
nctoolkit

Robert Wilson

Jan 03, 2021

GETTING STARTED

1	Fixing plotting problem due to xarray bug	3
2	Documentation	5
	Python Module Index	81
	Index	83

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

FIXING PLOTTING PROBLEM DUE TO XARRAY BUG

There is currently a bug in xarray caused by the update of pandas to version 1.1. As a result some plots will fail in nctoolkit. To fix this ensure pandas version 1.0.5 is installed. Do this after installing nctoolkit. This can be done as follows:

```
$ conda install -c conda-forge pandas=1.0.5
```

or:

```
$ pip install pandas==1.0.5
```


DOCUMENTATION

Quick overview

- *Installation*
- *Introduction tutorial*
- *Ensemble methods*
- *Speeding up code*

2.1 Installation

2.1.1 Python dependencies

- Python (3.6 or later)
- `numpy` (1.14 or later)
- `pandas` (0.24 or later)
- `xarray` (0.14 or later)
- `hvplot` (0.5 or later)
- `NetCDF4` (1.53 or later)
- `panel` (0.9.1 or later)

2.1.2 How to install nctoolkit

The easiest way to install the package is using conda, which will install nctoolkit and all system dependencies:

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

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

```
$ pip install nctoolkit
```

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

To install the development version from GitHub:

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

2.1.3 Fixing plotting problem due to xarray bug

There is currently a bug in xarray caused by the update of pandas to version 1.1. As a result some plots will fail in nctoolkit. To fix this ensure pandas version 1.0.5 is installed. Do this after installing nctoolkit. This can be done as follows:

```
$ conda install -c conda-forge pandas=1.0.5
```

or:

```
$ pip install pandas==1.0.5
```

2.1.4 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](#).

2.2 Introduction tutorial

nctoolkit is designed for the efficient analysis and manipulation of netCDF files. This tutorial provides an overview of how to work with individual files.

2.2.1 Opening netcdf data

This tutorial will illustrate the basic usage using a dataset of average global sea surface temperature from NOAA, which is available [here](#).

nctoolkit should be imported using the nc shorthand:

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

```
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
```

Reading in a dataset is straightforward:

```
[2]: ff = "sst.mon.ltm.1981-2010.nc"
sst = nc.open_data(ff)
```

We might want to know some basic information about the file. This can be done easily. Listing the available variables can be found quickly:

The current state of the dataset can be found as follows:

```
[3]: sst.variables
[3]: ['sst', 'valid_yr_count']
```

The months available can be found using:

```
[4]: sst.months
[4]: [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12]
```

We have 12 months available. In this case it is the monthly average temperature from 1981-2010.

2.2.2 Modifying datasets

Each time nctoolkit executes a command that modifies a dataset, it will generate a new NetCDF file, which becomes the current file in the dataset. Before any modification this is as follows:

```
[5]: sst.current
[5]: 'sst.mon.ltm.1981-2010.nc'
```

We have seen that there are two variables in the dataset. But we only really care about `sst`. So let's select that variable:

```
[6]: sst.select_variables("sst")
```

We can now see that there is only one variable in the `sst` dataset

```
[7]: sst.variables
[7]: ['sst']
```

We can also that a temporary file has been created with only this variable in it

```
[8]: sst.current
[8]: '/tmp/nctoolkitesugmpemnctoolkittmpxmrohbp.nc'
```

We have data for 12 months. But what we might really want is an average of those values. This can be quickly calculated:

```
[9]: sst.mean()
```

Once again a new temporary file has been generated.

```
[10]: sst.current
[10]: '/tmp/nctoolkitesugmpemnctoolkittmpgz_hzyoq.nc'
```

Do not worry about the temporary folder getting clogged up. nctoolkit cleans it up automatically.

Quick visualization of netCDF data is always a good thing. So nctoolkit provides an easy autoplot feature.

```
[11]: sst.plot()
```

```
[11]: :DynamicMap    [time]
      :Image      [lon,lat]    (sst)
```

What we have seen so far is not computationally efficient. In the code below nctoolkit has generated temporary files twice:

```
[12]: sst = nc.open_data(ff)
      sst.select_variables("sst")
      sst.mean()
```

We can see what went on behind the scenes by accessing history:

```
[13]: sst.history
[13]: ['cdo -L -selname,sst sst.mon.ltm.1981-2010.nc /tmp/nctoolkitesugmpemncnctoolkittmpxpb_
      ↪323a.nc',
      'cdo -L -timmean /tmp/nctoolkitesugmpemncnctoolkittmpxpb_323a.nc /tmp/
      ↪nctoolkitesugmpemncnctoolkittmp5agj679e.nc']
```

nctoolkit uses CDO. You do not understand how CDO works to use nctoolkit. But one nice feature of CDO is method chaining, which works like Python's. To get it working you just need to set evaluation to lazy in nctoolkit. This means nothing is evaluated until you force it to or it has to be.

```
[14]: nc.options(lazy = True)
```

Now, let's run the code again:

```
[15]: sst = nc.open_data(ff)
      sst.select_variables("sst")
      sst.mean()
      sst.plot()
[15]: :DynamicMap    [time]
      :Image      [lon,lat]    (sst)
```

When we look at history, we now see that only one temporary file was generated:

```
[16]: sst.history
[16]: ['cdo -L -timmean -selname,sst sst.mon.ltm.1981-2010.nc /tmp/
      ↪nctoolkitesugmpemncnctoolkittmpoqilxou.nc']
```

In the example, above the commands were only executed when plot was called. If we want to force commands to run we use run:

```
[17]: sst = nc.open_data(ff)
      sst.select_variables("sst")
      sst.mean()
      sst.run()
```

User Guide

- [Datasets](#)

2.3 Datasets

nctoolkit works with what it calls datasets. Each dataset is made up of a single or multiple NetCDF files. Each time you apply a method to a dataset the NetCDF file or files within the dataset will be modified.

2.3.1 Opening datasets

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

If the data you want to analyze is already available on your computer use `open_data`. This will accept either a path to a single file or a list of files to create a dataset.

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, then use `open_thredds`. The file paths should end with `.nc`.

2.3.2 Dataset attributes

We can find out key information about a dataset using its attributes. Here we will use a sea surface temperature file that is available via thredds.

```
[2]: import nctoolkit as nc
sst = nc.open_thredds("https://psl.noaa.gov/thredds/dodsC/Datasets/COBE/sst.mon.ltm.
↪1981-2010.nc")
```

```
1 file was created by nctoolkit in prior or current sessions. Consider running deep_
↪clean!
```

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

If we want to know a dataset's variables:

```
[3]: sst.variables
[3]: ['sst', 'valid_yr_count']
```

If we want to know a dataset's variables, we use the following. In this case there is only one because the file only shows the sea surface.

```
[4]: sst.levels
```

```
[4]: [0.0]
```

If we want to know where the dataset's NetCDF files are stored we can do the following:

```
[5]: sst.current
```

```
[5]: 'https://psl.noaa.gov/thredds/dodsC/Datasets/COBE/sst.mon.ltm.1981-2010.nc'
```

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

```
[6]: sst.times
```

```
[6]: ['0001-01-01T00:00:00',  
      '0001-02-01T00:00:00',  
      '0001-03-01T00:00:00',  
      '0001-04-01T00:00:00',  
      '0001-05-01T00:00:00',  
      '0001-06-01T00:00:00',  
      '0001-07-01T00:00:00',  
      '0001-08-01T00:00:00',  
      '0001-09-01T00:00:00',  
      '0001-10-01T00:00:00',  
      '0001-11-01T00:00:00',  
      '0001-12-01T00:00:00']
```

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

```
[7]: sst.months
```

```
[7]: [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12]
```

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

```
[8]: sst.years
```

```
[8]: [1]
```

If we do anything to the dataset, things will change. Let's calculate the average temperature:

```
[9]: sst.mean()
```

We can see that there is now a new temporary file associated with the dataset:

```
[10]: sst.current
```

```
[10]: '/tmp/nctoolkityjoudsyznctoolkittmpm3zaqqam.nc'
```

We can also access the history of operations carried out on the dataset:

```
[11]: sst.history
```

```
[11]: ['cdo -L -timmean https://psl.noaa.gov/thredds/dodsC/Datasets/COBE/sst.mon.ltm.1981-  
      ↪2010.nc /tmp/nctoolkityjoudsyznctoolkittmpm3zaqqam.nc']
```

Behind the scenes, nctoolkit mostly uses Climate Data Operators. If you are not familiar with Climate Data Operators, you can almost certainly just ignore the operations history.

We can also see that the times have changed. The only month now available is June, which is the mid-point of the year.

```
[12]: sst.months
```

```
[12]: [6]
```

2.3.3 Lazy evaluation of datasets

The code below will calculate the average sea surface temperature for a region in the North Atlantic for January. It does not do it efficiently.

```
[13]: sst = nc.open_thredds("https://psl.noaa.gov/thredds/dodsC/Datasets/COBE/sst.mon.ltm.
↪1981-2010.nc")
sst.select_months(1)
sst.clip(lon = [-80, 20], lat = [30, 70])
sst.spatial_mean()
```

If we look at the operation history, we see that temporary files have been created 3 times. Why not just once? We can do this by setting evaluation to lazy and then using `run` to evaluate everything when we need to.

```
[14]: sst.history
[14]: ['cdo -L -selmonth,1 https://psl.noaa.gov/thredds/dodsC/Datasets/COBE/sst.mon.ltm.
↪1981-2010.nc /tmp/nctoolkityjoudsyznctoolkittmpwcm57jx.nc',
'cdo -L -sellonlatbox,-80,20,30,70 /tmp/nctoolkityjoudsyznctoolkittmpwcm57jx.nc /
↪tmp/nctoolkityjoudsyznctoolkittmp34c020st.nc',
'cdo -L -fldmean /tmp/nctoolkityjoudsyznctoolkittmp34c020st.nc /tmp/
↪nctoolkityjoudsyznctoolkittmpwey97q4f.nc']
```

```
[15]: nc.options(lazy = True)
sst = nc.open_thredds("https://psl.noaa.gov/thredds/dodsC/Datasets/COBE/sst.mon.ltm.
↪1981-2010.nc")
sst.select_months(1)
sst.clip(lon = [-80, 20], lat = [30, 70])
sst.spatial_mean()
sst.run()
```

We can now see that only one temporary file was created

```
[16]: sst.history
[16]: ['cdo -L -fldmean -sellonlatbox,-80,20,30,70 -selmonth,1 https://psl.noaa.gov/
↪thredds/dodsC/Datasets/COBE/sst.mon.ltm.1981-2010.nc /tmp/
↪nctoolkityjoudsyznctoolkittmpdix9jo3.nc']
```

2.3.4 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:

```
[17]: sst = nc.open_thredds("https://psl.noaa.gov/thredds/dodsC/Datasets/COBE/sst.mon.ltm.
↪1981-2010.nc")
sst.select_months(1)
sst.clip(lon = [-80, 20], lat = [30, 70])
sst.plot()
```

```

Unable to decode time axis into full numpy.datetime64 objects, continuing using_
↳cftime.datetime objects instead, reason: dates out of range
Unable to decode time axis into full numpy.datetime64 objects, continuing using_
↳cftime.datetime objects instead, reason: dates out of range
Unable to decode time axis into full numpy.datetime64 objects, continuing using_
↳cftime.datetime objects instead, reason: dates out of range
Unable to decode time axis into full numpy.datetime64 objects, continuing using_
↳cftime.datetime objects instead, reason: dates out of range
The global colormaps dictionary is no longer considered public API.
The global colormaps dictionary is no longer considered public API.
The global colormaps dictionary is no longer considered public API.

```

```

[17]: :DynamicMap    [Variable,time]
      :Image      [lon,lat]    (sst)

```

To see how to use all of nctool's methods, check out the options on the left panel.

2.4 Importing and exporting data

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

2.4.1 Opening single files and ensembles

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

```

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

```

2.4.2 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
data = nc.open_url(example_url)

```


2.4.3 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
data = nc.open_thredds(example_url)
```

2.4.4 Exporting datasets

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

2.4.5 Save as a NetCDF

The method *write_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 *write_nc*. An example of usage is as follows:

```
data = nc.open_data(infile)
data.mean()
data.zip()
data.write_nc(outfile)
```

2.4.6 Convert to xarray Dataset

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

```
data = nc.open_data(infile)
data.annual_mean()
ds = data.to_xarray()
```

2.4.7 Convert to pandas dataframe

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

```
data = nc.open_data(infile)
data.annual_mean()
df = data.to_dataframe()
```

2.5 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
data = nc.open_data("sst.mon.mean.nc")
data.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:

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

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

```
data.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.

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

2.5.1 Calculating climatologies

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

```
data.tmean("season")
```

These methods all 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:

```
data.tmean("Seas")
```

Calculating a climatological monthly mean would require the following:

```
data.tmean("month")
```

and daily would be the following:

```
data.tmean("day")
```

2.5.2 Cumulative sums

We can calculate the cumulative sum as follows:

```
data.tcumsum()
```

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

2.6 Ensemble methods

2.6.1 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
```

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

```
[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 nans 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.annual_mean()
      all_sst.rolling_mean(10)
```

```
[12]: all_sst.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:

```
[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.annual_mean()
      all_sst.rolling_mean(10)
      all_sst.plot()
```

Data type cannot be displayed:

```
[13]: :Curve    [time]    (x)
```

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

2.6.2 Merging files with different times

TBC

2.6.3 Ensemble averaging

TBC

2.7 Speeding up code

2.7.1 Lazy evaluation

Under the hood nctoolkit relies mostly on CDO to carry out the specified manipulation of netcdf files. Each time CDO is called a new temporary file is generated. This has the potential to result in slower than necessary processing chains, as IO takes up far too much time.

I will demonstrate this using a netcdf file of sea surface temperature. To download the file we can just use wget:

```
[1]: import nctoolkit as nc
import warnings
warnings.filterwarnings('ignore')
from IPython.display import clear_output
!wget ftp://ftp.cdc.noaa.gov/Datasets/COBE2/sst.mon.ltm.1981-2010.nc
clear_output()
```

We can then set up the dataset which we will use for manipulating the SST climatology.

```
[2]: ff = "sst.mon.ltm.1981-2010.nc"
sst = nc.open_data(ff)
```

Now, let's select the variable sst, clip the file to the northern hemisphere, calculate the mean value in each grid cell for the first half of the year, and then calculate the spatial mean.

```
[3]: sst.select_variables("sst")
sst.clip(lat = [0,90])
sst.select_months(list(range(1,7)))
sst.mean()
sst.spatial_mean()
```

The dataset's history is as follows:

```
[4]: sst.history
[4]: ['cdo -L -selname,sst sst.mon.ltm.1981-2010.nc /tmp/
↪nctoolkitqhgujflsnctoolkittmpipj7up11.nc',
'cdo -L -sellonlatbox,-180,180,0,90 /tmp/nctoolkitqhgujflsnctoolkittmpipj7up11.nc /
↪tmp/nctoolkitqhgujflsnctoolkittmp920v1_r7.nc',
'cdo -L -selmonth,1,2,3,4,5,6 /tmp/nctoolkitqhgujflsnctoolkittmp920v1_r7.nc /tmp/
↪nctoolkitqhgujflsnctoolkittmpbnck_dy2.nc',
'cdo -L -timmean /tmp/nctoolkitqhgujflsnctoolkittmpbnck_dy2.nc /tmp/
↪nctoolkitqhgujflsnctoolkittmpjmzt1167.nc',
'cdo -L -fldmean /tmp/nctoolkitqhgujflsnctoolkittmpjmzt1167.nc /tmp/
↪nctoolkitqhgujflsnctoolkittmpdus63y8i.nc']
```

In total, there are 5 operations, with temporary files created each time. However, we only want to generate one temporary file. So, can we do that? Yes, thanks to CDO's method chaining ability. If we want to utilize this we need to set the session's evaluation to lazy, using options. Once this is done nctoolkit will only evaluate things either when it needs to, e.g. you call a method that cannot possibly be chained, or if you evaluate it using run. This works as follows:

```
[5]: ff = "sst.mon.ltm.1981-2010.nc"
nc.options(lazy = True)
sst = nc.open_data(ff)
sst.select_variables("sst")
sst.clip(lat = [0,90])
sst.select_months(list(range(1,7)))
sst.mean()
sst.spatial_mean()
sst.run()
```

We can now see that the history is much cleaner, with only one command.

```
[6]: sst.history
[6]: ['cdo -L -fldmean -timmean -selmonth,1,2,3,4,5,6 -sellonlatbox,-180,180,0,90 -
↳selname,sst sst.mon.ltm.1981-2010.nc /tmp/nctoolkitqhgujflsnctoolkittmpkdkiewey2.nc']
```

How does this impact run time? Let's time the original, unchained method.

```
[7]: %%time
nc.options(lazy = False)
ff = "sst.mon.ltm.1981-2010.nc"
sst = nc.open_data(ff)
sst.select_variables("sst")
sst.clip(lat = [0,90])
sst.select_months(list(range(1,7)))
sst.mean()
sst.spatial_mean()

CPU times: user 37.2 ms, sys: 61.6 ms, total: 98.7 ms
Wall time: 667 ms
```

```
[8]: %%time
nc.options(lazy = True)
ff = "sst.mon.ltm.1981-2010.nc"
sst = nc.open_data(ff)
sst.select_variables("sst")
sst.clip(lat = [0,90])
sst.select_months(list(range(1,7)))
sst.mean()
sst.spatial_mean()
sst.run()

CPU times: user 17.3 ms, sys: 4.28 ms, total: 21.6 ms
Wall time: 161 ms
```

This was almost 4 times faster. Exact speed improvements, will of course depend on specific IO requirements, and some times using lazy evaluation will make negligible impact, but in others can make code over 10 times faster. Exact speed improvements, will of course depend on specific IO requirements, and some times using lazy evaluation will make negligible impact, but in others can make code over 10 times faster.

2.7.2 Processing files in parallel

When processing a dataset made up of multiple files, it is possible carry out the processing in parallel for more or less all of the methods available in nctoolkit. To carry out processing in parallel with 6 cores, we would use options as follows:

```
[9]: nc.options(cores = 6)
```

By default the number of cores in use is 1. Of course, this can result in you crashing your computer if the total RAM in use is excessive, so it's best practise to check RAM used with one core first.

2.7.3 Using thread-safe libraries

If the CDO installation being called by nctoolkit is compiled with threadsafe hdf5, then you can achieve potentially significant speed ups with the following command:

```
[10]: nc.options(thread_safe = True)
```

If you are not sure, if hdf5 has been built thread safe, a simple way to find this out is to run the code below. If it fails, you can be more or less certain it is not threadsafe.

```
[11]: nc.options(lazy = True)
nc.options(thread_safe = True)
ff = "sst.mon.ltm.1981-2010.nc"
sst = nc.open_data(ff)
sst.select_variables("sst")
sst.clip(lat = [0,90])
sst.select_months(list(range(1,7)))
sst.mean()
sst.spatial_mean()
sst.run()
```

Reference and help

- [An A-Z guide to nctoolkit methods](#)
- [API Reference](#)
- [How to guide](#)
- [Package info](#)

2.8 An A-Z guide to nctoolkit methods

This guide will provide examples of how to use almost every method available in nctoolkit.

2.8.1 add

This method can add to a dataset. You can add a constant, another dataset or a NetCDF file. In the case of datasets or NetCDF files the grids etc. must be of the same structure as the original dataset.

For example, if we had a temperature dataset where temperature was in Celsius, we could convert it to Kelvin by adding 273.15.

```
data.add(273.15)
```

If we have two datasets, we add one to the other as follows:

```
data1 = nc.open_data(infile1)
data2 = nc.open_data(infile2)
data1.add(data2)
```

In the above example, all we are doing is adding infile2 to data2, so instead we could simply do this:

```
data1.add(infile2)
```

2.8.2 annual_anomaly

This method will calculate the annual anomaly for each variable (and in each grid cell) compared with a baseline. This is a standard anomaly calculation where first the mean value is calculated for the baseline period, and the difference between the values is calculated.

For example, if we wanted to calculate the anomalies in a dataset compared with a baseline period of 1900-1919 we would do the following:

```
data.annual_anomaly(baseline=[1900, 1919])
```

We may be more interested in the rolling anomaly, in particular when there is a lot of annual variation. In the above case, if you wanted a 20 year rolling mean anomaly, you would do the following:

```
data.annual_anomaly(baseline=[1900, 1919], window=20)
```

By default this method works out the absolute anomaly. However, in some cases the relative anomaly is more interesting. To calculate this we set the metric argument to “relative”:

```
data.annual_anomaly(baseline=[1900, 1919], metric = "relative")
```

2.8.3 annual_max

This method will calculate the maximum value in each available year and for each grid cell of dataset.

```
data.annual_max()
```


2.8.4 annual_mean

This method will calculate the maximum value in each available year and for each grid cell of dataset.

```
data.annual_mean()
```

2.8.5 annual_min

This method will calculate the minimum value in each available year and for each grid cell of dataset.

```
data.annual_min()
```

2.8.6 annual_range

This method will calculate the range of values in each available year and for each grid cell of dataset.

```
data.annual_range()
```

2.8.7 annual_sum

This method will calculate the sum of values in each available year and for each grid cell of dataset.

```
data.annual_sum()
```

2.8.8 append

This method will let you append individual or multiple files to your dataset. Usage is straightforward. Note that this will not perform any merging on the dataset.

```
data.append(newfile)
```

2.8.9 bottom

This method will extract the bottom vertical level from a dataset. This is useful for some oceanographic datasets, where the method can let you select the seabed. Note that this method will not work with all data types. For example, in ocean data with fixed depth levels, the bottom cell in the NetCDF data is not the actual seabed. See `bottom_mask` for these cases.

```
data.bottom()
```

2.8.10 bottom_mask

This method will identify the bottommost level in each grid with a non-NA value.

```
data.bottom_mask()
```

2.8.11 cdo_command

This method let's you run a cdo command. CDO commands are generally of the form “cdo {command} infile outfile”. cdo_command therefore only requires the command portion of this. If we wanted to run the following CDO command

```
cdo -timmean -selmon,4 infile outfile
```

we would do the following:

```
data.cdo_command("-timmean -selmon,4")
```

2.8.12 cell_areas

This method either adds the areas of each grid cell to the dataset or converts the dataset to a new dataset showing only the grid cell areas. By default it adds the cell areas (in square metres) to the dataset.

```
data.cell_areas()
```

If we only want the cell areas we can set join to False:

```
data.cell_areas(join=False)
```

2.8.13 centre

This method calculates the longitudinal or latitudinal centre of a dataset. There is one argument, which should either be “latitude” or “longitude”. If you want to calculate the latitudinal centre:

```
data.centre("longitude")
```

2.8.14 crop

This method will crop a region to a specified longitude and latitude box. For example, if we wanted to crop a dataset to the North Atlantic, we could do this:

```
data.crop(lon = [-80, 20], lat = [40, 70])
```

2.8.15 compare_all

This method let's us compare all variables in a dataset with a constant. If we wanted to identify the grid cells with values above 20, we could do the following:

```
data.compare_all(">20")
```

Similarly, if we wanted to identify grid cells with negative values we would do this:

```
data.compare_all("<0")
```

2.8.16 cor_space

This method calculates the correlation coefficients between two variables in space for each time step. So, if we wanted to work out the correlation between the variables var1 and var2, we would do this:

```
data.cor_space("var1", "var2")
```

2.8.17 cor_time

This method calculates the correlation coefficients between two variables in time for each grid cell. If we wanted to work out the correlation between two variables var1 and var2 we would do the following:

```
data.cor_time("var1", "var2")
```

2.8.18 cum_sum

This method will calculate the cumulative sum, over time, for all variables. Usage is simple:

```
data.cum_sum()
```

2.8.19 daily_max

This method will calculate the maximum value in each available day and for each grid cell of dataset.

```
data.daily_max()
```

2.8.20 daily_mean

This method will calculate the maximum value in each available day and for each grid cell of dataset.

```
data.daily_mean()
```

2.8.21 daily_min

This method will calculate the minimum value in each available day and for each grid cell of dataset.

```
data.daily_min()
```

2.8.22 daily_range

This method will calculate the range of values in each available day and for each grid cell of dataset.

```
data.daily_range()
```

2.8.23 daily_sum

This method will calculate the sum of values in each available day and for each grid cell of dataset.

```
data.daily_sum()
```

2.8.24 daily_max_climatology

This method will calculate the maximum value that is observed on each day of the year over time. So, for example, if you had 100 years of daily temperature data, it will calculate the maximum value ever observed on each day.

```
data.daily_max_climatology()
```

2.8.25 daily_mean_climatology

This method will calculate the mean value that is observed on each day of the year over time. So, for example, if you had 100 years of daily temperature data, it will calculate the mean value ever observed on each day.

```
data.daily_mean_climatology()
```

2.8.26 daily_min_climatology

This method will calculate the minimum value that is observed on each day of the year over time. So, for example, if you had 100 years of daily temperature data, it will calculate the minimum value ever observed on each day.

```
data.daily_min_climatology()
```

2.8.27 daily_range_climatology

This method will calculate the value range that is observed on each day of the year over time. So, for example, if you had 100 years of daily temperature data, it will calculate the difference between the maximum and minimum observed values each day.

```
:: data.daily_range_climatology()
```

2.8.28 divide

This method will divide a dataset by a constant, or the values in another dataset of NetCDF file. If we wanted to divide everything in a dataset by 2, we would do the following:

```
data.divide(2)
```

If we want to divide a dataset by another, we can do this easily. Note that the datasets must be comparable, i.e. they must have the same grid. The second dataset must have either the same number of variables or only one variable. In the latter case everything is divided by that variable. The same holds for vertical levels.

```
data1 = nc.open_data(infile1)
data2 = nc.open_data(infile2)
data1.divide(data2)
```

2.8.29 ensemble_max, ensemble_min, ensemble_range and ensemble_mean

These methods will calculate the ensemble statistic, when a dataset is made up of multiple files. Two methods are available. First, the statistic across all available time steps can be calculated. For this `ignore_time` must be set to `False`. For example:

```
data = nc.open_data(file_list)
data.ensemble_max(ignore_time = True)
```

The second method is to calculate the maximum value in each given time step. For example, if the ensemble was made up of 100 files where each file contains 12 months of data, `ensemble_max` will work out the maximum monthly value. By default `ignore_time` is `False`.

```
data = nc.open_data(file_list)
data.ensemble_max(ignore_time = False)
```

2.8.30 ensemble_percentile

This method works in the same way as `ensemble_mean` etc. above. However, it requires an additional term `p`, which is the percentile. For example, if we had to calculate the 75th ensemble percentile, we would do the following:

```
data = nc.open_data(file_list)
data = nc.ensemble_percentile(75)
```

2.8.31 format

This method will change the format of the files within a dataset. For example if you wanted to convert to NetCDF4:

```
data.format("nc4")
```

2.8.32 invert_levels

This method will invert the vertical levels of a dataset.

```
data.invert_levels()
```

2.8.33 mask_box

This method will set everything outside a specified longitude/latitude box to NA. The code below illustrates how to mask the North Atlantic in the SST dataset.

```
data.mask_box(lon = [-80, 20], lat = [40, 70])
```

2.8.34 max

This method will calculate the maximum value of all variables in all grid cells. If we wanted to calculate the maximum observed monthly sea surface temperature in the SST dataset we would do the following:

```
data.max()
```

2.8.35 mean

This method will calculate the mean value (averaged across all time steps) of all variables in all grid cells. Usage is simple:

```
data.mean()
```

2.8.36 median

This method will calculate the median value (averaged across all time steps) of all variables in all grid cells. Usage is simple:

```
data.median()
```

2.8.37 merge and merge_time

nctoolkit offers two methods for merging the files within a multi-file dataset. These methods operate in a similar way to column based joining and row-based binding in dataframes.

The merge method is suitable for merging files that have different variables, but the same time steps. The merge_time method is suitable for merging files that have the same variables, but have different time steps.

Usage for merge_time is as simple as:

```
data = nc.open_data(file_list)
data.merge_time()
```

Merging NetCDF files with different variables is potentially risky, as it is possible you can merge files that have the same number of time steps but have different times. nctoolkit's merge method therefore offers some security against a major error when merging. It requires a match argument to be supplied. This ensures that the times in each file is comparable to the others. By default match = ["year", "month", "day"], i.e. it checks if the times in each file all have the same year, month and day. The match argument must be some subset of ["year", "month", "day"]. For example, if you wanted to only make sure the files had the same year, you would do the following:

```
data = nc.open_data(file_list)
data.merge(match = ["year", "month", "day"])
```

2.8.38 meridional statistics

Calculate the following meridional statistics: mean, min, max and range:

```
data.meridional_mean()
data.meridional_min()
data.meridional_max()
data.meridional_range()
```

2.8.39 min

This method will calculate the minimum value (across all time steps) of all variables in all grid cells. Usage is simple:

```
data.min()
```

2.8.40 monthly_anomaly

This method will calculate the monthly anomaly compared with the mean value for a baseline period. For example, if we wanted the monthly anomaly compared with the mean for 1990-1999 we would do the below.

```
data.monthly_anomaly(baseline = [1990, 1999])
```

2.8.41 monthly_max

This method will calculate the maximum value in the month of each year of a dataset. This is useful for daily time series. If you want to calculate the mean value in each month across all available years, use `monthly_max_climatology`. Usage is simple:

```
data.monthly_max()
```

2.8.42 monthly_max_climatology

This method will calculate, for each month, the maximum value of each variable over all time steps.

```
data.monthly_max_climatology()
```

2.8.43 monthly_mean

This method will calculate the mean value of each variable in each month of a dataset. Note that this is calculated for each year. See `monthly_mean_climatology` if you want to calculate a climatological monthly mean.

```
data.monthly_mean()
```

2.8.44 monthly_mean_climatology

This method will calculate, for each month, the maximum value of each variable over all time steps. Usage is simple:

```
data.monthly_mean_climatology()
```

2.8.45 monthly_min

This method will calculate the minimum value in the month of each year of a dataset. This is useful for daily time series. If you want to calculate the mean value in each month across all available years, use `monthly_max_climatology`. Usage is simple:

```
data.monthly_min()
```

2.8.46 monthly_min_climatology

This method will calculate, for each month, the minimum value of each variable over all time steps. Usage is simple:

```
data.monthly_min_climatology()
```


2.8.47 monthly_range

This method will calculate the value range in the month of each year of a dataset. This is useful for daily time series. If you want to calculate the value range in each month across all available years, use `monthly_range_climatology`. Usage is simple:

```
data.monthly_range()
```

2.8.48 monthly_range_climatology

This method will calculate, for each month, the value range of each variable over all time steps. Usage is simple:

```
data.monthly_range_climatology()
```

2.8.49 multiply

This method will multiply a dataset by a constant, another dataset or a NetCDF file. If multiplied by a dataset or NetCDF file, the dataset must have the same grid and can only have one variable.

If you want to multiply a dataset by 2, you can do the following:

```
data.multiply(2)
```

If you wanted to multiply a dataset `data1` by another, `data2`, you can do the following:

```
data1 = nc.open_data(infile1)
data2 = nc.open_data(infile2)
data1.multiply(data2)
```

2.8.50 mutate

This method can be used to generate new variables using arithmetic expressions. New variables are added to the dataset. The method requires a dictionary, where the key-value pairs are the new variables and expression required to generate it.

For example, if had a temperature dataset, with temperature in Celsius, we might want to convert that to Kelvin. We can do this easily:

```
data.mutate({"temperature_k": "temperature+273.15"})
```

2.8.51 percentile

This method will calculate a given percentile for each variable and grid cell. This will calculate the percentile using all available timesteps.

We can calculate the 75th percentile of sea surface temperature as follows:

```
data.percentile(75)
```

2.8.52 phenology

A number of phenological indices can be calculated. These are based on the plankton metrics listed by [Ji et al. 2010](#). These methods require datasets or the files within a dataset to only be made up of individual years, and ideally every day of year is available. At present this method can only calculate the phenology metric for a single variable.

The available metrics are: peak - the time of year when the maximum value of a variable occurs. middle - the time of year when 50% of the annual cumulative sum of a variable is first exceeded start - the time of year when a lower threshold (which must be defined) of the annual cumulative sum of a variable is first exceeded end - the time of year when an upper threshold (which must be defined) of the annual cumulative sum of a variable is first exceeded

For example, if you wanted to calculate timing of the peak, you set metric to “peak”, and define the variable to be analyzed:

```
data.phenology(metric = "peak", var = "var_chosen")
```

2.8.53 plot

This method will plot the contents of a dataset. It will either show a map or a time series, depending on the data type. While it should work on at least 90% of NetCDF data, there are some data types that remain incompatible, but will be added to nctoolkit over time. Usage is simple:

```
data.plot()
```

2.8.54 range

This method calculates the range for all variables in each grid cell across all steps.

We can calculate the range of sea surface temperatures in the SST dataset as follows:

```
data.range()
```

2.8.55 regrid

This method will remap a dataset to a new grid. This grid must be either a pandas data frame, a NetCDF file or a single file nctoolkit dataset.

For example, if we wanted to regrid a dataset to a single location, we could do the following:

```
import pandas as pd
data = nc.open_data(infile)
grid = pd.DataFrame({"lon": [-20], "lat": [50]})
data.regrid(grid, method = "nn")
```

If we wanted to regrid one dataset, dataset1, to the grid of another, dataset2, using bilinear interpolation, we would do the following:

```
data1 = nc.open_data(infile1)
data2 = nc.open_data(infile2)
data1.regrid(data2, method = "bil")
```

2.8.56 remove_variables

This method will remove variables from a dataset. Usage is simple, with the method only requiring either a str of a single variable or a list of variables to remove:

```
data.remove_variables(vars)
```

2.8.57 rename

This method allows you to rename variables. It requires a dictionary, with key-value pairs representing the old variable names and new variables. For example, if we wanted to rename a variable old to new, we would do the following:

```
data.rename({"old": "new"})
```

2.8.58 resample_grid

This method let's you resample the horizontal grid. It takes one argument. If you wanted to only take every other grid cell, you would do the following:

```
data.resample_grid(2)
```

2.8.59 rolling_max

This method will calculate the rolling maximum over a specified window. For example, if you needed to calculate the rolling maximum with a window of 10, you would do the following:

```
data.rolling_max(window = 10)
```

2.8.60 rolling_mean

This method will calculate the rolling mean over a specified window. For example, if you needed to calculate the rolling mean with a window of 10, you would do the following:

```
data.rolling_mean(window = 10)
```

2.8.61 rolling_min

This method will calculate the rolling minimum over a specified window. For example, if you needed to calculate the rolling minimum with a window of 10, you would do the following:

```
data.rolling_min(window = 10)
```

2.8.62 rolling_range

This method will calculate the rolling range over a specified window. For example, if you needed to calculate the rolling range with a window of 10, you would do the following:

```
data.rolling_range(window = 10)
```

2.8.63 rolling_sum

This method will calculate the rolling sum over a specified window. For example, if you needed to calculate the rolling sum with a window of 10, you would do the following:

```
data.rolling_sum(window = 10)
```

2.8.64 run

This method will evaluate all of a dataset's unevaluated commands. Evaluation should be set to lazy. Usage is simple:

```
nc.options(lazy = True)
data = nc.open_data(infile)
#.... apply some methods to the dataset
data.run()
```

2.8.65 seasonal_max

This method will calculate the maximum value observed in each season. Note this is worked out for the seasons of each year. See `seasonal_max_climatology` for climatological seasonal maximums.

```
data.seasonal_max()
```

2.8.66 seasonal_max_climatology

This method calculates the maximum value observed in each season across all years. Usage is simple:

```
data.seasonal_max_climatology()
```

2.8.67 seasonal_mean

This method will calculate the mean value observed in each season. Note this is worked out for the seasons of each year. See `seasonal_mean_climatology` for climatological seasonal means.

```
data.seasonal_mean()
```

2.8.68 seasonal_mean_climatology

This method calculates the mean value observed in each season across all years. Usage is simple:

```
data.seasonal_mean_climatology()
```

2.8.69 seasonal_min

This method will calculate the minimum value observed in each season. Note this is worked out for the seasons of each year. See `seasonal_min_climatology` for climatological seasonal minimums.

```
data.seasonal_min()
```

2.8.70 seasonal_min_climatology

This method calculates the minimum value observed in each season across all years. Usage is simple:

```
data.seasonal_min_climatology()
```

2.8.71 seasonal_range

This method will calculate the value range observed in each season. Note this is worked out for the seasons of each year. See `seasonal_range_climatology` for climatological seasonal ranges.

```
data.seasonal_range()
```

2.8.72 seasonal_range_climatology

This method calculates the value range observed in each season across all years. Usage is simple:

```
data.seasonal_range_climatology()
```

2.8.73 select

A method to subset a dataset based on multiple criteria. This acts as a wrapper for `select_variables`, `select_months`, `select_years`, `select_seasons`, and `select_timesteps`, with the args used being `variables`, `months`, `years`, `seasons`, and `timesteps`. Subsetting will occur in the order given. For example, if you want to select the years 1990 and 1991 and months June and July, you would do the following:

```
data.select(years = [1990, 1991], months = [6, 7])
```

2.8.74 select_months

This method allows you to subset a dataset to specific months. This can either be a single month, a list of months or a range. For example, if we wanted the first half of a year, we would do the following:

```
data.select_months(range(1, 7))
```

2.8.75 select_variables

This method allows you to subset a dataset to specific variables. This either accepts a single variable or a list of variables. For example, if you wanted two variables, var1 and var2, you would do the following:

```
data.select_variables(["var1", "var2"])
```

2.8.76 select_years

This method subsets datasets to specified years. It will accept either a single year, a list of years, or a range. For example, if you wanted to subset a dataset the 1990s, you would do the following:

```
data.select_years(range(1990, 2000))
```

2.8.77 set_missing

This method allows you to set a range to missing values. It either accepts a single variable or two variables, specifying the range to be set to missing values. For example, if you wanted all values between 0 and 10 to be set to missing, you would do the following:

```
data.set_missing([0, 10])
```

2.8.78 shift_days

This method allows you to shift time by a set number of hours, days, months or years. This acts as a wrapper for *shift_hours*, *shift_days*, *shift_months* and *shift_years*. Use the args *hours*, *days*, *months*, or *years*. This takes any number of arguments. So, if you wanted to shift time forward by 1 year, 1 month and 1 days you would do the following:

```
data.shift(years = 1, months = 1, days = 1)
```

2.8.79 shift_days

This method allows you to shift time by a set number of days. For example, if you want time moved forward by 2 hours you would do the following:

```
data.shift_days(2)
```

2.8.80 shift_hours

This method allows you to shift time by a set number of hours. For example, if you want time moved back by 1 hour you would do the following:

```
data.shift_hours(-1)
```

2.8.81 shift_months

This method allows you to shift time by a set number of months. For example, if you want time moved back by 2 months you would do the following:

```
data.shift_months(2)
```

2.8.82 shift_years

This method allows you to shift time by a set number of years. For example, if you want time moved back by 10 years you would do the following:

```
data.shift_years(10)
```

2.8.83 spatial_max

This method will calculate the maximum value observed in space for each variable and time step. Usage is simple:

```
data.spatial_max()
```

2.8.84 spatial_mean

This method will calculate the spatial mean for each variable and time step. If the grid cell area can be calculated, this will be an area weighted mean. Usage is simple:

```
data.spatial_mean()
```

2.8.85 spatial_min

This method will calculate the minimum observed in space for each variable and time step. Usage is simple:

```
data.spatial_min()
```

2.8.86 spatial_percentile

This method will calculate the percentile of variable across space for time step. For example, if you wanted to calculate the 75th percentile, you would do the following:

```
data.spatial_percentile(p=75)
```

2.8.87 spatial_range

This method will calculate the value range observed in space for each variable and time step. Usage is simple:

```
data.spatial_range()
```

2.8.88 spatial_sum

This method will calculate the spatial sum for each variable and time step. In some cases, for example when variables are concentrations, it makes more sense to multiply the value in each grid cell by the grid cell area, when doing a spatial sum. This method therefore has an argument `by_area` which defines whether to multiply the variable value by the area when doing the sum. By default `by_area` is `False`.

Usage is simple:

```
data.spatial_sum()
```

2.8.89 split

Except for methods that begin with `merge` or `ensemble`, all `nctoolkit` methods operate on individual files within a dataset. There are therefore cases when you might want to be able to split a dataset into separate files for analysis. This can be done using `split`, which let's you split a file into separate years, months or year/month combinations. For example, if you want to split a dataset into files of different years, you can do this:

```
data.split("year")
```

2.8.90 subtract

This method can subtract from a dataset. You can subtract a constant, another dataset or a NetCDF file. In the case of datasets or NetCDF files the grids etc. must be of the same structure as the original dataset.

For example, if we had a temperature dataset where temperature was in Kelvin, we could convert it to Celsius by subtracting 273.15.

```
data.subtract(273.15)
```


2.8.91 sum

This method will calculate the sum of values of all variables in all grid cells. Usage is simple:

```
data.sum()
```

sum_all —

This method will calculate the sum of all variables separately for each time cell and grid cell. Usage is simple:

```
data.sum_all()
```

2.8.92 surface

This method will extract the surface level from a multi-level dataset. Usage is simple:

```
data.surface()
```

2.8.93 to_dataframe

This method will return a pandas dataframe with the contents of the dataset. This has a decode_times argument to specify whether you want the times to be decoded. Defaults to True. Usage is simple:

```
data.to_dataframe()
```

2.8.94 to_latlon

This method will regrid a dataset to a regular latlon grid. The minimum and maximum longitudes and latitudes must be specified, along with the horizontal and vertical resolutions.

```
data.to_latlon(lon = [-80, 20], lat = [30, 80], res = [1,1])
```

2.8.95 to_xarray

This method will return an xarray dataset with the contents of the dataset. This has a decode_times argument to specify whether you want the times to be decoded. Defaults to True. Usage is simple:

```
data.to_xarray()
```

2.8.96 transmute

This method can be used to generate new variables using arithmetic expressions. Existing will be removed from the dataset. See mutate if you want to keep existing variables. The method requires a dictionary, where the key-value pairs are the new variables and expression required to generate it.

For example, if had a temperature dataset, with temperature in Celsius, we might want to convert that to Kelvin. We can do this easily:

```
data.transmute({"temperature_k": "temperature+273.15"})
```

2.8.97 var

This method calculates the variance of each variable in the dataset. This is calculate across all time steps. Usage is simple:

```
data.var()
```

2.8.98 vertical_interp

This method interpolates variables vertically. It requires a list of vertical levels, for example depths, you want to interpolate. For example, if you had an ocean dataset and you wanted to interpolate to 10 and 20 metres you would do the following:

```
data.vertical_interp(levels = [10, 20])
```

2.8.99 vertical_max

This method calculates the maximum value of each variable across all vertical levels. Usage is simple:

```
data.vertical_max()
```

2.8.100 vertical_mean

This method calculates the mean value of each variable across all vertical levels. Usage is simple:

```
data.vertical_mean()
```

2.8.101 vertical_min

This method calculates the minimum value of each variable across all vertical levels. Usage is simple:

```
data.vertical_min()
```

2.8.102 vertical_range

This method calculates the value range of each variable across all vertical levels. Usage is simple:

```
data.vertical_range()
```

2.8.103 vertical_sum

This method calculates the sum each variable across all vertical levels. Usage is simple:

```
data.vertical_sum()
```

2.8.104 to_nc

This method allows you to write the contents of a dataset to a NetCDF file. If the target file exists and you want to overwrite it set `overwrite` to `True`. Usage is simple:

```
data.to_nc(outfile)
```

2.8.105 zip

This method will zip the contents of a dataset. This is mostly useful for processing chains where you want to minimize disk space usage by the output. Please note this method works lazily. In the code below only one file is generated, a zipped “outfile”.

```
nc.options(lazy = True)
data = nc.open_data(infile)
data.select_years(1990)
data.zip()
data.write_nc(outfile)
```

2.8.106 zonal statistics

Calculate the following zonal statistics: mean, min, max and range:

```
data.zonal_mean()
data.zonal_min()
data.zonal_max()
data.zonal_range()
```

2.9 How to guide

This guide will show how to carry out key nctoolkit operations. We will use a sea surface temperature data set and a depth-resolved ocean temperature data set. The data set can be downloaded from [here](#).

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

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

2.9.1 How to select years and months

If we want to select specific years and months we can use the `select_years` and `select_months` method

```
[2]: sst = nc.open_data("sst.mon.mean.nc")
      sst.select_years(1960)
      sst.select_months(1)
      sst.times

[2]: ['1960-01-01T00:00:00']
```

2.9.2 How to mean, mean, max etc.

If you want to calculate the mean value of a variable over all time steps you can use `mean`:

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

[3]: :DynamicMap      [time]
      :Image         [lon,lat]      (sst)
```

Similarly, if you want to calculate the minimum, maximum, sum and range of values over time just use `min`, `max`, `sum` and `range`.

2.9.3 How to copy a data set

If you want to make a deep copy of a data set, use the built in `copy` method. This method will return a new data set. This method should be used because of nctoolkit's built in methods to automatically delete temporary files that are no longer required. Behind the scenes, using `copy` will result in nctoolkit registering that it needs the NetCDF file for both the original dataset and the new copied one. So if you copy a dataset, and then delete the original, nctoolkit knows to not remove any NetCDF files related to the dataset.

```
[4]: sst = nc.open_data("sst.mon.mean.nc")
      sst.select_years(1960)
      sst.select_months(1)
      sst1 = sst.copy()
      del sst
      os.path.exists(sst1.current)

[4]: True
```

2.9.4 How to clip to a region

If you want to clip the data to a specific longitude and latitude box, we can use `clip`, with the longitude and latitude range given by `lon` and `lat`.

```
[5]: sst = nc.open_data("sst.mon.mean.nc")
      sst.select_months(1)
      sst.select_years(1980)
      sst.clip(lon = [-80, 20], lat = [40, 70])
      sst.plot()

[5]: :DynamicMap      [time]
      :Image          [lon,lat]      (sst)
```

2.9.5 How to rename a variable

If we want to rename a variable we use the `rename` method, and supply a dictionary where the key-value pairs are the original and new names

```
[6]: sst = nc.open_data("sst.mon.mean.nc")
      sst.variables

[6]: ['sst']
```

The original dataset had only one variable called `sst`. We can now rename it, and display the new variables.

```
[7]: sst.rename({"sst": "temperature"})
      sst.variables

[7]: ['temperature']
```

2.9.6 How to create new variables

New variables can be created using arithmetic operations using either `mutate` or `transmute`. The `mutate` method will maintain the original variables, whereas `transmute` will not. This method requires a dictionary, where the key, values pairs are the names of the new variables and the arithmetic operations to perform. The example below shows how to create a new variable with

```
[8]: sst = nc.open_data("sst.mon.mean.nc")
      sst.mutate({"sst_k": "sst+273.15"})
      sst.variables

[8]: ['sst', 'sst_k']
```

2.9.7 How to calculate a spatial average

You can calculate a spatial average using the `spatial_mean` method. There are additional methods for maximums etc.

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

[9]: :Curve      [time]      (x)
```

2.9.8 How to calculate an annual mean

You can calculate an annual mean using the `annual_mean` method.

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

```
[10]: :Curve    [time]    (x)
```

2.9.9 How to calculate a rolling average

You can calculate a rolling mean using the `rolling_mean` method, with the `window` argument providing the number of time steps to average over. There are additional methods for rolling sums etc. The code below will calculate a rolling mean of global SST using a 20 year window.

```
[11]: sst = nc.open_data("sst.mon.mean.nc")
      sst.spatial_mean()
      sst.annual_mean()
      sst.rolling_mean(20)
      sst.plot()
```

```
[11]: :Curve    [time]    (x)
```

2.9.10 How to calculate temporal anomalies

You can calculate annual temporal anomalies using the `annual_anomaly` method. This requires a baseline period.

```
[12]: sst = nc.open_data("sst.mon.mean.nc")
      sst.spatial_mean()
      sst.annual_anomaly(baseline = [1960, 1979])
      sst.plot()
```

```
[12]: :Curve    [time]    (x)
```

2.9.11 How to split data by year etc

Files within a dataset can be split by year, day, year and month or season using the `split` method. If we wanted to split by year, we do the following:

```
[13]: sst = nc.open_data("sst.mon.mean.nc")
      sst.split("year")
```

2.9.12 How to merge files in time

We can merge files based on time using `merge_time`. We can do this by merging the dataset that results from splitting the original sst dataset. If we split the dataset by year, we see that there are 169 files, one for each year.

```
[14]: sst = nc.open_data("sst.mon.mean.nc")
      sst.split("year")
```

We can then merge them together to get a single file dataset:

```
[15]: sst.merge_time()
```

2.9.13 How to do variables-based merging

If we have two more more files that have the same time steps, but different variables, we can merge them using `merge`. The code below will first create a dataset with a NetCDF file with SST in K, and it will then create a new dataset with this netcd file and the original, and then merge them.

```
[16]: sst1 = nc.open_data("sst.mon.mean.nc")
      sst2 = nc.open_data("sst.mon.mean.nc")
      sst2.transmute({"sst_k": "sst+273.15"})
      new_sst = nc.open_data([sst1.current, sst2.current])
      new_sst.current
      new_sst.merge()
```

In some cases we will have two or more datasets we want to merge. In this case we can use the `merge` function as follows:

```
[17]: sst1 = nc.open_data("sst.mon.mean.nc")
      sst2 = nc.open_data("sst.mon.mean.nc")
      sst2.transmute({"sst_k": "sst+273.15"})
      new_sst = nc.merge(sst1, sst2)
      new_sst.variables
```

```
[17]: ['sst', 'sst_k']
```

2.9.14 How to horizontally regrid data

Variables can be regridded horizontally using `regrid`. This method requires the new grid to be defined. This can either be a pandas data frame, with lon/lat as columns, an xarray object, a NetCDF file or an nctoolkit dataset. I will demonstrate all three methods by regridding SST to the North Atlantic. Let's begin by getting a grid for the North Atlantic.

```
[18]: new_grid = nc.open_data("sst.mon.mean.nc")
      new_grid.clip(lon = [-80, 20], lat = [30, 70])
      new_grid.select_months(1)
      new_grid.select_years(2000)
```

First, we will use the new dataset itself to do the regridding. I will calculate mean SST using the original data, and then regrid to the North Atlantic.

```
[19]: sst = nc.open_data("sst.mon.mean.nc")
      sst.mean()
      sst.regrid(grid = new_grid)
      sst.plot()
```

```
[19]: :DynamicMap    [time]
      :Image      [lon,lat]    (sst)
```

We can also do this using the NetCDF, which is `new_grid.current`

```
[20]: sst = nc.open_data("sst.mon.mean.nc")
      sst.mean()
      sst.regrid(grid = new_grid.current)
      sst.plot()
```

```
[20]: :DynamicMap    [time]
      :Image      [lon,lat]    (sst)
```

or we can use a pandas data frame. In this case I will convert the xarray data set to a data frame.

```
[21]: na_grid = xr.open_dataset(new_grid.current)
      na_grid = na_grid.to_dataframe().reset_index().loc[:,["lon", "lat"]]
      sst = nc.open_data("sst.mon.mean.nc")
      sst.mean()
      sst.regrid(grid = na_grid)
      sst.plot()
```

```
[21]: :DynamicMap    [time]
      :Image      [lon,lat]    (sst)
```

2.9.15 How to temporally interpolate

Temporal interpolation can be carried out using `time_interp`. This method requires a start date (start) of the format YYYY/MM/DD and an end date (end), and a temporal resolution (resolution), which is either 1 day (“daily”), 1 week (“weekly”), 1 month (“monthly”), or 1 year (“yearly”).

```
[22]: sst = nc.open_data("sst.mon.mean.nc")
      sst.time_interp(start = "1990/01/01", end = "1990/12/31", resolution = "daily")
```

2.9.16 How to calculate a monthly average from daily data

If you have daily data, you can calculate a month average using `monthly_mean`. There are also methods for maximums etc.

```
[23]: sst = nc.open_data("sst.mon.mean.nc")
      sst.time_interp(start = "1990/01/01", end = "1990/12/31", resolution = "daily")
      sst.monthly_mean()
```

2.9.17 How to calculate a monthly climatology

If we want to calculate the mean value of variables for each month in a given dataset, we can use the `monthly_mean_climatology` method as follows:

```
[24]: sst = nc.open_data("sst.mon.mean.nc")
      sst.monthly_mean_climatology()
      sst.select_months(1)
      sst.plot()
```



```
[24]: :DynamicMap    [time]
      :Image      [lon,lat]    (sst)
```

2.9.18 How to calculate a seasonal climatology

```
[25]: sst = nc.open_data("sst.mon.mean.nc")
      sst.seasonal_mean_climatology()
      sst.select_timesteps(0)
      sst.plot()
```

```
[25]: :DynamicMap    [time]
      :Image      [lon,lat]    (sst)
```

```
[26]: ## How to read a dataset using pandas or xarray
```

To read the dataset to an xarray Dataset use `to_xarray`:

```
[27]: sst = nc.open_data("sst.mon.mean.nc")
      sst.to_xarray()
```

```
[27]: <xarray.Dataset>
Dimensions:  (lat: 180, lon: 360, time: 2028)
Coordinates:
  * lat      (lat) float32 89.5 88.5 87.5 86.5 85.5 ... -86.5 -87.5 -88.5 -89.5
  * lon      (lon) float32 0.5 1.5 2.5 3.5 4.5 ... 355.5 356.5 357.5 358.5 359.5
  * time      (time) datetime64[ns] 1850-01-01 1850-02-01 ... 2018-12-01
Data variables:
  sst        (time, lat, lon) float32 ...
Attributes:
  title:      created 12/2013 from data provided by JRA
  history:    Created 12/2012 from data obtained from JRA by ESRL/PSD
  platform:   Analyses
  citation:   Hirahara, S., Ishii, M., and Y. Fukuda,2014: Centennial...
  institution: NOAA ESRL/PSD
  Conventions: CF-1.2
  References: http://www.esrl.noaa.gov/psd/data/gridded/cobe2.html
  dataset_title: COBE-SST2 Sea Surface Temperature and Ice
  original_source: https://climate.mri-jma.go.jp/pub/ocean/cobe-sst2/
```

To read the dataset in as a pandas dataframe use `to_dataframe`:

```
[28]: sst.to_dataframe()
```

```
[28]:
```

	lat	lon	time	sst
	89.5	0.5	1850-01-01	-1.712
			1850-02-01	-1.698
			1850-03-01	-1.707
			1850-04-01	-1.742
			1850-05-01	-1.725
...				...
	-89.5	359.5	2018-08-01	NaN
			2018-09-01	NaN
			2018-10-01	NaN
			2018-11-01	NaN
			2018-12-01	NaN

(continues on next page)

```
[131414400 rows x 1 columns]
```

2.9.19 How to calculate cell areas

If we want to calculate the area of each cell in a dataset, we use the `cell_area` method. The `join` argument let's you choose whether to join the cell areas to the existing dataset, or to only include cell areas in the dataset.

```
[29]: sst = nc.open_data("sst.mon.mean.nc")
      sst.cell_areas(join=False)
      sst.plot()
```

```
[29]: :Image      [lon,lat]      (cell_area)
```

2.9.20 How to use urls

If a file is located at a url, we can send it to `open_data`:

```
[30]: url = "ftp://ftp.cdc.noaa.gov/Datasets/COBE2/sst.mon.ltm.1981-2010.nc"
      sst = nc.open_data(url)
```

```
Downloading ftp://ftp.cdc.noaa.gov/Datasets/COBE2/sst.mon.ltm.1981-2010.nc
```

This will download the file from the url and save it as a temp file. We can then work with it as usual. A future release of nctoolkit will have thredds support.

2.9.21 How to calculate an ensemble average

nctoolkit has built in methods for working with ensembles. Let's start by splitting the 1850-2019 sst dataset into an ensemble, where each file is a separate year:

```
[31]: sst = nc.open_data("sst.mon.mean.nc")
      sst.split("year")
```

An ensemble mean can be calculated in two ways. First, we can calculate the mean in each time step. So here the files have temperature from 1850 onwards. We can calculate the monthly mean temperature over that time period as follows, and from there we can calculate the global mean:

```
[32]: sst.ensemble_mean()
      sst.spatial_mean()
      sst.plot()
```

```
[32]: :Curve      [time]      (x)
```

We might want to calculate the average over all time steps, i.e. calculating mean temperature since 1850. We do this by changing the `ignore_time` argument:

```
[33]: sst = nc.open_data("sst.mon.mean.nc")
      sst.split("year")
      sst.ensemble_mean(ignore_time=True)
      sst.plot()
```

```
[33]: :DynamicMap    [time]
      :Image      [lon,lat]    (sst)
```

2.10 API Reference

2.10.1 Session options

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

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.

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

2.10.2 Reading/copying data

<code>open_data(x, suppress_messages, 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

nctoolkit.open_data

`nctoolkit.open_data(x=None, suppress_messages=False, 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 url is given the file will be downloaded before processing. If a `*.nc` style wildcard is supplied, `open_data` will use all files available.
- **thredds** (*boolean*) – Are you accessing a thredds server? Must end with `.nc`.
- **checks** (*boolean*) – Do you want basic checks to ensure cdo can read files?

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.

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.

nctoolkit.DataSet.copy

`DataSet.copy` (*self*)

Make a deep copy of an DataSet object

2.10.3 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.

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"]`.

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

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

2.10.4 Adding file(s) to a dataset

append

nctoolkit.append

Functions

`append(self[, x])`

Add new file(s) to a dataset

2.10.5 Accessing attributes

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

nctoolkit.DataSet.variables

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

nctoolkit.DataSet.years

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

nctoolkit.DataSet.months

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

nctoolkit.DataSet.times

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

nctoolkit.DataSet.levels

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

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.

nctoolkit.DataSet.current**property** `DataSet.current`

The current file or files in the DataSet object

nctoolkit.DataSet.history**property** `DataSet.history`

The history of operations on the DataSet

nctoolkit.DataSet.start**property** `DataSet.start`

The starting file or files of the DataSet object

2.10.6 Plotting

<code>DataSet.plot(self[, vars])</code>	
<code>DataSet.view(self)</code>	Open the current dataset's file in ncview

nctoolkit.DataSet.plot

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

nctoolkit.DataSet.view

`DataSet.view (self)`

Open the current dataset's file in ncview

2.10.7 Variable modification

<code>DataSet.mutate(self[, operations])</code>	Create new variables using mathematical expressions, and keep original variables
<code>DataSet.transmute(self[, operations])</code>	Create new variables using mathematical expressions, and drop original variables
<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

nctoolkit.DataSet.mutate

`DataSet.mutate(self, operations=None)`

Create new variables using mathematical expressions, and keep original variables

Parameters `operations` (*dict*) – operations to apply. The keys are the new variables to generate. The values are the mathematical operations to carry out. Both must be strings.

nctoolkit.DataSet.transmute

`DataSet.transmute(self, operations=None)`

Create new variables using mathematical expressions, and drop original variables

Parameters `operations` (*dict*) – operations to apply. The keys are the new variables to generate. The values are the mathematical operations to carry out. Both must be strings.

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

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.

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?

2.10.8 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

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

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.

2.10.9 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)</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_cum_sum(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.

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.

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.

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.

nctoolkit.DataSet.vertical_mean

DataSet.**vertical_mean** (*self*)

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

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

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

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

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

nctoolkit.DataSet.vertical_cum_sum

`DataSet.vertical_cum_sum(self)`

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

nctoolkit.DataSet.invert_levels

`DataSet.invert_levels(self)`

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

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

2.10.10 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

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

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

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

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

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

2.10.11 Evaluation setting

<code>DataSet.run(self)</code>	Run all stored commands in a dataset
--------------------------------	--------------------------------------

nctoolkit.DataSet.run

`DataSet.run(self)`

Run all stored commands in a dataset

2.10.12 Cleaning functions

2.10.13 Ensemble creation

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

nctoolkit.create_ensemble

`nctoolkit.create_ensemble(path="", var=None, recursive=True)`

Generate an ensemble

Parameters

- **path** (*str*) – The system 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

2.10.14 Arithmetic methods

<code>DataSet.mutate(self[, operations])</code>	Create new variables using mathematical expressions, and keep original variables
<code>DataSet.transmute(self[, operations])</code>	Create new variables using mathematical expressions, and drop original variables
<code>DataSet.add(self[, x, var])</code>	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.
<code>DataSet.subtract(self[, x, var])</code>	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.
<code>DataSet.multiply(self[, x, var])</code>	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.
<code>DataSet.divide(self[, x, var])</code>	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.

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

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

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

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

2.10.15 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.

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.

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.

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.

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. $0 \leq p \leq 100$.

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.

2.10.16 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.select_variables(self[, vars])</code>	Select variables from a dataset
<code>DataSet.remove_variables(self[, vars])</code>	Remove variables This will remove stated variables from files in the dataset.
<code>DataSet.select_years(self[, years])</code>	Select years from a dataset This method will subset the dataset to only contain years within the list given.
<code>DataSet.select_months(self[, months])</code>	Select months from a dataset This method will subset the dataset to only contain months within the list given.
<code>DataSet.select_seasons(self[, season])</code>	Select season from a dataset
<code>DataSet.select_timesteps(self[, times])</code>	Select timesteps from a dataset

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 clipping? 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

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.

lon [list] A list [lon_min, lon_max] giving the maximum and minimum longitude required.

lat [list] A list [lat_min, lat_max] giving the maximum and minimum latitude required.

A usage example is the following:

```
data.select(variables = “var”)
```

nctoolkit.DataSet.select_variables

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

Select variables from a dataset

Parameters **vars** (*list or str*) – Variable(s) to select.

nctoolkit.DataSet.remove_variables

`DataSet.remove_variables (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

nctoolkit.DataSet.select_years

`DataSet.select_years (self, years=None)`

Select years from a dataset This method will subset the dataset to only contain years within the list given. A warning message will be provided when there are missing years. :param years: Years(s) to select. These should be integers :type years: list, range or int

nctoolkit.DataSet.select_months

`DataSet.select_months(self, months=None)`

Select months from a dataset This method will subset the dataset to only contain months within the list given. A warning message will be provided when there are missing months.

Parameters `months` (*list, range or int*) – Month(s) to select.

nctoolkit.DataSet.select_seasons

`DataSet.select_seasons(self, season=None)`

Select season from a dataset

Parameters `season` (*str*) – Season to select. One of “DJF”, “MAM”, “JJA”, “SON”.

nctoolkit.DataSet.select_timesteps

`DataSet.select_timesteps(self, times=None)`

Select timesteps from a dataset

Parameters `times` (*list or int*) – time step(s) to select. For example, if you wanted the first time step set times=0.

2.10.17 Time-based methods

<code>DataSet.set_date(self[, year, month, day, ...])</code>	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.
<code>DataSet.shift_hours(self[, shift])</code>	Shift times in dataset by a number of hours
<code>DataSet.shift_days(self[, shift])</code>	Shift times in dataset by a number of days
<code>DataSet.shift_months(self[, shift])</code>	Shift times in dataset by a number of months
<code>DataSet.shift_years(self[, shift])</code>	Shift times in dataset by a number of years
<code>DataSet.shift(self, **kwargs)</code>	Shift method.

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.

nctoolkit.DataSet.shift_hours

`DataSet.shift_hours(self, shift=None)`

Shift times in dataset by a number of hours

Parameters `shift` (*int*) – Number of hours, positive or negative, to shift the time by.

nctoolkit.DataSet.shift_days

`DataSet.shift_days(self, shift=None)`

Shift times in dataset by a number of days

Parameters `shift` (*int*) – Number of days, positive or negative, to shift the time by.

nctoolkit.DataSet.shift_months

`DataSet.shift_months(self, shift=None)`

Shift times in dataset by a number of months

Parameters `shift` (*int*) – Number of days, positive or negative, to shift the time by.

nctoolkit.DataSet.shift_years

`DataSet.shift_years(self, shift=None)`

Shift times in dataset by a number of years

Parameters `shift` (*int*) – Number of days, positive or negative, to shift the time by.

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.

2.10.18 Interpolation and resampling methods

<code>DataSet.regrid(self[, grid, method])</code>	Regrid a dataset to a target grid
<code>DataSet.to_latlon(self[, lon, lat, res, method])</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

nctoolkit.DataSet.regrid

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

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”; “bic” - “bicubic interpolation”.

nctoolkit.DataSet.to_latlon

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

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”. Bilinear: “bil”; Nearest neighbour: “nn”.

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

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.

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`.

2.10.19 Masking methods

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

Mask a lon/lat box

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]

2.10.20 Summary 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.

`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.

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.

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.

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.

2.10.21 Statistical methods

<code>DataSet.mean(self)</code>	Calculate the temporal mean of all variables
<code>DataSet.min(self)</code>	Calculate the temporal minimum of all variables
<code>DataSet.median(self)</code>	Calculate the temporal median of all variables
<code>DataSet.percentile(self[, p])</code>	Calculate the temporal percentile of all variables
<code>DataSet.max(self)</code>	Calculate the temporal maximum of all variables
<code>DataSet.sum(self)</code>	Calculate the temporal sum of all variables
<code>DataSet.range(self)</code>	Calculate the temporal range of all variables
<code>DataSet.var(self)</code>	Calculate the temporal variance of all variables
<code>DataSet.cum_sum(self)</code>	Calculate the temporal cumulative sum of all variables
<code>DataSet.cor_space(self[, var1, var2])</code>	Calculate the correlation correct between two variables in space This is calculated for each time step.
<code>DataSet.cor_time(self[, var1, var2])</code>	Calculate the correlation correct in time between two variables The correlation is calculated for each grid cell, ignoring missing values.
<code>DataSet.spatial_mean(self)</code>	Calculate the area weighted spatial mean for all variables This is performed for each time step.

continues on next page

Table 22 – continued from previous page

<code>DataSet.spatial_min(self)</code>	Calculate the spatial minimum for all variables This is performed for each time step.
<code>DataSet.spatial_max(self)</code>	Calculate the spatial maximum for all variables This is performed for each time step.
<code>DataSet.spatial_percentile(self[, p])</code>	Calculate the spatial sum for all variables This is performed for each time step.
<code>DataSet.spatial_range(self)</code>	Calculate the spatial range for all variables This is performed for each time step.
<code>DataSet.spatial_sum(self[, by_area])</code>	Calculate the spatial sum for all variables This is performed for each time step.
<code>DataSet.centre(self[, by, by_area])</code>	Calculate the latitudinal or longitudinal centre for each year/month combination in files. This applies to each file in an ensemble. <code>by</code> : str Set to 'latitude' if you want the latitudinal centre calculated. 'longitude' for longitudinal. <code>by_area</code> : bool If the variable is a value/m2 type variable, set to True, otherwise set to False.
<code>DataSet.monthly_mean(self)</code>	Calculate the monthly mean for each year/month combination in files.
<code>DataSet.monthly_min(self)</code>	Calculate the monthly minimum for each year/month combination in files.
<code>DataSet.monthly_max(self)</code>	Calculate the monthly maximum for each year/month combination in files.
<code>DataSet.monthly_range(self)</code>	Calculate the monthly range for each year/month combination in files.
<code>DataSet.daily_mean(self)</code>	Calculate the daily mean for each variable
<code>DataSet.daily_min(self)</code>	Calculate the daily minimum for each variable
<code>DataSet.daily_max(self)</code>	Calculate the daily maximum for each variable
<code>DataSet.daily_mean(self)</code>	Calculate the daily mean for each variable
<code>DataSet.daily_range(self)</code>	Calculate the daily range for each variable
<code>DataSet.daily_sum(self)</code>	Calculate the daily sum for each variable
<code>DataSet.daily_mean_climatology(self)</code>	Calculate a daily mean climatology
<code>DataSet.daily_min_climatology(self)</code>	Calculate a daily minimum climatology
<code>DataSet.daily_max_climatology(self)</code>	Calculate a daily maximum climatology
<code>DataSet.daily_mean_climatology(self)</code>	Calculate a daily mean climatology
<code>DataSet.daily_range_climatology(self)</code>	Calculate a daily range climatology
<code>DataSet.monthly_mean_climatology(self)</code>	Calculate the monthly mean climatologies Defined as the minimum value in each month across all years.
<code>DataSet.monthly_min_climatology(self)</code>	Calculate the monthly minimum climatologies Defined as the minimum value in each month across all years.
<code>DataSet.monthly_max_climatology(self)</code>	Calculate the monthly maximum climatologies Defined as the maximum value in each month across all years.
<code>DataSet.monthly_range_climatology(self)</code>	Calculate the monthly range climatologies Defined as the range of value in each month across all years.
<code>DataSet.annual_mean(self)</code>	Calculate the annual mean for each variable
<code>DataSet.annual_min(self)</code>	Calculate the annual minimum for each variable
<code>DataSet.annual_max(self)</code>	Calculate the annual maximum for each variable
<code>DataSet.annual_sum(self)</code>	Calculate the annual sum for each variable
<code>DataSet.annual_range(self)</code>	Calculate the annual range for each variable
<code>DataSet.seasonal_mean(self)</code>	Calculate the seasonal mean for each year.
<code>DataSet.seasonal_min(self)</code>	Calculate the seasonal minimum for each year.
<code>DataSet.seasonal_max(self)</code>	Calculate the seasonal maximum for each year.

continues on next page

Table 22 – continued from previous page

<i>DataSet.seasonal_range</i> (self)	Calculate the seasonal range for each year.
<i>DataSet.seasonal_mean_climatology</i> (self)	Calculate a climatological seasonal mean
<i>DataSet.seasonal_min_climatology</i> (self)	Calculate a climatological seasonal min This is defined as the minimum value in each season across all years.
<i>DataSet.seasonal_max_climatology</i> (self)	Calculate a climatological seasonal max This is defined as the maximum value in each season across all years.
<i>DataSet.seasonal_range_climatology</i> (self)	Calculate a climatological seasonal range This is defined as the range of values in each season across all years.
<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.
<i>DataSet.zonal_max</i> (self)	Calculate the zonal maximum for each year/month combination in files.
<i>DataSet.zonal_range</i> (self)	Calculate the zonal range for each year/month combination in files.
<i>DataSet.meridional_mean</i> (self)	Calculate the meridional mean for each year/month combination in files.
<i>DataSet.meridional_min</i> (self)	Calculate the meridional minimum for each year/month combination in files.
<i>DataSet.meridional_max</i> (self)	Calculate the meridional maximum for each year/month combination in files.
<i>DataSet.meridional_range</i> (self)	Calculate the meridional range for each year/month combination in files.

nctoolkit.DataSet.mean`DataSet.mean(self)`

Calculate the temporal mean of all variables

nctoolkit.DataSet.min`DataSet.min(self)`

Calculate the temporal minimum of all variables

nctoolkit.DataSet.median

`DataSet.median(self)`

Calculate the temporal median of all variables

nctoolkit.DataSet.percentile

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

Calculate the temporal percentile of all variables

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

nctoolkit.DataSet.max

`DataSet.max(self)`

Calculate the temporal maximum of all variables

nctoolkit.DataSet.sum

`DataSet.sum(self)`

Calculate the temporal sum of all variables

nctoolkit.DataSet.range

`DataSet.range(self)`

Calculate the temporal range of all variables

nctoolkit.DataSet.var

`DataSet.var(self)`

Calculate the temporal variance of all variables

nctoolkit.DataSet.cum_sum

`DataSet.cum_sum(self)`

Calculate the temporal cumulative sum of all variables

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

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

nctoolkit.DataSet.spatial_mean

`DataSet.spatial_mean(self)`

Calculate the area weighted spatial mean for all variables This is performed for each time step.

nctoolkit.DataSet.spatial_min

`DataSet.spatial_min(self)`

Calculate the spatial minimum for all variables This is performed for each time step.

nctoolkit.DataSet.spatial_max

`DataSet.spatial_max(self)`

Calculate the spatial maximum for all variables This is performed for each time step.

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

nctoolkit.DataSet.spatial_range

`DataSet.spatial_range(self)`

Calculate the spatial range for all variables This is performed for each time step.

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.

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.

nctoolkit.DataSet.monthly_mean

`DataSet.monthly_mean(self)`

Calculate the monthly mean for each year/month combination in files. This applies to each file in an ensemble.

nctoolkit.DataSet.monthly_min

`DataSet.monthly_min(self)`

Calculate the monthly minimum for each year/month combination in files. This applies to each file in an ensemble.

nctoolkit.DataSet.monthly_max

`DataSet.monthly_max(self)`

Calculate the monthly maximum for each year/month combination in files. This applies to each file in an ensemble.

nctoolkit.DataSet.monthly_range

`DataSet.monthly_range(self)`

Calculate the monthly range for each year/month combination in files. This applies to each file in an ensemble.

nctoolkit.DataSet.daily_mean

`DataSet.daily_mean(self)`

Calculate the daily mean for each variable

nctoolkit.DataSet.daily_min

`DataSet.daily_min(self)`

Calculate the daily minimum for each variable

nctoolkit.DataSet.daily_max

`DataSet.daily_max(self)`
Calculate the daily maximum for each variable

nctoolkit.DataSet.daily_range

`DataSet.daily_range(self)`
Calculate the daily range for each variable

nctoolkit.DataSet.daily_sum

`DataSet.daily_sum(self)`
Calculate the daily sum for each variable

nctoolkit.DataSet.daily_mean_climatology

`DataSet.daily_mean_climatology(self)`
Calculate a daily mean climatology

nctoolkit.DataSet.daily_min_climatology

`DataSet.daily_min_climatology(self)`
Calculate a daily minimum climatology

nctoolkit.DataSet.daily_max_climatology

`DataSet.daily_max_climatology(self)`
Calculate a daily maximum climatology

nctoolkit.DataSet.daily_range_climatology

`DataSet.daily_range_climatology(self)`
Calculate a daily range climatology

nctoolkit.DataSet.monthly_mean_climatology

`DataSet.monthly_mean_climatology(self)`
Calculate the monthly mean climatologies Defined as the minimum value in each month across all years. This applies to each file in an ensemble.

nctoolkit.DataSet.monthly_min_climatology

`DataSet.monthly_min_climatology(self)`

Calculate the monthly minimum climatologies Defined as the minimum value in each month across all years.
This applies to each file in an ensemble.

nctoolkit.DataSet.monthly_max_climatology

`DataSet.monthly_max_climatology(self)`

Calculate the monthly maximum climatologies Defined as the maximum value in each month across all years.
This applies to each file in an ensemble.

nctoolkit.DataSet.monthly_range_climatology

`DataSet.monthly_range_climatology(self)`

Calculate the monthly range climatologies Defined as the range of value in each month across all years. This applies to each file in an ensemble.

nctoolkit.DataSet.annual_mean

`DataSet.annual_mean(self)`

Calculate the annual mean for each variable

nctoolkit.DataSet.annual_min

`DataSet.annual_min(self)`

Calculate the annual minimum for each variable

nctoolkit.DataSet.annual_max

`DataSet.annual_max(self)`

Calculate the annual maximum for each variable

nctoolkit.DataSet.annual_sum

`DataSet.annual_sum(self)`

Calculate the annual sum for each variable

nctoolkit.DataSet.annual_range

`DataSet.annual_range(self)`

Calculate the annual range for each variable

nctoolkit.DataSet.seasonal_mean

`DataSet.seasonal_mean(self)`

Calculate the seasonal mean for each year. Applies at the grid cell level.

nctoolkit.DataSet.seasonal_min

`DataSet.seasonal_min(self)`

Calculate the seasonal minimum for each year. Applies at the grid cell level.

nctoolkit.DataSet.seasonal_max

`DataSet.seasonal_max(self)`

Calculate the seasonal maximum for each year. Applies at the grid cell level.

nctoolkit.DataSet.seasonal_range

`DataSet.seasonal_range(self)`

Calculate the seasonal range for each year. Applies at the grid cell level.

nctoolkit.DataSet.seasonal_mean_climatology

`DataSet.seasonal_mean_climatology(self)`

Calculate a climatological seasonal mean

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

nctoolkit.DataSet.seasonal_min_climatology

`DataSet.seasonal_min_climatology(self)`

Calculate a climatological seasonal min This is defined as the minimum value in each season across all years.

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

nctoolkit.DataSet.seasonal_max_climatology

`DataSet.seasonal_max_climatology(self)`

Calculate a climatological seasonal max This is defined as the maximum value in each season across all years.

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

nctoolkit.DataSet.seasonal_range_climatology

`DataSet.seasonal_range_climatology(self)`

Calculate a climatological seasonal range This is defined as the range of values in each season across all years.

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

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.

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.

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.

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.

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.

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.

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.

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.

2.10.22 Seasonal methods

<code>DataSet.seasonal_mean(self)</code>	Calculate the seasonal mean for each year.
<code>DataSet.seasonal_min(self)</code>	Calculate the seasonal minimum for each year.
<code>DataSet.seasonal_max(self)</code>	Calculate the seasonal maximum for each year.
<code>DataSet.seasonal_range(self)</code>	Calculate the seasonal range for each year.
<code>DataSet.seasonal_mean_climatology(self)</code>	Calculate a climatological seasonal mean
<code>DataSet.seasonal_min_climatology(self)</code>	Calculate a climatological seasonal min This is defined as the minimum value in each season across all years.
<code>DataSet.seasonal_max_climatology(self)</code>	Calculate a climatological seasonal max This is defined as the maximum value in each season across all years.
<code>DataSet.seasonal_range_climatology(self)</code>	Calculate a climatological seasonal range This is defined as the range of values in each season across all years.
<code>DataSet.select_seasons(self[, season])</code>	Select season from a dataset

2.10.23 Merging methods

<code>DataSet.merge(self[, match])</code>	Merge a multi-file ensemble into a single file Merging will occur based on the time steps in the first file.
<code>DataSet.merge_time(self)</code>	Time-based merging of a multi-file ensemble into a single file This method is ideal if you have the same data split over multiple files covering different data sets.

nctoolkit.DataSet.merge

`DataSet.merge(self, match=['year', 'month', 'day'])`

Merge a multi-file ensemble into a single file Merging will occur based on the time steps in the first file. This will only be effective if you want to merge files with the same times, but with different variables.

Parameters `match` (*list*, *str*) – 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.

nctoolkit.DataSet.merge_time

`DataSet.merge_time(self)`

Time-based merging of a multi-file ensemble into a single file This method is ideal if you have the same data split over multiple files covering different data sets.

2.10.24 Climatology methods

<code>DataSet.daily_mean_climatology(self)</code>	Calculate a daily mean climatology
<code>DataSet.daily_min_climatology(self)</code>	Calculate a daily minimum climatology
<code>DataSet.daily_max_climatology(self)</code>	Calculate a daily maximum climatology
<code>DataSet.daily_mean_climatology(self)</code>	Calculate a daily mean climatology
<code>DataSet.daily_range_climatology(self)</code>	Calculate a daily range climatology
<code>DataSet.monthly_mean_climatology(self)</code>	Calculate the monthly mean climatologies Defined as the minimum value in each month across all years.
<code>DataSet.monthly_min_climatology(self)</code>	Calculate the monthly minimum climatologies Defined as the minimum value in each month across all years.
<code>DataSet.monthly_max_climatology(self)</code>	Calculate the monthly maximum climatologies Defined as the maximum value in each month across all years.
<code>DataSet.monthly_range_climatology(self)</code>	Calculate the monthly range climatologies Defined as the range of value in each month across all years.

2.10.25 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.
--	---

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.

2.10.26 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.

continues on next page

Table 27 – continued from previous page

<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/nc NetCDF4 (HDF5) = nc4 NetCDF4-classic = nc4c NetCDF version 5 (64-bit data) = nc5 :type ext: str.
--	---

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.

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

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

nctoolkit.DataSet.zip

`DataSet.zip(self)`

Zip the dataset This will compress the files within the dataset. This works lazily.

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-classic
= nc4c NetCDF version 5 (64-bit data) = nc5

2.10.27 Miscellaneous methods

<code>DataSet.cell_areas(self[, join])</code>	Calculate the area of grid cells.
<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_all(self[, expression])</code>	Compare all variables to a constant
<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.

nctoolkit.DataSet.cell_areas

`DataSet.cell_areas(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.

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.

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.

nctoolkit.DataSet.compare_all

`DataSet.compare_all(self, expression=None)`

Compare all variables to a constant

Parameters `expression` (*str*) – This a regular comparison such as “<0”, “>0”, “==0”

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.

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.

2.11 Package info

This package was created by Robert Wilson at Plymouth Marine Laboratory (PML).

2.11.1 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.

2.11.2 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`, [49](#)

A

add() (*nctoolkit.DataSet* method), 58
 annual_anomaly() (*nctoolkit.DataSet* method), 65
 annual_max() (*nctoolkit.DataSet* method), 73
 annual_mean() (*nctoolkit.DataSet* method), 73
 annual_min() (*nctoolkit.DataSet* method), 73
 annual_range() (*nctoolkit.DataSet* method), 73
 annual_sum() (*nctoolkit.DataSet* method), 73

B

bottom() (*nctoolkit.DataSet* method), 54
 bottom_mask() (*nctoolkit.DataSet* method), 55

C

cdo_command() (*nctoolkit.DataSet* method), 79
 cell_areas() (*nctoolkit.DataSet* method), 79
 centre() (*nctoolkit.DataSet* method), 71
 compare_all() (*nctoolkit.DataSet* method), 80
 copy() (*nctoolkit.DataSet* method), 48
 cor_space() (in module *nctoolkit*), 49
 cor_space() (*nctoolkit.DataSet* method), 69
 cor_time() (in module *nctoolkit*), 49
 cor_time() (*nctoolkit.DataSet* method), 70
 create_ensemble() (in module *nctoolkit*), 57
 crop() (*nctoolkit.DataSet* method), 60
 cum_sum() (*nctoolkit.DataSet* method), 69
 current() (*nctoolkit.DataSet* property), 51

D

daily_max() (*nctoolkit.DataSet* method), 72
 daily_max_climatology() (*nctoolkit.DataSet* method), 72
 daily_mean() (*nctoolkit.DataSet* method), 71
 daily_mean_climatology() (*nctoolkit.DataSet* method), 72
 daily_min() (*nctoolkit.DataSet* method), 71
 daily_min_climatology() (*nctoolkit.DataSet* method), 72
 daily_range() (*nctoolkit.DataSet* method), 72
 daily_range_climatology() (*nctoolkit.DataSet* method), 72
 daily_sum() (*nctoolkit.DataSet* method), 72

divide() (*nctoolkit.DataSet* method), 58

E

ensemble_max() (*nctoolkit.DataSet* method), 59
 ensemble_mean() (*nctoolkit.DataSet* method), 59
 ensemble_min() (*nctoolkit.DataSet* method), 59
 ensemble_percentile() (*nctoolkit.DataSet* method), 60
 ensemble_range() (*nctoolkit.DataSet* method), 60

F

format() (*nctoolkit.DataSet* method), 79

H

history() (*nctoolkit.DataSet* property), 51

I

invert_levels() (*nctoolkit.DataSet* method), 55

L

levels() (*nctoolkit.DataSet* property), 50

M

mask_box() (*nctoolkit.DataSet* method), 65
 max() (*nctoolkit.DataSet* method), 69
 mean() (*nctoolkit.DataSet* method), 68
 median() (*nctoolkit.DataSet* method), 69
 merge() (in module *nctoolkit*), 49
 merge() (*nctoolkit.DataSet* method), 76
 merge_time() (*nctoolkit.DataSet* method), 77
 meridional_max() (*nctoolkit.DataSet* method), 75
 meridional_mean() (*nctoolkit.DataSet* method), 75
 meridional_min() (*nctoolkit.DataSet* method), 75
 meridional_range() (*nctoolkit.DataSet* method), 76
 min() (*nctoolkit.DataSet* method), 68
 module
 nctoolkit.append, 49
 monthly_anomaly() (*nctoolkit.DataSet* method), 66
 monthly_max() (*nctoolkit.DataSet* method), 71
 monthly_max_climatology() (*nctoolkit.DataSet* method), 73

`monthly_mean()` (*nctoolkit.DataSet method*), 71
`monthly_mean_climatology()` (*nctoolkit.DataSet method*), 72
`monthly_min()` (*nctoolkit.DataSet method*), 71
`monthly_min_climatology()` (*nctoolkit.DataSet method*), 73
`monthly_range()` (*nctoolkit.DataSet method*), 71
`monthly_range_climatology()` (*nctoolkit.DataSet method*), 73
`months()` (*nctoolkit.DataSet property*), 50
`multiply()` (*nctoolkit.DataSet method*), 58
`mutate()` (*nctoolkit.DataSet method*), 52

N

`nco_command()` (*nctoolkit.DataSet method*), 79
`nctoolkit.append`
 module, 49

O

`open_data()` (*in module nctoolkit*), 47
`open_thredds()` (*in module nctoolkit*), 48
`open_url()` (*in module nctoolkit*), 48
`options()` (*in module nctoolkit*), 47

P

`percentile()` (*nctoolkit.DataSet method*), 69
`phenology()` (*nctoolkit.DataSet method*), 66
`plot()` (*nctoolkit.DataSet method*), 51

R

`range()` (*nctoolkit.DataSet method*), 69
`reduce_dims()` (*nctoolkit.DataSet method*), 80
`reduce_grid()` (*nctoolkit.DataSet method*), 80
`regrid()` (*nctoolkit.DataSet method*), 64
`remove_variables()` (*nctoolkit.DataSet method*), 61
`rename()` (*nctoolkit.DataSet method*), 52
`resample_grid()` (*nctoolkit.DataSet method*), 64
`rolling_max()` (*nctoolkit.DataSet method*), 56
`rolling_mean()` (*nctoolkit.DataSet method*), 55
`rolling_min()` (*nctoolkit.DataSet method*), 55
`rolling_range()` (*nctoolkit.DataSet method*), 56
`rolling_sum()` (*nctoolkit.DataSet method*), 56
`run()` (*nctoolkit.DataSet method*), 56

S

`seasonal_max()` (*nctoolkit.DataSet method*), 74
`seasonal_max_climatology()` (*nctoolkit.DataSet method*), 74
`seasonal_mean()` (*nctoolkit.DataSet method*), 74
`seasonal_mean_climatology()` (*nctoolkit.DataSet method*), 74
`seasonal_min()` (*nctoolkit.DataSet method*), 74

`seasonal_min_climatology()` (*nctoolkit.DataSet method*), 74
`seasonal_range()` (*nctoolkit.DataSet method*), 74
`seasonal_range_climatology()` (*nctoolkit.DataSet method*), 75
`select()` (*nctoolkit.DataSet method*), 61
`select_months()` (*nctoolkit.DataSet method*), 62
`select_seasons()` (*nctoolkit.DataSet method*), 62
`select_timesteps()` (*nctoolkit.DataSet method*), 62
`select_variables()` (*nctoolkit.DataSet method*), 61
`select_years()` (*nctoolkit.DataSet method*), 61
`set_date()` (*nctoolkit.DataSet method*), 62
`set_longnames()` (*nctoolkit.DataSet method*), 53
`set_missing()` (*nctoolkit.DataSet method*), 52
`set_units()` (*nctoolkit.DataSet method*), 53
`shift()` (*nctoolkit.DataSet method*), 63
`shift_days()` (*nctoolkit.DataSet method*), 63
`shift_hours()` (*nctoolkit.DataSet method*), 63
`shift_months()` (*nctoolkit.DataSet method*), 63
`shift_years()` (*nctoolkit.DataSet method*), 63
`size()` (*nctoolkit.DataSet property*), 51
`spatial_max()` (*nctoolkit.DataSet method*), 70
`spatial_mean()` (*nctoolkit.DataSet method*), 70
`spatial_min()` (*nctoolkit.DataSet method*), 70
`spatial_percentile()` (*nctoolkit.DataSet method*), 70
`spatial_range()` (*nctoolkit.DataSet method*), 70
`spatial_sum()` (*nctoolkit.DataSet method*), 70
`split()` (*nctoolkit.DataSet method*), 77
`start()` (*nctoolkit.DataSet property*), 51
`subtract()` (*nctoolkit.DataSet method*), 58
`sum()` (*nctoolkit.DataSet method*), 69
`sum_all()` (*nctoolkit.DataSet method*), 52
`surface()` (*nctoolkit.DataSet method*), 54

T

`time_interp()` (*nctoolkit.DataSet method*), 64
`times()` (*nctoolkit.DataSet property*), 50
`timestep_interp()` (*nctoolkit.DataSet method*), 65
`to_dataframe()` (*nctoolkit.DataSet method*), 78
`to_latlon()` (*nctoolkit.DataSet method*), 64
`to_nc()` (*nctoolkit.DataSet method*), 78
`to_xarray()` (*nctoolkit.DataSet method*), 78
`transmute()` (*nctoolkit.DataSet method*), 52

V

`var()` (*nctoolkit.DataSet method*), 69
`variables()` (*nctoolkit.DataSet property*), 50
`vertical_cum_sum()` (*nctoolkit.DataSet method*), 55
`vertical_interp()` (*nctoolkit.DataSet method*), 54
`vertical_max()` (*nctoolkit.DataSet method*), 54

`vertical_mean()` (*nctoolkit.DataSet method*), 54
`vertical_min()` (*nctoolkit.DataSet method*), 54
`vertical_range()` (*nctoolkit.DataSet method*), 54
`vertical_sum()` (*nctoolkit.DataSet method*), 55
`view()` (*nctoolkit.DataSet method*), 51

Y

`years()` (*nctoolkit.DataSet property*), 50

Z

`zip()` (*nctoolkit.DataSet method*), 78
`zonal_max()` (*nctoolkit.DataSet method*), 75
`zonal_mean()` (*nctoolkit.DataSet method*), 75
`zonal_min()` (*nctoolkit.DataSet method*), 75
`zonal_range()` (*nctoolkit.DataSet method*), 75