctoolkit.DataSet method), 90
spatial_range() (nctoolkit.DataSet method), 90
spatial_sum() (nctoolkit.DataSet method), 90
split() (nctoolkit.DataSet method), 94
sqrt() (nctoolkit.DataSet method), 70
square() (nctoolkit.DataSet method), 70
start (nctoolkit.DataSet property), 57
subtract() (nctoolkit.DataSet method), 71
sum_all() (nctoolkit.DataSet method), 58
surface() (nctoolkit.DataSet method), 60

T

tcumsum() (nctoolkit.DataSet method), 87
time_interp() (nctoolkit.DataSet method), 78
times (nctoolkit.DataSet property), 56

timestep_interp() (nctoolkit.DataSet method), 79
tmax() (nctoolkit.DataSet method), 84
tmean() (nctoolkit.DataSet method), 82
tmedian() (nctoolkit.DataSet method), 83
tmin() (nctoolkit.DataSet method), 83
to_dataframe() (nctoolkit.DataSet method), 96
to_latlon() (nctoolkit.DataSet method), 78
to_nc() (nctoolkit.DataSet method), 95
to_xarray() (nctoolkit.DataSet method), 95
tpercentile() (nctoolkit.DataSet method), 84
trange() (nctoolkit.DataSet method), 85
tstdev() (nctoolkit.DataSet method), 86
tsum() (nctoolkit.DataSet method), 85
tvar() (nctoolkit.DataSet method), 87
tvariance() (nctoolkit.DataSet method), 86

V

variables (nctoolkit.DataSet property), 56
vertical_cumsum() (nctoolkit.DataSet method), 63
vertical_integration() (nctoolkit.DataSet method), 62
vertical_interp() (nctoolkit.DataSet method), 60
vertical_max() (nctoolkit.DataSet method), 61
vertical_mean() (nctoolkit.DataSet method), 61
vertical_min() (nctoolkit.DataSet method), 61
vertical_range() (nctoolkit.DataSet method), 62
vertical_sum() (nctoolkit.DataSet method), 62

Y

years (nctoolkit.DataSet property), 56

Z

zip() (nctoolkit.DataSet method), 96
zonal_max() (nctoolkit.DataSet method), 92
zonal_mean() (nctoolkit.DataSet method), 91
zonal_min() (nctoolkit.DataSet method), 91
zonal_range() (nctoolkit.DataSet method), 92