
xhistogram Documentation

Release .0.3.2+0.g002d6bd.dirty

xhistogram developers

Sep 20, 2022

CONTENTS

1	Why a new histogram package?	3
2	Contents	5
	Python Module Index	23
	Index	25

Histograms (a.k.a “binning”) are much more than just a visualization tool. They are the foundation of a wide range of scientific analyses including [joint] probability distributions and coordinate transformations. Xhistogram makes it easier to calculate flexible, complex histograms with multi-dimensional data. It integrates (optionally) with Dask, in order to scale up to very large datasets and with Xarray, in order to consume and produce labelled, annotated data structures. It is useful for a wide range of scientific tasks.

WHY A NEW HISTOGRAM PACKAGE?

The main problem with the standard `histogram` function in `numpy` and `dask` is that they automatically act over the entire input array (i.e. they “flatten” the data). `Xhistogram` allows you to choose which axes / dimensions you want to preserve and which you want to flatten. It also allows you to combine `N` arbitrary inputs to produce `N`-dimensional histograms. A good place to start is the [Xhistogram Tutorial](#).

CONTENTS

2.1 Installation

2.1.1 Requirements

xhistogram is compatible with python 3. It requires numpy and, optionally, xarray.

2.1.2 Installation from Conda Forge

The easiest way to install xhistogram along with its dependencies is via conda forge:

```
conda install -c conda-forge xhistogram
```

2.1.3 Installation from Pip

An alternative is to use pip:

```
pip install xhistogram
```

This will install the latest release from [pypi](#).

2.1.4 Installation from GitHub

xhistogram is under active development. To obtain the latest development version, you may clone the [source repository](#) and install it:

```
git clone https://github.com/xgcm/xhistogram.git
cd xhistogram
python setup.py install
```

or simply:

```
pip install git+https://github.com/xgcm/xhistogram.git
```

Users are encouraged to [fork](#) xhistogram and submit [issues](#) and [pull requests](#).

2.2 Xhistogram Tutorial

Histograms are the foundation of many forms of data analysis. The goal of xhistogram is to make it easy to calculate weighted histograms in multiple dimensions over n-dimensional arrays, with control over the axes. Xhistogram builds on top of xarray, for automatic coordinates and labels, and dask, for parallel scalability.

2.2.1 Toy Data

We start by showing an example with toy data. First we use xarray to create some random, normally distributed data.

1D Histogram


```
[1]: import xarray as xr
import numpy as np
%matplotlib inline

nt, nx = 100, 30
da = xr.DataArray(np.random.randn(nt, nx), dims=['time', 'x'],
                  name='foo') # all inputs need a name

display(da)
da.plot()

<xarray.DataArray 'foo' (time: 100, x: 30)>
array([[ 1.42275651, -0.51171556, -0.35861319, ...,  0.07637755,
         1.08344715,  0.16911172],
       [ 0.12762657,  1.37398101,  0.8555511 , ..., -1.18853823,
        -2.18461814, -1.64191352],
       [ 0.07222905,  0.12172286, -0.08736727, ...,  0.51628303,
         0.49562015, -0.29984425],
       ...,
       [ 0.29185079, -0.23180969, -0.51562499, ...,  1.18842907,
        -0.09055815,  2.2021528 ],
       [ 0.54692395, -0.17940362, -1.11082126, ..., -2.01808662,
         0.15032889,  1.88620498],
       [ 1.66456683, -0.07788249,  0.71951318, ...,  1.74048813,
        -0.26434053, -1.11108024]])
Dimensions without coordinates: time, x

[1]: <matplotlib.collections.QuadMesh at 0x7f070e310910>
```



nbsphinx-code-borderwhite


By default xhistogram operates on all dimensions of an array, just like numpy. However, it operates on xarray DataArrays, taking labels into account.

```
[2]: from xhistogram.xarray import histogram
```

```
bins = np.linspace(-4, 4, 20)
h = histogram(da, bins=[bins])
display(h)
h.plot()
```

```
<xarray.DataArray 'histogram_foo' (foo_bin: 19)>
array([[ 0,  3,  7, 37, 51, 113, 234, 333, 469, 488, 445, 381, 217,
        135, 54, 25,  5,  2,  1]])
Coordinates:
  * foo_bin  (foo_bin) float64 -3.789 -3.368 -2.947 -2.526 ... 2.947 3.368 3.789
```

```
[2]: [<matplotlib.lines.Line2D at 0x7f07061638e0>]
```



nbsphinx-code-borderwhite


TODO: - Bins needs to be a list; this is annoying, would be good to accept single items - The `foo_bin` coordinate is the estimated bin center, not the bounds. We need to add the bounds to the coordinates, but we can as long as we are returning `DataArray` and not `Dataset`.

Both of the above need GitHub Issues

Histogram over a single axis

```
[3]: h_x = histogram(da, bins=[bins], dim=['time'])  
     h_x.plot()
```

```
[3]: <matplotlib.collections.QuadMesh at 0x7f070610f1c0>
```




nbsphinx-code-borderwhite

TODO: - Relax / explain requirement that dims is always a list

```
[4]: h_x.mean(dim='x').plot()
```

```
[4]: [<matplotlib.lines.Line2D at 0x7f070604a220>]
```



nbsphinx-code-borderwhite

Weighted Histogram

Weights can be the same shape as the input:

```
[5]: weights = 0.4 * xr.ones_like(da)
      histogram(da, bins=[bins], weights=weights)

[5]: <xarray.DataArray 'histogram_foo' (foo_bin: 19)>
      array([ 0. ,  1.2,  2.8, 14.8, 20.4, 45.2, 93.6, 133.2, 187.6,
             195.2, 178. , 152.4, 86.8, 54. , 21.6, 10. ,  2. ,  0.8,
             0.4])
      Coordinates:
        * foo_bin  (foo_bin) float64 -3.789 -3.368 -2.947 -2.526 ... 2.947 3.368 3.789
```

Or can use Xarray broadcasting:

```
[6]: weights = 0.2 * xr.ones_like(da.x)
      histogram(da, bins=[bins], weights=weights)

[6]: <xarray.DataArray 'histogram_foo' (foo_bin: 19)>
      array([ 0. ,  0.6,  1.4,  7.4, 10.2, 22.6, 46.8, 66.6, 93.8, 97.6, 89. ,
             76.2, 43.4, 27. , 10.8,  5. ,  1. ,  0.4,  0.2])
      Coordinates:
        * foo_bin  (foo_bin) float64 -3.789 -3.368 -2.947 -2.526 ... 2.947 3.368 3.789
```


2.2.2 2D Histogram

Now let's say we have multiple input arrays. We can calculate their joint distribution:

```
[7]: db = xr.DataArray(np.random.randn(nt, nx), dims=['time', 'x'],
                       name='bar') - 2

      histogram(da, db, bins=[bins, bins]).plot()

[7]: <matplotlib.collections.QuadMesh at 0x7f0705fd0070>
```



nbsphinx-code-borderwhite

2.2.3 Real Data Example

Ocean Volume Census in TS Space

Here we show how to use xhistogram to do a volume census of the ocean in Temperature-Salinity Space

First we open the World Ocean Atlas dataset from the opendap dataset (http://apdrc.soest.hawaii.edu/dods/public_data/WOA/WOA13/1_deg/annual).

Here we read the annual mean Temperature, Salinity and Oxygen on a 5 degree grid.

```
[8]: # Read WOA using opendap
Temp_url = 'http://apdrc.soest.hawaii.edu:80/dods/public_data/WOA/WOA13/5_deg/annual/temp'
Salt_url = 'http://apdrc.soest.hawaii.edu:80/dods/public_data/WOA/WOA13/5_deg/annual/salt'
Oxy_url = 'http://apdrc.soest.hawaii.edu:80/dods/public_data/WOA/WOA13/5_deg/annual/doxy'

ds = xr.merge([
    xr.open_dataset(Temp_url).tmn.load(),
    xr.open_dataset(Salt_url).smn.load(),
    xr.open_dataset(Oxy_url).omn.load()])
ds
```

/home/docs/checkouts/readthedocs.org/user_builds/xhistogram/conda/stable/lib/python3.9/
 ↳ site-packages/xarray/coding/times.py:150: SerializationWarning: Ambiguous reference_
 ↳ date string: 1-1-1 00:00:0.0. The first value is assumed to be the year hence will be_
 ↳ padded with zeros to remove the ambiguity (the padded reference date string is: 0001-1-
 ↳ 1 00:00:0.0). To remove this message, remove the ambiguity by padding your reference_

(continues on next page)

(continued from previous page)

```

↳date strings with zeros.
  warnings.warn(warning_msg, SerializationWarning)
/home/docs/checkouts/readthedocs.org/user_builds/xhistogram/conda/stable/lib/python3.9/
↳site-packages/xarray/coding/times.py:201: CFWarning: this date/calendar/year zero.
↳convention is not supported by CF
  cftime.num2date(num_dates, units, calendar, only_use_cftime_datetimes=True)
/home/docs/checkouts/readthedocs.org/user_builds/xhistogram/conda/stable/lib/python3.9/
↳site-packages/xarray/coding/times.py:682: SerializationWarning: Unable to decode time.
↳axis into full numpy.datetime64 objects, continuing using cftime.datetime objects.
↳instead, reason: dates out of range
  dtype = _decode_cf_datetime_dtype(data, units, calendar, self.use_cftime)
/home/docs/checkouts/readthedocs.org/user_builds/xhistogram/conda/stable/lib/python3.9/
↳site-packages/numpy/core/_asarray.py:83: SerializationWarning: Unable to decode time.
↳axis into full numpy.datetime64 objects, continuing using cftime.datetime objects.
↳instead, reason: dates out of range
  return array(a, dtype, copy=False, order=order)

```

```

[8]: <xarray.Dataset>
Dimensions: (time: 1, lev: 102, lat: 36, lon: 72)
Coordinates:
  * time      (time) object -001-01-15 00:00:00
  * lev      (lev) float64 0.0 5.0 10.0 15.0 ... 5.2e+03 5.3e+03 5.4e+03 5.5e+03
  * lat      (lat) float64 -87.5 -82.5 -77.5 -72.5 -67.5 ... 72.5 77.5 82.5 87.5
  * lon      (lon) float64 -177.5 -172.5 -167.5 -162.5 ... 167.5 172.5 177.5
Data variables:
  tmn      (time, lev, lat, lon) float32 nan nan nan nan ... nan nan nan nan
  smn      (time, lev, lat, lon) float32 nan nan nan nan ... nan nan nan nan
  omn      (time, lev, lat, lon) float32 nan nan nan nan ... nan nan nan nan
Attributes:
  long_name:  statistical mean sea water temperature [degc]

```

Use histogram to bin data points. Use canonical ocean T/S ranges, and bin size of $0.1^{\circ}C$, and $0.025psu$. Similar ranges and bin size as this review paper on Mode Waters: <https://doi.org/10.1016/B978-0-12-391851-2.00009-X>.

```

[9]: sbins = np.arange(31,38, 0.025)
     tbins = np.arange(-2, 32, 0.1)

```

```

[10]: # histogram of number of data points
     # histogram of number of data points
     hTS = histogram(ds.smn, ds.tmn, bins=[sbins, tbins])
     np.log10(hTS.T).plot(levels=31)

```

```


/home/docs/checkouts/readthedocs.org/user_builds/xhistogram/conda/stable/lib/python3.9/
↳site-packages/xarray/core/computation.py:771: RuntimeWarning: divide by zero.
↳encountered in log10
  result_data = func(*input_data)

```

```

[10]: <matplotlib.collections.QuadMesh at 0x7f0705472c70>

```

nbsphinx-code-borderwhite

However, we would like to do a volume census, which requires the data points to be weighted by volume of the grid box.

$$dV = dz * dx * dy \quad (2.1)$$


```
[11]: # histogram of number of data points weighted by volume resolution
# Note that depth is a non-uniform axis

# Create a dz variable
dz = np.diff(ds.lev)
dz = np.insert(dz, 0, dz[0])
dz = xr.DataArray(dz, coords= {'lev':ds.lev}, dims='lev')

# weight by volume of grid cell (resolution = 5degree, 1degree=110km)
dVol = dz * (5*110e3) * (5*110e3*np.cos(ds.lat*np.pi/180))

# Note: The weights are automatically broadcast to the right size
hTSw = histogram(ds.smn, ds.tmn, bins=[sbins, tbins], weights=dVol)
np.log10(hTSw.T).plot(levels=31, vmin=11.5, vmax=16, cmap='brg')
```

```
[11]: <matplotlib.collections.QuadMesh at 0x7f07053b5af0>
```



nbsphinx-code-borderwhite

The ridges of this above plot are indicative of T/S classes with a lot of volume, and some of them are indicative of Mode Waters (example Eighteen Degree water with $T \sim 18^{\circ}C$, and $S \sim 36.5psu$).

Averaging a variable

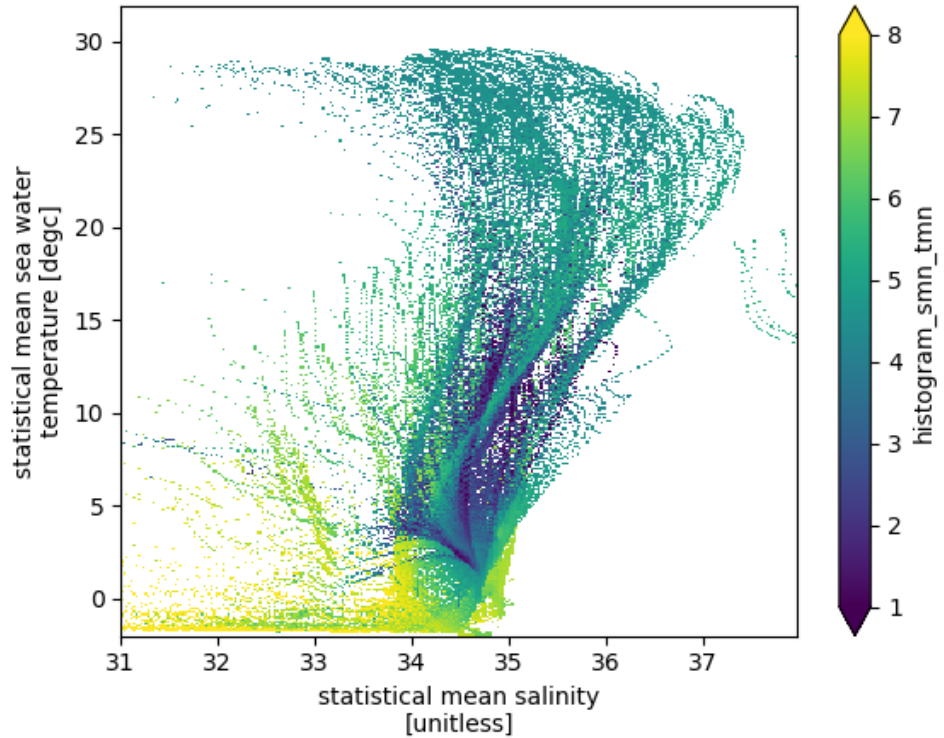
Next we calculate the mean oxygen value in each TS bin.

$$\bar{A}(m, n) = \frac{\sum_{T(x,y,z)=m, S(x,y,z)=n} (A(x, y, z) dV)}{\sum_{T(x,y,z)=m, S(x,y,z)=n} dV}. \quad (2.2)$$

```
[12]: hTSO2 = (histogram(ds.smn.where(~np.isnan(ds.omn)),
    ds.tmn.where(~np.isnan(ds.omn)),
    bins=[sbins, tbins],
    weights=ds.omn.where(~np.isnan(ds.omn))*dVol)/
    histogram(ds.smn.where(~np.isnan(ds.omn)),
    ds.tmn.where(~np.isnan(ds.omn)),
    bins=[sbins, tbins],
    weights=dVol))

(hTSO2.T).plot(vmin=1, vmax=8)
```

```
[12]: <matplotlib.collections.QuadMesh at 0x7f07042f4ac0>
```



nbsphinx-code-borderwhite

Some interesting patterns in average oxygen emerge. Convectively ventilated cold water have the highest oxygen and mode waters have relatively high oxygen. Oxygen minimum zones are interspersed in the middle of volumetric ridges (high volume waters).

NOTE: NaNs in weights will make the weighted sum as nan. To avoid this, call `.fillna(0.)` on your weights input data before calling `histogram()`.

2.2.4 Dask Integration

Should just work, but need examples.

2.3 API

2.3.1 Core Module

Numpy API for xhistogram.

`xhistogram.core.histogram(*args, bins=None, range=None, axis=None, weights=None, density=False, block_size='auto')`

Histogram applied along specified axis / axes.

Parameters

args

[array_like] Input data. The number of input arguments determines the dimensionality of the histogram. For example, two arguments produce a 2D histogram. All args must have the same size.

bins

[int, str or numpy array or a list of ints, str and/or arrays, optional] If a list, there should be one entry for each item in `args`. The bin specifications are as follows:

- If int; the number of bins for all arguments in `args`.
- If str; the method used to automatically calculate the optimal bin width for all arguments in `args`, as defined by numpy `histogram_bin_edges`.
- If numpy array; the bin edges for all arguments in `args`.
- If a list of ints, str and/or arrays; the bin specification as above for every argument in `args`.

When bin edges are specified, all but the last (righthand-most) bin include the left edge and exclude the right edge. The last bin includes both edges.

A `TypeError` will be raised if `args` or `weights` contains dask arrays and `bins` are not specified explicitly as an array or list of arrays. This is because other bin specifications trigger computation.

range

[(float, float) or a list of (float, float), optional] If a list, there should be one entry for each item in `args`. The range specifications are as follows:

- If (float, float); the lower and upper range(s) of the bins for all arguments in `args`. Values outside the range are ignored. The first element of the range must be less than or equal to the second. `range` affects the automatic bin computation as well. In this case, while bin width is computed to be optimal based on the actual data within `range`, the bin count will fill the entire range including portions containing no data.
- If a list of (float, float); the ranges as above for every argument in `args`.
- If not provided, range is simply `(arg.min(), arg.max())` for each arg.

axis

[None or int or tuple of ints, optional] Axis or axes along which the histogram is computed. The default is to compute the histogram of the flattened array

weights

[array_like, optional] An array of weights, of the same shape as `a`. Each value in `a` only contributes its associated weight towards the bin count (instead of 1). If `density` is `True`, the weights are normalized, so that the integral of the density over the range remains 1.

density

[bool, optional] If `False`, the result will contain the number of samples in each bin. If `True`, the result is the value of the probability *density* function at the bin, normalized such that the *integral* over the range is 1. Note that the sum of the histogram values will not be equal to 1 unless bins of unity width are chosen; it is not a probability *mass* function.

block_size

[int or 'auto', optional] A parameter which governs the algorithm used to compute the histogram. Using a nonzero value splits the histogram calculation over the non-histogram axes into blocks of size `block_size`, iterating over them with a loop (numpy inputs) or in parallel (dask inputs). If 'auto', blocks will be determined either by the underlying dask chunks (dask inputs) or an experimental built-in heuristic (numpy inputs).

Returns**hist**

[array] The values of the histogram.

bin_edges

[list of arrays] Return the bin edges for each input array.

See also:

`numpy.histogram`, `numpy.bincount`, `numpy.searchsorted`

2.3.2 Xarray Module

Xarray API for xhistogram.

`xhistogram.xarray.histogram(*args, bins=None, range=None, dim=None, weights=None, density=False, block_size='auto', keep_coords=False, bin_dim_suffix='_bin')`

Histogram applied along specified dimensions.

Parameters

args

[xarray.DataArray objects] Input data. The number of input arguments determines the dimensionality of the histogram. For example, two arguments produce a 2D histogram. All args must be aligned and have the same dimensions.

bins

[int, str or numpy array or a list of ints, str and/or arrays, optional] If a list, there should be one entry for each item in args. The bin specifications are as follows:

- If int; the number of bins for all arguments in args.
- If str; the method used to automatically calculate the optimal bin width for all arguments in args, as defined by `numpy.histogram_bin_edges`.
- If numpy array; the bin edges for all arguments in args.
- If a list of ints, str and/or arrays; the bin specification as above for every argument in args.

When bin edges are specified, all but the last (righthand-most) bin include the left edge and exclude the right edge. The last bin includes both edges.

A `TypeError` will be raised if args or weights contains dask arrays and bins are not specified explicitly as an array or list of arrays. This is because other bin specifications trigger computation.

range

[(float, float) or a list of (float, float), optional] If a list, there should be one entry for each item in args. The range specifications are as follows:

- If (float, float); the lower and upper range(s) of the bins for all arguments in args. Values outside the range are ignored. The first element of the range must be less than or equal to the second. *range* affects the automatic bin computation as well. In this case, while bin width is computed to be optimal based on the actual data within *range*, the bin count will fill the entire range including portions containing no data.
- If a list of (float, float); the ranges as above for every argument in args.
- If not provided, range is simply `(arg.min(), arg.max())` for each arg.

dim

[tuple of strings, optional] Dimensions over which which the histogram is computed. The default is to compute the histogram of the flattened array.

weights

[array_like, optional] An array of weights, of the same shape as *a*. Each value in *a* only contributes its associated weight towards the bin count (instead of 1). If *density* is True, the weights are normalized, so that the integral of the density over the range remains 1. NaNs in the weights input will fill the entire bin with NaNs. If there are NaNs in the weights input call `.fillna(0.)` before running `histogram()`.

density

[bool, optional] If False, the result will contain the number of samples in each bin. If True, the result is the value of the probability *density* function at the bin, normalized such that the *integral* over the range is 1. Note that the sum of the histogram values will not be equal to 1 unless bins of unity width are chosen; it is not a probability *mass* function.

block_size

[int or 'auto', optional] A parameter which governs the algorithm used to compute the histogram. Using a nonzero value splits the histogram calculation over the non-histogram axes into blocks of size `block_size`, iterating over them with a loop (numpy inputs) or in parallel (dask inputs). If 'auto', blocks will be determined either by the underlying dask chunks (dask inputs) or an experimental built-in heuristic (numpy inputs).

keep_coords

[bool, optional] If True, keep all coordinates. Default: False

bin_dim_suffix

[str, optional] Suffix to append to input arg names to define names of output bin dimensions

Returns**hist**

[xarray.DataArray] The values of the histogram. For each bin, the midpoint of the bin edges is given along the bin coordinates.

2.4 Contributor Guide

This package is in very early stages. Lots of work is needed.

You can help out just by using `xhistogram` and reporting [issues](#).

The following sections cover some general guidelines for maintainers and contributors wanting to help develop `xhistogram`.

2.4.1 Feature requests, suggestions and bug reports

We are eager to hear about any bugs you have found, new features you would like to see and any other suggestions you may have. Please feel free to submit these as [issues](#).

When suggesting features, please make sure to explain in detail how the proposed feature should work and to keep the scope as narrow as possible. This makes features easier to implement in small PRs.

When report bugs, please include:

- Any details about your local setup that might be helpful in troubleshooting, specifically the Python interpreter version, installed libraries, and `xhistogram` version.
- Detailed steps to reproduce the bug, ideally a [Minimal, Complete and Verifiable Example](#).
- If possible, a demonstration test that currently fails but should pass when the bug is fixed.

2.4.2 Write documentation

Adding documentation is always helpful. This may include:

- More complementary documentation. Have you perhaps found something unclear?
- Docstrings.
- Example notebooks of `xhistogram` being used in real analyses.

The `xhistogram` documentation is written in reStructuredText. You can follow the conventions in already written documents. Some helpful guides can be found [here](#) and [here](#).

When writing and editing documentation, it can be useful to see the resulting build without having to push to Github. You can build the documentation locally by running:

```
$ # Install the packages required to build the docs in a conda environment
$ conda env create -f doc/environment.yml
$ conda activate xhistogram_doc_env
$ # Install the latest xhistogram
$ pip install --no-deps -e .
$ cd doc/
$ make html
```

This will build the documentation locally in `doc/_build/`. You can then open `_build/html/index.html` in your web browser to view the documentation. For example, if you have `xdg-open` installed:

```
$ xdg-open _build/html/index.html
```

To lint the reStructuredText documentation files run:

```
$ doc8 doc/*.rst
```

2.4.3 Preparing Pull Requests

1. Fork the `xhistogram` [GitHub repository](#). It's fine to use `xhistogram` as your fork repository name because it will live under your username.
2. Clone your fork locally, connect your repository to the upstream (main project), and create a branch to work on:

```
$ git clone git@github.com:YOUR_GITHUB_USERNAME/xhistogram.git
$ cd xhistogram
$ git remote add upstream git@github.com:xgcm/xhistogram.git
$ git checkout -b your-bugfix-feature-branch-name master
```

If you need some help with Git, follow [this quick start guide](#)

3. Install dependencies into a new conda environment:

```
$ conda env create -f ci/environment-3.9.yml
$ conda activate xhistogram_test_env
```

4. Install `xhistogram` using the editable flag (meaning any changes you make to the package will be reflected directly in your environment):

```
$ pip install --no-deps -e .
```

5. Start making your edits. Please try to type annotate your additions as much as possible. Adding type annotations to existing unannotated code is also very welcome. You can read about Python typing [here](#).
6. Break your edits up into reasonably sized commits:

```
$ git commit -a -m "<commit message>"
$ git push -u
```

It can be useful to manually run `pre-commit` as you make your edits. `pre-commit` will run checks on the format and typing of your code and will show you where you need to make changes. This will mean your code is more likely to pass the CI checks when you push it:

```
$ pip install pre_commit # you only need to do this once
$ pre-commit run --all-files
```

7. Run the tests (including those you add to test your edits!):

```
$ pytest xhistogram
```

You can also test that your contribution and tests increased the test coverage:

```
$ coverage run --source xhistogram -m py.test
$ coverage report
```

8. Add a new entry describing your contribution to the *Release History* in `doc/contributing.rst`. Please try to follow the format of the existing entries.
9. Submit a pull request through the GitHub [website](#).

Note that you can create the Pull Request while you're working on your PR. The PR will update as you add more commits. `xhistogram` developers and contributors can then review your code and offer suggestions.

2.4.4 Release History

v0.3.2

- Fix bug producing `TypeError` when `weights` is provided with `keep_coords=True` [GH78](#). By [Dougie Squire](#).
- Raise `TypeError` when `weights` is a dask array and bin edges are not explicitly provided [GH12](#). By [Dougie Squire](#).

v0.3.1

- Add DOI badge and `CITATION.cff`. By [Julius Busecke](#).

v0.3.0

- Add support for histograms over non-float dtypes (e.g. datetime objects) [GH25](#). By [Dougie Squire](#).
- Refactor histogram calculation to use `dask.array.blockwise` when input arguments are dask arrays, resulting in significant performance improvements [GH49](#). By [Ryan Abernathy](#), [Tom Nicholas](#) and [Gabe Joseph](#).
- Fixed bug with density calculation when NaNs are present [GH51](#). By [Dougie Squire](#).
- Implemented various options for users for providing bins to `xhistogram` that mimic the numpy histogram API. This included adding a range argument to the `xhistogram` API [GH13](#). By [Dougie Squire](#).

- Added a function to check if the object passed to xhistogram is an `xarray.DataArray` and if not, throw an error. [GH14](#). By [Yang Yunyi](#).

v0.2.0

- Added `FutureWarning` for upcoming changes to core API [GH13](#). By [Dougie Squire](#).
- Added documentation on how to deal with NaNs in weights [GH26](#). By [Shanice Bailey](#).
- Move CI to GitHub actions [GH32](#). By [James Bourbeau](#).
- Add documentation for contributors. By [Dougie Squire](#).
- Add type checking with `mypy` [GH32](#). By [Dougie Squire](#).

v0.1.3

- Update dependencies to exclude incompatible dask version [GH27](#). By [Ryan Abernathey](#).

v0.1.2

- Aligned definition of `bins` with `numpy.histogram` [GH18](#). By [Dougie Squire](#).

v0.1.1

Minor bugfix release

- Improved documentation examples. By [Dhruv Balwada](#).
- Fixed issue [GH5](#) related to incorrect dimension order and dropping of dimension coordinates. By [Ryan Abernathey](#).

v0.1

First release

PYTHON MODULE INDEX

X

`xhistogram.core`, 15

`xhistogram.xarray`, 17

INDEX

H

histogram() (*in module xhistogram.core*), 15
histogram() (*in module xhistogram.xarray*), 17

M

module
 xhistogram.core, 15
 xhistogram.xarray, 17

X

xhistogram.core
 module, 15
xhistogram.xarray
 module, 17