

Paleoceanography and Paleoclimatology[®]

RESEARCH ARTICLE

10.1029/2022PA004509

Key Points:

- `Pyleoclim` makes timeseries analysis tools accessible to practicing scientists, via a user-friendly Python package
- Three Jupyter Notebooks illustrate how `Pyleoclim` facilitates common and advanced tasks
- `Pyleoclim` can enhance reproducibility and rigor of paleogeoscientific workflows involving timeseries

Correspondence to:

D. Khider,
khider@usc.edu

Citation:

Khider, D., Emile-Geay, J., Zhu, F., James, A., Landers, J., Ratnakar, V., & Gil, Y. (2022). `Pyleoclim`: Paleoclimate timeseries analysis and visualization with Python. *Paleoceanography and Paleoclimatology*, 37, e2022PA004509. <https://doi.org/10.1029/2022PA004509>

Received 11 JUL 2022

Accepted 27 SEP 2022

Pyleoclim: Paleoclimate Timeseries Analysis and Visualization With Python

Deborah Khider¹ , Julien Emile-Geay² , Feng Zhu³, Alexander James², Jordan Landers², Varun Ratnakar¹, and Yolanda Gil¹ 

¹University of Southern California, Information Sciences Institute, Marina Del Rey, CA, USA, ²Department of Earth Sciences, University of Southern California, Los Angeles, CA, USA, ³School of Atmospheric Sciences, Nanjing University of Information Science and Technology, Nanjing, China

Abstract We present a Python package geared toward the intuitive analysis and visualization of paleoclimate timeseries, `Pyleoclim`. The code is open-source, object-oriented, and built upon the standard scientific Python stack, allowing users to take advantage of a large collection of existing and emerging techniques. We describe the code's philosophy, structure, and base functionalities and apply it to three paleoclimate problems: (a) orbital-scale climate variability in a deep-sea core, illustrating spectral, wavelet, and coherency analysis in the presence of age uncertainties; (b) correlating a high-resolution speleothem to a climate field, illustrating correlation analysis in the presence of various statistical pitfalls (including age uncertainties); (c) model-data confrontations in the frequency domain, illustrating the characterization of scaling behavior. We show how the package may be used for transparent and reproducible analysis of paleoclimate and paleoceanographic datasets, supporting Findable, Accessible, Interoperable, and Reusable software and an open science ethos. The package is supported by an extensive documentation and a growing library of tutorials shared publicly as videos and cloud-executable Jupyter notebooks, to encourage adoption by new users.

Plain Language Summary This article describes a software application called `Pyleoclim` meant to help scientists analyze datasets of ordered observations, particularly applicable to the study of past climates, environments, and ecology. `Pyleoclim` is meant to be used by domain scientists as well as students interested in learning more about Earth's climate through examples provided in the documentation and online tutorials. `Pyleoclim` is intended to help scientists save time with their analyses, documenting the steps for better transparency, and as such, allows other scientists to reproduce results from previous studies.

1. Introduction

As paleoclimate and paleoceanographic data continue to increase in size, diversity, and quality, it remains a longstanding challenge to adequately extract and visualize the quantitative information present in such records so as to constrain model estimates of past and future change (National Academies of Sciences, Engineering, and Medicine, 2021). Indeed, these datasets often violate basic statistical assumptions (i.e., normality, independence, even sampling in time, high signal-to-noise ratio), requiring specific tools and workflows that go beyond what can be found in standard software libraries. In addition to recent efforts in *R* (McKay et al., 2021) and Matlab (Greene et al., 2019), a similar offering in the Python research ecosystem was heretofore lacking. Python's popularity among physical and data scientists has been on the rise (Perkel, 2015), with a growing collection of libraries for data analysis (e.g., `pandas` (McKinney, 2010), `statsmodels` (Seabold & Perktold, 2010), `SciPy` (Virtanen et al., 2020)) and visualization (e.g., `matplotlib` (Hunter, 2007), `seaborn` (Waskom, 2021) and `Cartopy` (Elson et al., 2022)), including libraries tailored to climate research (e.g., `xarray` (Hoyer & Hamman, 2017) and `climlab` (Rose, 2018)). However, none of the existing packages can natively handle the challenges of paleoclimatological and paleoceanographic datasets (i.e., observations are often unevenly-spaced in time, uncertainties are present in both abscissa and ordinate, proxies hold an often complex relationship to dynamically-relevant variables). As such, standard analysis methods do not work "out-of-the-box," often requiring time-consuming adaptation by users. In addition, several well-established statistical techniques (e.g., controlling for spurious null hypothesis rejection with the False Discovery Rate (Benjamini & Hochberg, 1995) or performing wavelet analysis on unevenly-spaced data (Foster, 1996)) are not currently implemented in a widely-available, well-documented, and user-friendly package in a major programming language. Lastly, there is a persistent language barrier between data generated by paleo-observations and model simulations, which few frameworks address

© 2022 The Authors.

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial License](#), which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

explicitly, particularly from the viewpoint of uncertainty quantification (Dee et al., 2015). To remedy this situation, we present `Pyleoclim`, a Python package specifically designed for scientific studies in paleoceanography and paleoclimatology, using data generated from both observations or models. While it is impossible to anticipate all user needs, the package is meant to provide a one-stop shop for the most common tasks encountered in the analysis of timeseries in our field, like interpolation, filtering, spectral and wavelet analysis, correlation analysis, principal component analysis (PCA), and many more. It has been, and will continue to be, used for research and teaching.

The remainder of this paper is organized as follows: Section 2 describes the `Pyleoclim` codebase and its re-use of emerging data standards for paleoclimate datasets; Section 3 describes three case studies, highlighting how `Pyleoclim` allows for Findable, Accessible, Interoperable, and Reusable (FAIR) paleoclimate research; Section 4 provides a conclusion and outlook toward future versions and scientific uses of the package.

2. The `Pyleoclim` Codebase

2.1. Philosophy

`Pyleoclim` was designed to harness the power of various Python libraries for data science (e.g., NumPy (Harris et al., 2020), Pandas (McKinney, 2010), SciPy (Virtanen et al., 2020), and scikit-learn (Pedregosa et al., 2011)) and visualization (Matplotlib (Hunter, 2007), seaborn (Waskom, 2021), and Cartopy (Elson et al., 2022)) for paleoclimatology and paleoceanography. The user application programming interface (API) is designed around manipulating objects (such as a time series) for analysis. This design, called object-oriented programming (OOP), places the data at the center of the analysis, rather than the functions. The objects contain both data and metadata in the form of fields that can be entered by a user (e.g., a timeseries would require at least values for time and the quantity being measured in time, but optionally allow for labels and units) and code that represents procedures that are applicable to each object. The number of data and metadata fields is dictated by the procedures (and their desired level of automation). OOP is ubiquitous in Python and presents several advantages over method-oriented programming: it follows the natural relationship between an object and a method, with each call representing a clearly defined action that helps constructing workflows through method chaining (for an example, see Section 2.3).

`Pyleoclim` is supported by extensive documentation (<https://pyleoclim-util.readthedocs.io/>) that provides minimal usage examples for the code. Scientific examples in the form of Jupyter notebooks (Kluyver et al., 2016) are available on several GitHub repositories (Emile-Geay et al., 2019; Khider, Emile-Geay, James, et al., 2022; Khider, Emile-Geay, & Zhu, 2022; Khider, Emile-Geay, Zhu, & James, 2022). Tutorials are also provided on YouTube (https://www.youtube.com/playlist?list=PL93NbaRnKAuF4WpIQf-4y_U4lo-Gqrcw) and in the form of a Jupyter Book (<http://linked.earth/PyleoTutorials/>). The LinkedEarth Discourse forum (<https://discourse.linked.earth>) also provides an avenue to discuss the science applications of `Pyleoclim`.

The package is open-source and follows the principle of Open Development. As such, the code is available on GitHub under an open-source license. A contributing guide (https://pyleoclim-util.readthedocs.io/en/master/contribution_guide.html) details how the community can engage in `Pyleoclim`'s development. The simplest level of engagement is to report bugs as GitHub issues and starting community discussions about scientific use cases on the LinkedEarth Discourse forum (<https://discourse.linked.earth>). More proficient programmers can also contribute by upgrading existing functionalities or creating new ones through GitHub pull requests.

Finally, publishers and funding agencies are increasingly promoting the principles of FAIR science, not only for data (Wilkinson et al., 2016) but also software (Lamprecht et al., 2020) and scientific workflows (Goble et al., 2020). `Pyleoclim` follows the guidelines set forth for FAIR software: it is available and versioned on GitHub, licensed under a GNU public license, registered on the Python Package Index (Pypi), and citable from a Zenodo Digital Object Identifier. Various versions of the software are available through Docker containers stored on quay.io. As such, `Pyleoclim` supports the development of FAIR scientific workflows (Goble et al., 2020).

2.2. Functionalities

`Pyleoclim` contains functionalities designed to help users customize their own workflows from data pre-processing (such as standardizing, detrending, removing outliers, placing time series on a common time axis) to analysis (spectral and wavelet analysis, paleo-aware correlation, spatial and temporal decomposition) and visualization of the results. Most `Pyleoclim` functionalities leverage existing and well-documented software packages:

Visualizations were built upon the `Matplotlib` (Hunter, 2007) and `seaborn` packages (Waskom, 2021). Mapping capabilities are provided through `Cartopy` (Elson et al., 2022).

Signal processing and statistics: the `SciPy` package (Virtanen et al., 2020) supports signal processing functionalities, including methods for digital filtering and spectral analysis (namely the basic periodogram, Welch's periodogram, and the Lomb-Scargle periodogram (VanderPlas, 2018)). `Pyleoclim` also allows for the use of the multi-taper method (Thomson, 1982) as implemented in `nitime` (Millman & Brett, 2007), many types of interpolation (e.g., linear, quadratic, natural splines), statistics (e.g., kernel density estimation, quantile estimation) and various optimization functions used internally by `Pyleoclim`.

Machine Learning: the `scikit-learn` (Pedregosa et al., 2011) package supports clustering for outlier detection.

Timeseries modeling `statsmodels` (Seabold & Perktold, 2010) supports principal component analysis (PCA (Hannachi et al., 2007)), parametric timeseries modeling, and Granger causality estimation.

Wavelet analysis via the continuous wavelet transform, as implemented in Matlab by Torrence and Compo (1998), was recently ported to Python (Predybaylo et al., 2022).

These basic functionalities were adjusted for paleoclimate data either by changing the default parameter values to ones more appropriate for the data characteristics, raising errors when appropriate (e.g., when trying to apply a method meant for evenly-spaced series on an unevenly-spaced series), or performing regridding within the analysis function at the user's request.

In addition, some functionalities were coded in Python specifically for the package, such as the Weighted Wavelet Z-Transform (Foster, 1996; Kirchner & Neal, 2013) and Liang-Kleeman causality (Liang, 2013, 2014, 2015, 2016, 2018). Because of the nonlinear and nonstationary nature of many paleoclimate timeseries (Ghil et al., 2002), `Pyleoclim` features advanced detrending techniques such as empirical mode decomposition (Huang et al., 1998) and Savitzky-Golay filtering (Savitzky & Golay, 1964). On the analysis side, `Pyleoclim` enables Singular Spectrum Analysis (SSA) (Ghil et al., 2002; Vautard et al., 1992; Vautard & Ghil, 1989), including significance testing for “red” timeseries (Allen & Smith, 1996) and tolerance for missing values (Schoellhamer, 2001), which enables SSA to be used as an interpolant.

All these functionalities are available through the `Pyleoclim` utilities APIs, which are meant for developers and apply to `NumPy` (Harris et al., 2020) arrays. This means that those methods, which often are not specific to observational paleoclimate data, can easily be repurposed by other packages that rely on arrayed data (e.g., climate model output). However, most users are expected to interact with the `Pyleoclim` user APIs, which group these functionalities into a common interface attached to specific objects, which we now describe.

2.3. User API

The main interface for `Pyleoclim` revolves around objects that can be manipulated for analysis (Figure 1). The functionalities described in Section 2.2 are grouped into object methods that offer a common interface to call the various functions from the supporting libraries and internally handle the data transformation for these functions. At the user level, `Pyleoclim` allows scientists to concentrate on their workflows rather than handle data transformations among the various Python data objects and types.

The main object in `Pyleoclim` is the `Series` object, which takes as arguments the values for time and the variable of interest, as well as their names and units. These `Series` objects can be easily created from various file formats, for example, csv files:

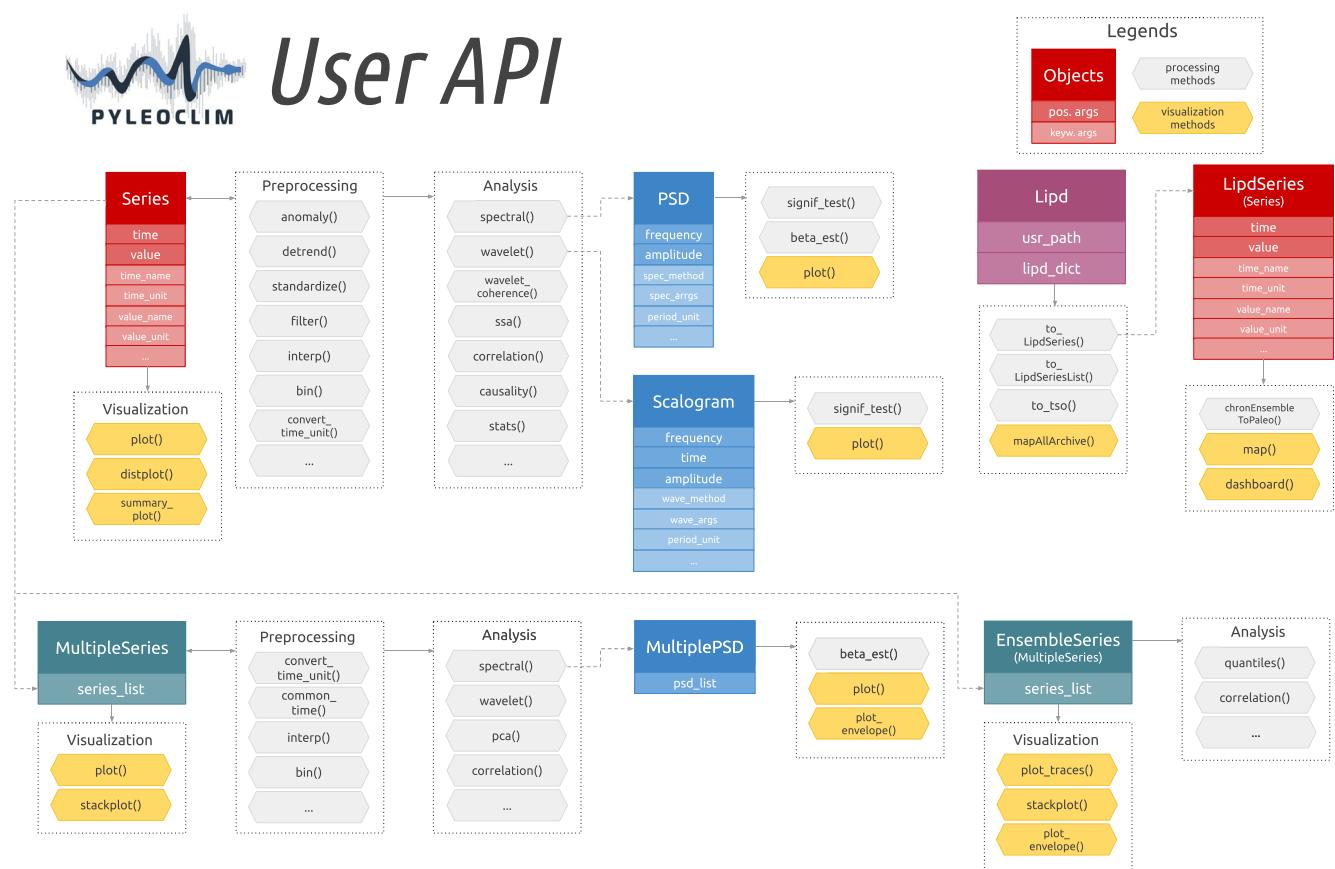


Figure 1. Diagram of the objects and associated functions in the Pyleoclim user application programming interfaces (APIs).

```

[1] import pandas as pd
[2] import pyleoclim as pyleo
[3] url = 'https://raw.githubusercontent.com/LinkedEarth/Pyleoclim_util/' + \
'master/example_data/oni.csv'
[4] df = pd.read_csv(url,header=0)
[5] ts = pyleo.Series(time=df['Dec year'],value=df['NINO34_ANOM'],
                     time_name='Year', value_name='SST anomaly',
                     time_unit='CE', value_unit='$^\circ C',
                     label='Ni\~{n}o 3.4', clean_ts=True)

```

The `Series` object `ts` contains both the data in the `time` and `value` arguments as well as relevant metadata, such as the name and units of each variable. The metadata become especially relevant for plotting; however, Pyleoclim has a rudimentary understanding of paleo-relevant time and attempts to correct time units when two series are compared (for instance one in kyr BP and the other in yr BP). The `label` metadata is used to build the legend on figures. The argument `clean_ts` is used here to remove NaNs and sort the timeseries in increasing time.

Once the data are loaded into a `Series` object, complex analyses can be made through simple commands. For illustrative purposes, we run it through spectral and wavelet analysis:

```
[6] ts_detrend = ts.detrend() # remove trends
[7] ts_interp = ts_detrend.interp() # interpolate over missing values
[8] ts_std = ts_interp.standardize() # standardizing
[9] PSD = ts_std.spectral(method='mtm') #spectral analysis
[10] PSD_signif = PSD.signif_test() #run AR(1) significance test
```

Code lines [6]-[8] correspond to pre-processing steps (in this case, detrending, interpolation, and standardizing) using the default methods in `Pyleoclim`. The spectral density is computed through the MTM method, and the result stored in a new `PSD` object, from which a significance test against an AR(1) benchmark (Emile-Geay, 2017) can be performed.

One advantage of OOP is method chaining: since each method returns a `Pyleoclim` object, the calls can be chained together in a single statement without having to store the intermediate results. With method chaining, the block code above can be rewritten as a single line:

```
PSD_signif = ts.detrend().interp().standardize().spectral(method='mtm').signif_test()
```

It can be beneficial to limit the chaining to the pre-processing steps so the resulting `Series` can be used with other methods like wavelet analysis, which produces a `Scalogram` object:

```
[11] scal = ts_std.wavelet(method='cwt') #wavelet analysis
[12] scal_signif=scal.signif_test(method='ar1asym') #run AR(1) significance test
```

The wavelet analysis presented here follows the method of Torrence and Compo (1998) to obtain the scalogram and significance level. `Pyleoclim` contains various methods to visualize timeseries, periodograms, and scalograms. Here, we will generate a summary of our analysis through a single method:

```
[13] fig, ax = ts.summary_plot(PSD_signif, scal_signif,
                               time_lim=[1871,2022],
                               value_lim=[-3.5,3.5],
                               psd_label='PSD',
                               time_label='',
                               ts_plot_kwarg={'lgd_kwarg':{'loc':'upper right',
                               'bbox_to_anchor':(1.4,0.95)}},
                               gridspec_kwarg={'hspace':0,'wspace':0}) #plot
[14] ax['cb'].set_xlabel('Amplitude')
```

The resulting figure is shown in Figure 2. All figures generated by `Pyleoclim` are highly customizable, either directly through our APIs or `Matplotlib/Cartopy`. Let's examine the code above, which provides examples of the various options. Line [13] is for the direct customization of the resulting plot through `Pyleoclim` with the following information: the limits for the time axis through the `time_lim` argument, the limits for the y-axis of the timeseries plot (`value_lim` argument), a new `x`-axis label for the periodogram (`psd_label` argument), removal of the time axis label (`time_label` argument), a dictionary of `Matplotlib` arguments to deal with legend placement for the timeseries plot, and another dictionary to deal with the spacing between the various plots.

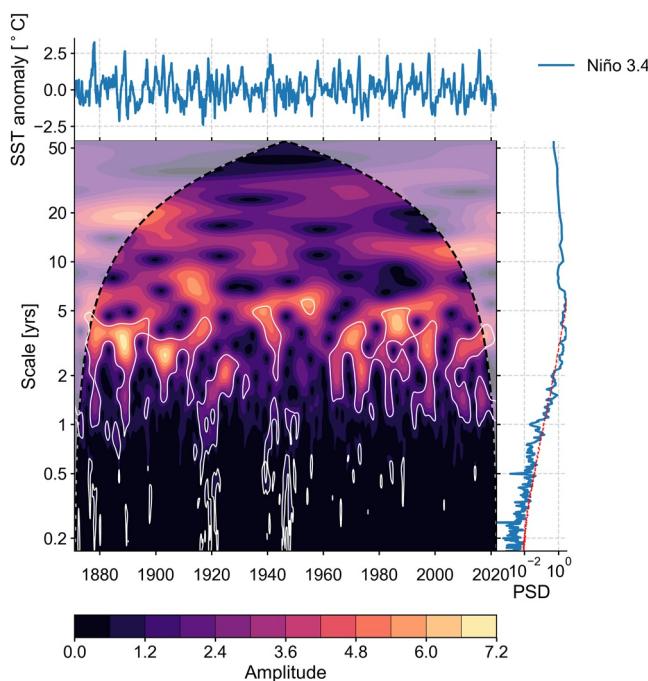


Figure 2. Summary of the spectral and wavelet analysis performed on the Niño 3.4 sea surface temperature (SST) anomalies timeseries as encoded in `Pyleoclim`. The series displays significant power in the 2–7 years band, consistent with the El Niño Southern Oscillation.

Line [14] sets an appropriate label for the colorbar.

Note that these plots can also be obtained individually:

```
ts.plot()  
PSD_signif.plot()  
scal_signif.plot()
```

Even though plotting methods are available for the `Series`, `PSD`, and `Scalogram` objects, the behavior depends on the object to which it is attached. This is another advantage of OOP: since the methods are attached to objects, they can share a name for a similar action (e.g., plotting) while behaving in a manner appropriate for each object.

Although we expect that users will be creating `Series` objects from an existing file (e.g., `xls`, `csv`, `NOAA`, `PANGAEA`, `netCDF`), many `Pyleoclim` objects are generated as results of the analysis. For instance `PSD` is generated by spectral analysis methods, `Scalogram` by wavelet analysis methods, `Coherence` by cross-wavelet analysis methods, and `Corr` by correlation methods. Object creation in the development of `Pyleoclim` was motivated by the need to attach specific methods with specific behavior to particular objects (e.g., significance testing for spectral and wavelet analysis or plotting methods).

Several objects use the prefix `Multiple` (e.g., `MultipleSeries`, `MultiplePSD`), which signal that this object is comprised of a list of the basic `Pyleoclim` objects. For instance, the `MultipleSeries` object contains several `Series` objects, with dedicated plotting (e.g., `stackplot()`) and analysis (e.g., principal component analysis) methods that are applicable to collections of paleoclimate timeseries.

2.4. Leveraging Paleoclimate Data Standards

In addition to the data science and visualization libraries mentioned above, `Pyleoclim` is compatible with the Linked Paleo Data (LiPD (McKay & Emile-Geay, 2016)) format. LiPD is a universally-readable data container that stores metadata in a JSON-LD file (JavaScript Object Notation for Linked Data) and the data in tables saved in CSV format. Utilities have been written in Matlab, Python, and *R* to manipulate these metadata-rich files. Consequently, we created two objects in `Pyleoclim` that take advantage of the additional, standardized metadata: the `LiPD` object, which allows users to deal with one file or a collection of files and have mapping capabilities, and the `LiPDSeries` object, a child of the `Series` object. As such, `LiPDSeries` inherits all the methods available for `Series` with additional functionalities that take advantage of the richness of the metadata, such as dashboards for displaying relevant information (Figure 3).

3. Three Paleoclimate Studies Enabled by `Pyleoclim`

To illustrate the use of `Pyleoclim` in research, we summarize three studies available as fully executable Jupyter Notebooks as companion to this manuscript (see the code availability statement in the Data Availability section). The first study walks through spectral, wavelet, and cross-wavelet analysis in the presence of age uncertainties. The second study is reproduced from Hu et al. (2017) and presents the pitfalls of using correlation analysis for the interpretation of a paleoclimate record. Finally, the last study shows how to reproduce the results of Zhu et al. (2019), using spectral analysis to assess whether current models can capture the continuum of climate variability.

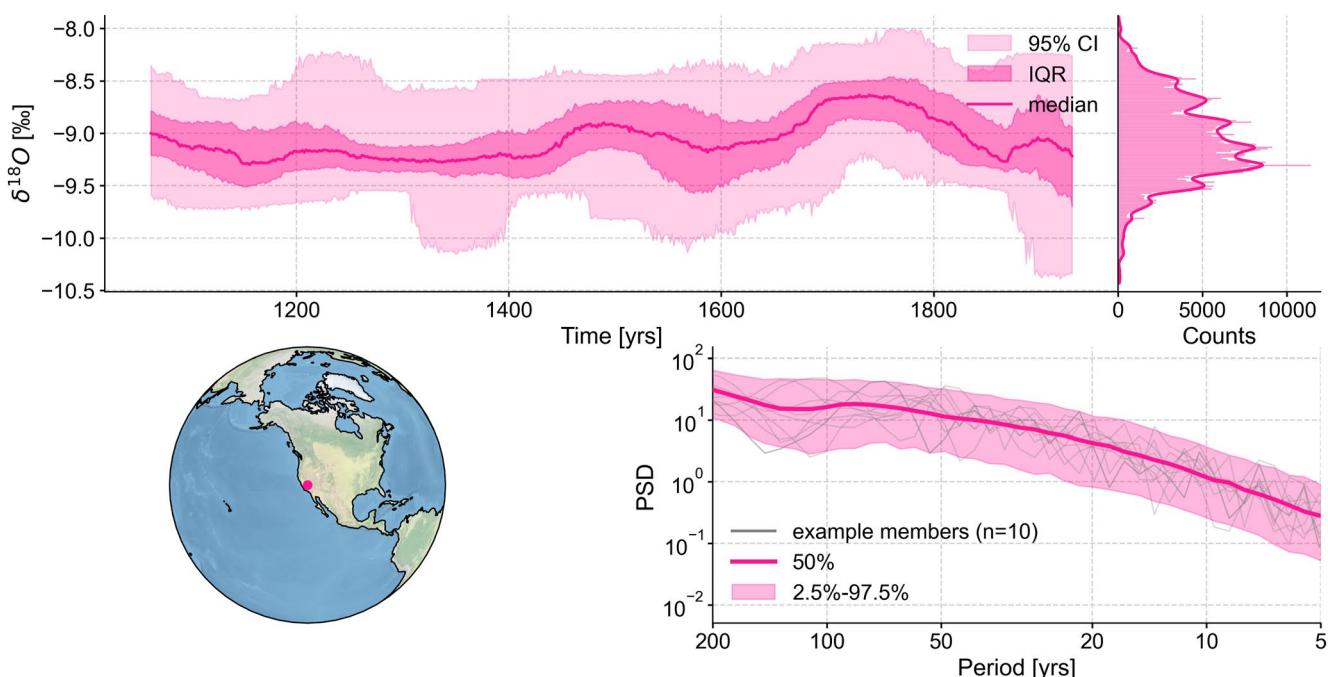


Figure 3. Example dashboard in Pyleoclim enabled by LiPD. The dashboard consists of four panels: the top left panel plots the timeseries, in this case the speleothem record from Crystal Cave (McCabe-Glynn et al., 2013). Note that axis labels and legend are automatically generated from the metadata in the file. The envelope represents the age uncertainty obtained from Bchron (Haslett & Parnell, 2008), a Bayesian age modeling software. The top right panel shows the distribution of values. The bottom left panel displays the location of the record while the bottom right displays the results of spectral analysis using the Lomb-Scargle method. To assess the effect of age uncertainty on the interpretation of the peaks in the record, the spectral analysis is performed on each of the members present in the age ensemble from Bchron.

3.1. Orbital-Scale Climate Variability in a Deep Sea Core

The first case study concerns the analysis of paleoclimate records in the frequency domain (specifically spectral, wavelet, and coherence analysis). This type of analysis is often performed to look at common periodicities among records or between a record and its hypothesized forcing. Analysis of paleoclimate time series in the frequency domain is complicated by several factors:

Irregular sampling: most spectral methods are designed for series that are evenly spaced in time. Hypothesizing over missing values can bias the statistical results and enhance the low-frequency components of the spectrum at the expense of the high-frequency components (Schulz & Mudelsee, 2002; Schulz & Stattegger, 1997). Methods that do not require interpolation, such as the Lomb-Scargle periodogram (Lomb, 1976; Scargle, 1982, 1989), also have known biases (Schulz & Mudelsee, 2002; VanderPlas, 2018). The trade-offs of the various options need to be carefully examined in light of the data.

Pre-processing steps: in addition to interpolation, detrending and removal of outliers can affect the results of the analysis. Whether to use these options needs to be evaluated for the specific data set and hypothesis to be tested.

Age uncertainties: age uncertainties affect the location of features in time, so methods need to allow for an ensemble of plausible chronologies (generated, for instance, by a Bayesian age model).

Pyleoclim offers a variety of pre-processing and spectral/wavelet analysis methods to allow for a robust assessment of the time series characteristics in the frequency domain. This section and accompanying notebook walks the reader through spectral, wavelet, and coherence analysis of a marine deep sea record (Ocean Drilling Program (ODP) Site 846) covering the past 5 million years and obtained from benthic $\delta^{18}\text{O}$ (Mix et al., 1995; Shackleton et al., 1995) and alkenone paleothermometry (Lawrence et al., 2006). The core location is in the Eastern tropical Pacific (3.1°S, 90.8°W, 3296 m). The age model (Khider et al., 2017) for the record was obtained by aligning the benthic record to the benthic stack of Lisiecki and Raymo (2005,

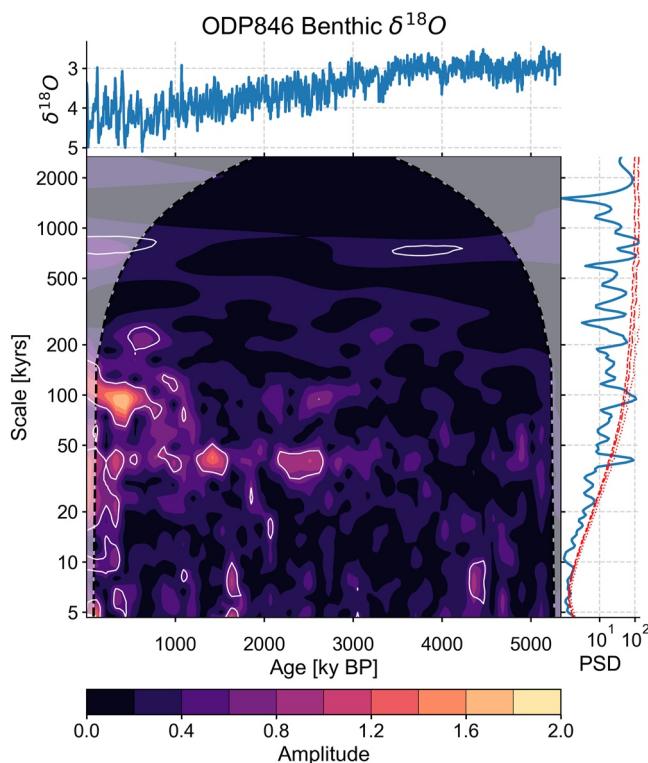


Figure 4. Summary of the spectral and wavelet analysis performed on the benthic $\delta^{18}\text{O}$ record of Ocean Drilling Program (ODP) Site 846 (Mix et al., 1995; Shackleton et al., 1995). Both analyses were performed using the Weighted Wavelet Z-Transform (Foster, 1996; Kirchner & Neal, 2013) method. The record displays significant periodicities in the 40 and 100 kyr bands with a drop in power in the 100 kyr band at the mid-Pleistocene transition.

SST record, the age uncertainty precludes any meaningful interpretation of specific peaks in power for periods shorter than 40–50 kyr.

Finally, we use `Pyleoclim` to perform wavelet coherence analysis (Grinsted et al., 2004) between the SST record from ODP Site 846 (Figure 6) and insolation at 5°S calculated using the `climlab` package (Rose, 2018). We limit the analysis to the first 3 million years of the record, when significant periodicities were apparent in the scalogram. The `wavelet_coherence` method returns a `Coherence` object, which contains the cross-wavelet transform (XWT) and the wavelet transform coherence (WTC). XWT informs about areas where there is high common power between the two series. The analysis reveals high common power in the precession band (23 kyr) but the phase angles are irregular. This is not surprising given the spectral analysis on the age ensemble, which shows large effects of age uncertainty at 20 kyr scales (compared to 40–100 kyr). Even if there was a regular behavior, the age uncertainty prevents us from capturing it in the analysis. WTC shows areas of common behavior between the two time series even if there is low power. The analysis reveals coherence in the 23, 40, 100 and 400 kyr bands, consistent with orbital forcing of climate. The phase angles in the two upper bands are also regular and show an in-phase behavior in the eccentricity band (particularly around 1 Ma) and nearly in phase quadrature in the 400 kyr band.

The example illustrates how `Pyleoclim` facilitates the use of sophisticated spectral and wavelet analysis methods to paleoclimate datasets, especially in regards to age uncertainties and irregular sampling. The package also offers a variety of pre-processing steps (i.e., detrending, removal of outliers and, if desired, interpolating schemes in the time domain) to construct workflows and easily assess the effect of each of these steps on the conclusions.

LR04) using the HMM-Match algorithm developed by Lin et al. (2014). HMM-Match is a $\delta^{18}\text{O}$ Bayesian alignment technique based on a hidden Markov model to develop age models and accompanying uncertainties for deep sea cores.

We first analyze the benthic $\delta^{18}\text{O}$ record using both spectral and wavelet analysis appropriate for uneven timeseries. In this example, we use the Lomb-Scargle periodogram for spectral analysis and the Weighted Wavelet Z-Transform (Foster, 1996; Kirchner & Neal, 2013, WWZ) for both spectral and wavelet analysis (Figure 4). In both cases, the significance is assessed against an AR(1) benchmark. Within `Pyleoclim`, we use the same functionalities as presented in Section 2.3. We find that the record displays significant periodicities in the 40 and 100 kyr bands. This result is hardly surprising considering that the age model was obtained through alignment to the orbitally-tuned LR04 record, which strongly oscillates at those frequencies. Furthermore, the scalogram reveals the non-stationary character of these periodicities, with a drop in power in the 100 kyr band at the mid-Pleistocene transition, ca 0.8 Ma (Paillard, 2001).

The sea surface temperature (SST) record (Lawrence et al., 2006) shows similar, albeit less defined, power in the orbital band (Figure 5). Since the age model returns an ensemble of posterior draws (Khider et al., 2017; Lin et al., 2014), we can perform spectral analysis on each ensemble member to assess the robustness of our conclusions. `Pyleoclim` allows to load an age ensemble as a `EnsembleSeries` object, equipped with its own plotting and analysis functions. As illustrated in the companion notebook, we make use of the `plot` method, which shows various traces based on individual realizations of the age model and the `plot_envelope` method, which uses confidence intervals to communicate age uncertainty. The `spectral` method as applied to `EnsembleSeries` computes the periodogram for each age model realization in the ensemble. `Pyleoclim` allows users to plot the resulting ensemble periodograms to assess the robustness of the spectral peaks in face of age uncertainty. In the case of the Site ODP846

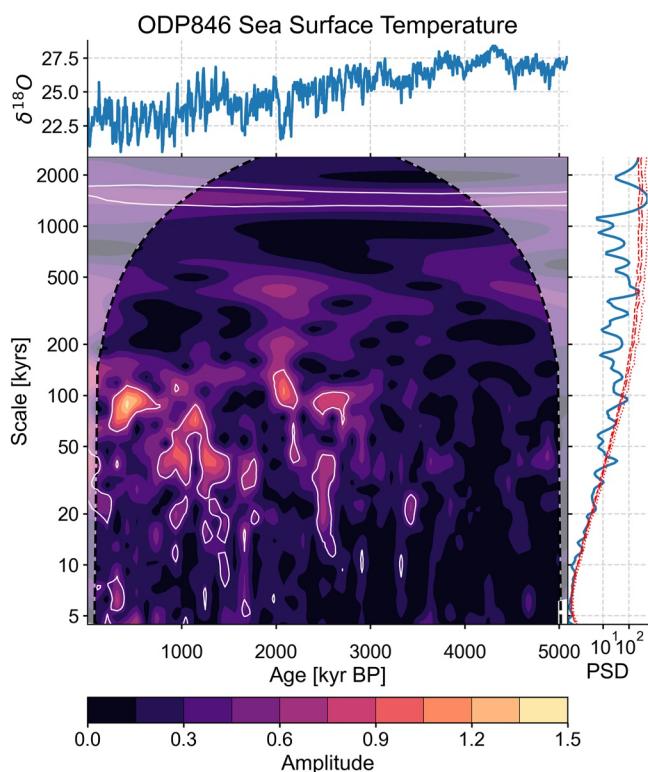


Figure 5. Summary of the spectral and wavelet analysis performed on the sea surface temperature record of Ocean Drilling Program (ODP) Site 846 (Lawrence et al., 2006). Both analyses were performed using the Weighted Wavelet Z-Transform (Foster, 1996; Kirchner & Neal, 2013) method. The record displays significant periodicities in the 40–100 kyr bands.

Extension region of the West Pacific. This interpretation was shown in Hu et al. (2017) to be invalid because of persistence, test multiplicity, and age uncertainties. This notebook repeats the analysis of Hu et al. (2017) leveraging `Pyleoclim` and the updated SST analysis of HadSST4 (Kennedy et al., 2019); in so doing, we extend the original work by showcasing three different methods for assessing the significance of linear correlations: (a) a *t* test with degrees of freedom adjusted for autocorrelation (Dawdy & Matalas, 1964), as used by Hu et al. (2017); (b) the phase-randomization procedure of Ebisuzaki (1997) (dubbed “isospectral” because it preserves a series’ amplitude spectrum) and (c) an “isopersistent” method that gauges the observed correlation against a large sample of AR(1) timeseries with identical persistence parameter as the target series.

In `Pyleoclim`, the `correlation()` method enables tests (a–c), with the default implementing the isospectral method with 1,000 surrogates. The method works between two series, between a series and an ensemble, or between two ensembles, with the same user experience. In the case of ensembles, the object holding the result (`CorrEns`) is equipped with a plotting method (Figure 7) that displays the histogram of correlations, the proportion of correlations with a *p*-value under the test level α (i.e., correlations deemed significant by this test), and the proportion of those that also meet the False Discovery Rate criterion of Benjamini and Hochberg (1995). In this case, we see that only 1 out of the 327 grid points displays a significant correlation with the published Crystal Cave $\delta^{18}\text{O}$ record (Figure 7, top). In addition, the published age model is simply the median of a broader ensemble, which was not made available by the authors. We therefore generated another ensemble of 1,000 draws from the posterior distribution of ages using the Bayesian age model `Bchron` (Haslett & Parnell, 2008) within the `GeoChronR` software (McKay et al., 2021)—the resulting ensemble of possible timeseries is shown in Figure 3 (top). For illustration, we show the result of correlating this ensemble with SST at a single grid point in the Kuroshio Extension region, where McCabe-Glynn et al. (2013) originally reported significant correlations (Figure 7, bottom). While the correlation between HadSST4 SST and the published $\delta^{18}\text{O}$ record was over 0.32, we see that

3.2. Speleothem Correlations With a Temperature Field

Correlation analysis, despite its many shortcomings, remains a centerpiece of empirical analysis in many fields, particularly the paleosciences. Computing correlations is trivial enough; the difficulty lies in properly assessing their significance. Of particular importance are four considerations:

Irregular sampling: comparing two records with different time axes, possibly unevenly spaced, is a challenge to standard methods, which assume concordant observations.

Persistence: persistence violates the standard assumption that the data are independent (which underlies the classical *t*-test of significance implemented in most software packages).

Age uncertainties: age uncertainties affect the location of features in time, so methods need to allow for an ensemble of plausible chronologies (generated, for instance, by a Bayesian age model).

Test multiplicity: test multiplicity, aka the “Look Elsewhere effect,” states that repeatedly performing the same test can result in unacceptably high type I error (accepting correlations as significant, when in fact they are not). This arises for example, when correlating a paleoclimate record with an instrumental field, assessing significance at thousands of grid points at once, or assessing significance within an age ensemble.

Accordingly, `Pyleoclim` facilitates an assessment of correlations that deals with all these challenges, makes the necessary data transformations transparent to the user, and allows for one-line plot commands to visualize the results.

This section and accompanying notebook use `Pyleoclim` to reproduce the study of Hu et al. (2017), particularly the example of their Section 4, which illustrates all the above challenges at once. The example uses the speleothem record of McCabe-Glynn et al. (2013) from Crystal Cave, California, in Sequoia National Park. Based on correlations with the instrumental sea-surface temperature (SST) field of Kaplan et al. (1997), McCabe-Glynn et al. (2013) interpreted their $\delta^{18}\text{O}$ record as a proxy for SST in the Kuroshio

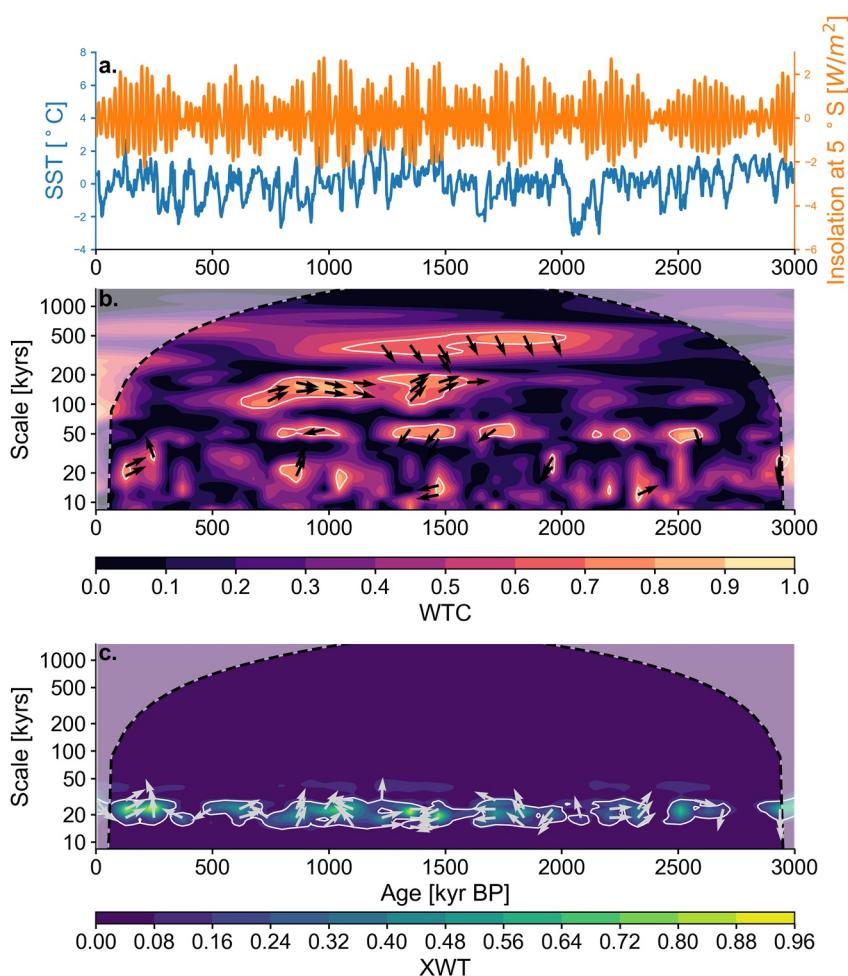


Figure 6. Coherence analysis in Pyleoclim. (a) Sea surface temperature (SST) over the past 3 million years obtained from alkenone paleothermometry at Ocean Drilling Program (ODP) Site 846 (blue) and insolation at 5°S (orange) calculated using the `climlab` package (Rose, 2018). (b) Wavelet transform coherency (WTC) obtained from WWZ between the two timeseries. Contours display WTC, which indicates the degree of resemblance between the signals at each time and scale. The angle of the phase arrows show the relative phasing at each time and scale (e.g., in-phase records are indicated by arrows pointing to the right, out-of-phase to the left, and in phase quadrature up and down). Phase angles are only shown for areas with significant coherence values, assessed against 1,000 random realizations of an AR(1) process. (c) Cross-wavelet transform, with contours displaying areas of high common power, and phase arrows as above. For details on the method, see Grinsted et al. (2004).

the bulk of the histogram is far below this value, with a substantial fraction of ensemble members exhibiting negative correlations. This is a powerful illustration that age uncertainties can go as far as reversing the sign of a correlation, and must be taken into account in this type of exercise. Once all three pitfalls (persistence, multiple comparisons, age uncertainties) are considered, no significant correlation is found.

The example illustrates the risk of relying exclusively on correlations between a paleoclimate record and an instrumental field for interpretation. Historically, this has not been an isolated incident (Hu et al., 2017), so this case study should not be viewed as an indictment of a particular study or group of authors. Rather, it is a reminder of how easy it is to be fooled by spurious correlations, and how easy it is to avoid them with proper methods, such as those made available in Pyleoclim.

3.3. Model-Data Confrontations in the Frequency Domain

The third case study tackles an emerging need in the paleoclimate community: quantitatively comparing paleoclimate observations with transient climate model simulations. In addition to technical challenges (model output

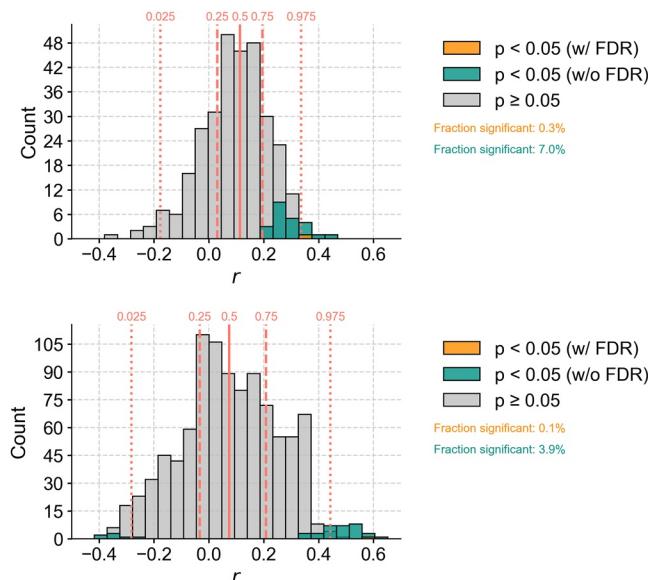


Figure 7. Ensemble correlations in Pyleoclim. Top: histogram of Pearson correlations (r) between the published Crystal Cave record of McCabe-Glynn et al. (2013) with the HadCRUT4 Sea surface temperature (SST) field over the North Pacific (327 grid points). Bottom: histogram of Pearson correlations (r) between the Crystal Cave record of McCabe-Glynn et al. (2013) with a 1000-member Bchron (Haslett & Parnell, 2008) age model ensemble with the HadCRUT4 SST at 32.5°N, 142.5°W in the Kuroshio Extension region. On both panels, “FDR” denotes the False Discovery Rate criterion of Benjamini and Hochberg (1995).

is evenly spaced; observations typically are not), a conceptual difficulty is due to sensitive dependence to initial conditions (chaos): slight changes in initial conditions can result in wildly different climate trajectories despite identical (or even constant) boundary conditions. In paleoclimatology, those initial conditions are unknown, as there typically is no reliable estimate of the 3D state of the climate system at a given point in time. Thus, except when one seeks to compare the expression of external forcings (e.g., Zhu et al., 2020, 2022), it is often sensible to discard phase information altogether and to restrict the comparison to spectral features (peaks, scaling exponents) (Dee et al., 2017; C. L. E. Franzke et al., 2020; Laepple & Huybers, 2014).

This section and accompanying notebook use Pyleoclim to reproduce the comparative study of Zhu et al. (2019), which used several paleoclimate observational datasets to test the ability of a hierarchy of climate models to simulate the continuum of climate variability. Figure 8 emulates part of the original study’s Figure 2, and compares the spectral scaling exponents from 3 transient simulations and five observational datasets, estimated using the WWZ method. The notebook illustrates how few function calls are needed to perform this complex comparison with Pyleoclim, including uncertainty estimates of the scaling exponents.

Zhu et al. (2019) concluded that these models produced simulations of the continuum of climate variability consistent with what can be estimated from paleoclimate observations, provided information about the deglaciation was specified. Most remarkably, these 3 simulations show scaling exponents similar to those observed over the past millennium, despite the models having no knowledge of what are believed to be the leading causes of climate variability over this interval (solar and volcanic forcing). For more details and a discussion of the broader implications of this result, see the original study.

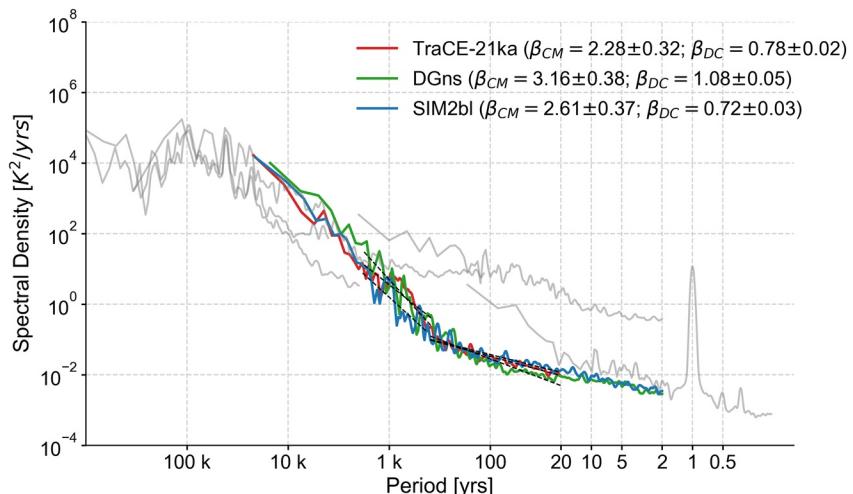


Figure 8. A spectral estimate of the global-average surface temperature variability as portrayed by transient model simulations (TraCE-21ka (Liu et al., 2009), DG_{ns} (Menviel et al., 2011), SIM2bl (Timm & Timmermann, 2007), colors) and observational datasets (gray): HadCRUT4, The Met Office Hadley Centre gridded data set of global historical surface temperature anomalies (Morice et al., 2012); PAGES2k/LMR, the Last Millennium Reanalysis framework (Hakim et al., 2016; Tardif et al., 2019) applied to the PAGES2k data set (PAGES 2k Consortium, 2017); the reconstruction of global average surface temperature of Snyder (2016); ProbStack: A probabilistic Pliocene-Pleistocene stack of benthic $\delta^{18}\text{O}$ (Ahn et al., 2017). The regional data set EPICA Dome C Ice Core 800 Kyr Deuterium Data and Temperature Estimates (Jouzel et al., 2007). β ’s denote the estimated scaling exponents over each appropriate frequency band: β_{CM} is the centennial-to-millennial scale exponent estimated over scales of 400–2,000 years, while β_{DC} is the decadal-to-centennial-scale exponent, estimated over 20–400 years.

4. Conclusion and Outlook

We have presented a new, Python-based toolkit for the analysis and visualization of paleoclimate and paleoceanographic data, whether from observations or models. As of publication, `Pyleoclim` supports a broad array of functionalities to load, process, analyze and visualize timeseries and their relationships to other variables.

Although `Pyleoclim` was primarily designed as a research tool, its extensive documentation makes it useful for established researchers and students alike. At the time of writing, `Pyleoclim` has been used in three virtual workshops (<http://linked.earth/paleoHackathon/>) to build data science capacity within the paleogeosciences communities, and an undergraduate course at the University of Southern California. An in-person training event is planned for the summer of 2023. As part of the PaleoCube grant (<https://medium.com/cyberpaleo/announcing-the-next-linkedearth-chapter-paleocube-790778b6ff0>), many videos (https://www.youtube.com/channel/UCo7yzNTM_4g5H-xyWV5KbA) and notebook tutorials (<https://github.com/LinkedEarth/PaleoBooks>) will be made available to the community to further disseminate and demystify these techniques.

`Pyleoclim` follows an open development model, accessible primarily through its GitHub repository (see data and software availability statement in the acknowledgment section). Interactions with developers and other users are facilitated by a community Slack channel and Discourse forum (<http://linked.earth/community.html>), to ensure knowledge dissemination and align development to the needs of the scientific community. Currently planned extensions include:

Pandas integration: The Pandas library (McKinney, 2010) contains many functionalities for timeseries data that had to be re-implemented for `Pyleoclim`, since the way time is encoded into Pandas is not appropriate for paleoscientific applications: timestamps are represented at nanosecond resolution, so the largest time span that can be represented by a 64-bit integer is limited to approximately 584 years (CE 1677 to 2262), an unacceptably short time for our field. Current work with the Pandas community aims at generalizing this representation to arbitrary intervals, and we expect `Pyleoclim` to soon make direct use of Pandas functionalities (e.g., slicing, aggregating, resampling and many other built-in methods), which will allow for closer integration with climate model output through the popular `xarray` library (Hoyer & Hamman, 2017).

Generalized surrogates: currently, the statistical significance of spectral and wavelet features in `Pyleoclim` can only be assessed against parametric AR(1) surrogates. While those are often reasonable first-order approximations to geophysical timeseries (Ghil et al., 2002), many geophysical phenomena are better emulated by long-range dependent processes (C. Franzke, 2010; Fredriksen & Rypdal, 2017; Samorodnitsky, 2007). We plan for the `SurrogateSeries` class to include more options, such as phase randomization (Ebisuzaki, 1997) (currently only available to correlation and causality methods), fractal and multifractal timeseries generation, and maximum entropy bootstrap (Vinod & de Lacalle, 2009).

Nonlinear Dynamics: Most of the methods currently available in `Pyleoclim` are linear methods. In the near future, we plan to leverage some recent advances in the analysis of nonlinear timeseries via recurrence networks (Zou et al., 2019), convergent cross-mapping (Sugihara et al., 2012) and causal discovery (Runge et al., 2019).

By making sophisticated and rigorous methods available to non-experienced programmers in a few keystrokes, and by providing extensive documentation and training, we expect the package to help streamline the work of many readers of this journal, and contribute to heightened statistical rigor in the analysis of paleoclimate and paleoceanographic data. Furthermore, the package is broadly applicable to any timeseries-based data, and has already been re-used in other fields like astronomy (Peñil et al., 2020)—a trend that we hope spreads to other fields of the geosciences and beyond.

Data Availability Statement

v0.9.1 of `Pyleoclim` (Khider, Emile-Geay, Zhu, James, Landers, et al., 2022) was used to generate all the examples in this study and the supporting Jupyter Notebooks and is available via a GPL-3.0 license and developed openly at https://github.com/LinkedEarth/Pyleoclim_util. v0.4 of the accompanying Jupyter Notebooks (Khider, Emile-Geay, & Zhu, 2022) that provide examples of how `Pyleoclim` can be used for scientific studies is available via an Apache2.0 license and developed openly at <https://github.com/LinkedEarth/PyleoclimPaper>. Tutorials v0.0.1 (Khider, Emile-Geay, James, et al., 2022) are available via an Apache2.0 license and devel-

oped openly at <https://github.com/LinkedEarth/PyleoTutorials> and viewable in the form of a Jupyter Book at <http://linked.earth/PyleoTutorials>.

Acknowledgments

Development of Pyleoclim and associated documentation and training materials has been supported by NSF Grants ICER 1541029, 2126510, AGS2002556, JP Morgan AI Research Awards, and ONR N00014-21-1-2437.

References

Ahn, S., Khider, D., Lisiecki, L. E., & Lawrence, C. E. (2017). A probabilistic Pliocene–Pleistocene stack of benthic $\delta^{18}\text{O}$ using a profile hidden Markov model. *Dynamics and Statistics of the Climate System*, 2(1), dxz002. <https://doi.org/10.1093/climsys/dxz002>

Allen, M. R., & Smith, L. A. (1996). Monte Carlo SSA: Detecting irregular oscillations in the presence of coloured noise. *Journal of Climate*, 9(12), 3373–3404. [https://doi.org/10.1175/1520-0442\(1996\)009<3373:mcssio>2.0.co;2](https://doi.org/10.1175/1520-0442(1996)009<3373:mcssio>2.0.co;2)

Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B*, 57(1), 289–300. <https://doi.org/10.1111/j.2517-6161.1995.tb02031.x>

Dawdy, D., & Matalas, N. (1964). *Statistical and probability analysis of hydrologic data, part III: Analysis of variance, covariance and time series*. McGraw-Hill.

Dee, S. G., Emile-Geay, J., Evans, M. N., Allam, A., Steig, E. J., & Thompson, D. M. (2015). PRYSM: An open-source framework for PROXY System Modeling, with applications to oxygen-isotope systems. *Journal of Advances in Modeling Earth Systems*, 7(3), 1220–1247. <https://doi.org/10.1002/2015MS000447>

Dee, S. G., Parsons, L. A., Loope, G. R., Overpeck, J. T., Ault, T. R., & Emile-Geay, J. (2017). Improved spectral comparisons of paleoclimate models and observations via proxy system modeling: Implications for multi-decadal variability. *Earth and Planetary Science Letters*, 476, 34–46. <https://doi.org/10.1016/j.epsl.2017.07.036>

Ebisuzaki, W. (1997). A method to estimate the statistical significance of a correlation when the data are serially correlated. *Journal of Climate*, 10(9), 2147–2153. do [https://doi.org/10.1175/1520-0442\(1997\)010<2147:amets>2.0.co;2](https://doi.org/10.1175/1520-0442(1997)010<2147:amets>2.0.co;2)

Elson, P., de Andrade, E. S., Lucas, G., May, R., Hattersley, R., Campbell, E., et al. (2022). Scitoools/cartopy: V0.20.3 [Software]. Zenodo. <https://doi.org/10.5281/zenodo.6775197>

Emile-Geay, J. (2017). Data analysis in the Earth & environmental sciences (Third ed.). *FigShare*. <https://doi.org/10.6084/m9.figshare.1014336>

Emile-Geay, J., Khider, D., & James, A. (2019). PaleoBooks: Doing science with pyleoclim [Software]. Zenodo. <https://doi.org/10.5281/zenodo.5771123>

Foster, G. (1996). Wavelets for period analysis of unevenly sampled time series. *The Astronomical Journal*, 112(4), 1709–1729. <https://doi.org/10.1086/118137>

Franzke, C. (2010). Long-range dependence and climate noise characteristics of Antarctic temperature data. *Journal of Climate*, 23(22), 6074–6081. <https://doi.org/10.1175/2010JCLI3654.1>

Franzke, C. L. E., Barbosa, S., Blender, R., Fredriksen, H.-B., Laepple, T., Lambert, F., et al. (2020). The structure of climate variability across scales. *Reviews of Geophysics*, 58(2), e2019RG000657. <https://doi.org/10.1029/2019rg000657>

Fredriksen, H.-B., & Rypdal, M. (2017). Long-range persistence in global surface temperatures explained by linear multibox energy balance models. *Journal of Climate*, 30(18), 7157–7168. <https://doi.org/10.1175/JCLI-D-16-0877.1>

Ghil, M., Allen, R. M., Dettinger, M. D., Ide, K., Kondrashov, D., Mann, M. E., et al. (2002). Advanced spectral methods for climatic time series. *Reviews of Geophysics*, 40(1), 1003–1052. <https://doi.org/10.1029/2000RG000092>

Goble, C., Cohen-Boulakia, S., Soiland-Reyes, S., Garjio, D., Gil, Y., Crusoe, M. R., et al. (2020). FAIR computational workflows. *FAIR Computational Workflows. Data Intelligence*, 2(1–2), 108–121. https://doi.org/10.1162/dint_a_00033

Greene, C. A., Thirumalai, K., Kearney, K. A., Delgado, J. M., Schwanghart, W., Wolfenbarger, N. S., et al. (2019). The climate data toolbox for MATLAB. *Geochemistry, Geophysics, Geosystems*, 20(7), 3774–3781. <https://doi.org/10.1029/2019gc008392>

Grinsted, A., Moore, J. C., & Jevremović, S. (2004). Application of the cross wavelet transform and wavelet coherence to geophysical time series. *Nonlinear Processes in Geophysics*, 11(5/6), 561–566. <https://doi.org/10.5194/npg-11-561-2004>

Hakim, G. J., Emile-Geay, J., Steig, E. J., Noone, D., Anderson, D. M., Tardif, R., et al. (2016). The last millennium climate reanalysis project: Framework and first results. *Journal of Geophysical Research: Atmospheres*, 121(12), 6745–6764. <https://doi.org/10.1002/2016JD024751>

Hannachi, A., Jolliffe, I. T., & Stephenson, D. B. (2007). Empirical orthogonal functions and related techniques in atmospheric science: A review. *International Journal of Climatology*, 27(9), 1119–1152. <https://doi.org/10.1002/joc.1499>

Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., et al. (2020). Array programming with NumPy. *Nature*, 585(7825), 357–362. <https://doi.org/10.1038/s41586-020-2649-2>

Haslett, J., & Parnell, A. (2008). A simple monotone process with application to radiocarbon-dated depth chronologies. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 57(4), 399–418. <https://doi.org/10.1111/j.1467-9876.2008.00623.x>

Hoyer, S., & Hamman, J. (2017). Xarray: N-D labeled arrays and datasets in Python. *Journal of Open Research Software*, 5(1), 10. <https://doi.org/10.5334/jors.148>

Hu, J., Emile-Geay, J., & Partin, J. (2017). Correlation-based interpretations of paleoclimate data – Where statistics meet past climates. *Earth and Planetary Science Letters*, 459, 362–371. <https://doi.org/10.1016/j.epsl.2016.11.048>

Huang, N. E., Shen, Z., Long, S. R., Wu, M. C., Shih, H. H., Zheng, Q., et al. (1998). The empirical mode decomposition and the hilbert spectrum for nonlinear and non-stationary time series analysis. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 454(1971), 903–995. <https://doi.org/10.1098/rspa.1998.0193>

Hunter, J. D. (2007). Matplotlib: A 2D graphics environment. *Computing in Science & Engineering*, 9(3), 90–95. <https://doi.org/10.1109/MCSE.2007.55>

Jouzel, J., Masson-Delmotte, V., Cattani, O., Dreyfus, G., Falourd, S., Hoffmann, G., et al. (2007). Orbital and millennial Antarctic climate variability over the past 800,000 years. *Science*, 317(5839), 793–796. <https://doi.org/10.1126/science.1141038>

Kaplan, A., Kushnir, Y., Cane, M. A., & Blumenthal, M. B. (1997). Reduced space optimal analysis for historical data sets: 136 years of Atlantic sea surface temperatures. *Journal of Geophysical Research*, 102(C13), 27835–27860. <https://doi.org/10.1029/97jc01734>

Kennedy, J. J., Rayner, N. A., Atkinson, C. P., & Killick, R. E. (2019). An ensemble data set of sea surface temperature change from 1850: The met office Hadley centre HadSST4.0.0.0 data set. *Journal of Geophysical Research: Atmospheres*, 124(14), 7719–7763. <https://doi.org/10.1029/2018jd029867>

Khider, D., Ahn, S., Lisiecki, L. E., Lawrence, C. E., & Kienast, M. (2017). The role of uncertainty in estimating lead/lag relationships in marine sedimentary archives: A case study from the tropical Pacific. *Paleoceanography*, 32(11), 1275–1290. <https://doi.org/10.1002/2016pa003057>

Khider, D., Emile-Geay, J., James, A., Landers, J., & Zhu, F. (2022). PyleoTutorials: A gentle introduction to the pyleoclim package [Software]. Zenodo. <https://doi.org/10.5281/zenodo.6999577>

Khider, D., Emile-Geay, J., & Zhu, F. (2022). Example scientific workflows using Pyleoclim [Software]. Zenodo. <https://doi.org/10.5281/zenodo.7093617>

Khider, D., Emile-Geay, J., Zhu, F., & James, A. (2022). PaleoHack: Building coding capacity in the paleogeosciences [Software]. Zenodo <https://doi.org/10.5281/zenodo.6365841>

Khider, D., Emile-Geay, J., Zhu, F., James, A., Landers, J., Kwan, M., & Athreya, P. (2022). Pyleoclim: A Python package for the analysis and visualization of paleoclimate data [Software]. Zenodo <https://doi.org/10.5281/zenodo.1205661>

Kirchner, J. W., & Neal, C. (2013). Universal fractal scaling in stream chemistry and its implications for solute transport and water quality trend detection. *Proceedings of the National Academy of Sciences of the United States of America*, 110(30), 12213–12218. <https://doi.org/10.1073/pnas.1304328110>

Kluyver, T., Ragan-Kelley, B., Pérez, F., Granger, B., Bussonnier, M., Frederic, J., et al. (2016). Jupyter notebooks—A publishing format for reproducible computational workflows. In F. Loizides & B. Schmidt (Eds.), *Positioning and power in academic publishing: Players, agents and agendas* (pp. 87–90).

Laepple, T., & Huybers, P. (2014). Ocean surface temperature variability: Large model–data differences at decadal and longer periods. *Proceedings of the National Academy of Sciences of the United States of America*, 111(47), 16682–16687. <https://doi.org/10.1073/pnas.1412077111>

Lampricht, A.-L., Garcia, L., Kuzak, M., Martinez, C., Arcila, R., Martin Del Pico, E., et al. (2020). Towards fair principles for research software. *Data Science*, 3(1), 37–59. <https://doi.org/10.3233/ds-190026>

Lawrence, K., Liu, Z., & Herbert, T. (2006). Evolution of the eastern tropical Pacific through plio-pleistocene glaciation. *Science*, 312(5770), 79–83. <https://doi.org/10.1126/science.1120395>

Liang, X. S. (2013). The Liang-Kleeman information flow: Theory and applications. *Entropy*, 15(1), 327–360. <https://doi.org/10.3390/e15010327>

Liang, X. S. (2014). Unraveling the cause-effect relation between time series. *Physical Review E - Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics*, 90(5), 052150. <https://doi.org/10.1103/PhysRevE.90.052150>

Liang, X. S. (2015). Normalizing the causality between time series. *Physical Review E - Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics*, 92(2), 022126. <https://doi.org/10.1103/PhysRevE.92.022126>

Liang, X. S. (2016). Information flow and causality as rigorous notions ab initio. *Physical Review E - Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics*, 94(5), 052201. <https://doi.org/10.1103/PhysRevE.94.052201>

Liang, X. S. (2018). Causation and information flow with respect to relative entropy. *Chaos: An interdisciplinary journal of nonlinear science*, 28(7), 073511. <https://doi.org/10.1063/1.5010253>

Lin, L., Khider, D., Lisiecki, L., & Lawrence, C. (2014). Probabilistic sequence alignment of stratigraphic records. *Paleoceanography*, 29(976–989), 976–989. <https://doi.org/10.1002/2014PA002713>

Lisiecki, L., & Raymo, M. (2005). A Pliocene-Pleistocene stack of 57 globally distributed benthic $\delta^{18}\text{O}$ records. *Paleoceanography*, 20(1), PA1003. <https://doi.org/10.1029/2004PA001071>

Liu, Z., Otto-Btiesner, B. L., He, F., Brady, E. C., Tomas, R., Clark, P. U., et al. (2009). Transient simulation of last deglaciation with a new mechanism for bølling-allerød warming. *Science*, 325(5938), 310–314. <https://doi.org/10.1126/science.1171041>

Lomb, N. (1976). Least-squares frequency analysis of unequally spaced data. *Astrophysics and Space Science*, 39(2), 447–462. <https://doi.org/10.1007/bf00648343>

McCabe-Glynn, S., Johnson, K. R., Strong, C., Berkhammer, M., Sinha, A., Cheng, H., & Edwards, R. L. (2013). Variable North Pacific influence on drought in southwestern North America since AD 854. *Nature Geoscience*, 6(8), 617–621. <https://doi.org/10.1038/NGEO1862>

McKay, N. P., & Emile-Geay, J. (2016). Technical note: The linked paleo data framework – A common tongue for paleoclimatology. *Climate of the Past*, 12(4), 1093–1100. <https://doi.org/10.5194/cp-12-1093-2016>

McKay, N. P., Emile-Geay, J., & Khider, D. (2021). geoChronR – An R package to model, analyze, and visualize age-uncertain data. *Geochronology*, 3(1), 149–169. <https://doi.org/10.5194/gchron-3-149-2021>

McKinney, W. (2010). Data structures for statistical computing in Python. In S. van der Walt & J. Millman (Eds.), *Proceedings of the 9th Python in science conference* (pp. 56–61). <https://doi.org/10.25080/Majora-92bf1922-004>

Menviel, L., Timmermann, A., Timm, O. E., & Mouchet, A. (2011). Deconstructing the last glacial termination: The role of millennial and orbital-scale forcings. *Quaternary Science Reviews*, 30(9), 1155–1172. <https://doi.org/10.1016/j.quascirev.2011.02.005>

Millman, K., & Brett, M. (2007). Analysis of functional magnetic resonance imaging in Python. *Computing in Science and Engineering*, 9(3), 52–55. <https://doi.org/10.1109/mcse.2007.46>

Mix, A. C., Le, J., & Shackleton, N. J. (1995). Benthic foraminiferal stable isotope stratigraphy from site 846: 0–1.8 Ma. *Proceedings of the Ocean Drilling Program. Scientific results*, 138, 839–847.

Morice, C. P., Kennedy, J. J., Rayner, N. A., & Jones, P. D. (2012). Quantifying uncertainties in global and regional temperature change using an ensemble of observational estimates: The HadCRUT4 data set. *Journal of Geophysical Research*, 117(D8), D08101. <https://doi.org/10.1029/2011JD017187>

National Academies of Sciences, Engineering, and Medicine. (2021). In R. Silvern & A. Skrivanek (Eds.), *Identifying new community-driven science themes for NSF's support of paleoclimate research: Proceedings of a workshop*. The National Academies Press. <https://doi.org/10.17226/26377>

PAGES 2k Consortium. (2017). A global multiproxy database for temperature reconstructions of the Common Era. *Scientific Data*, 4(1), 170088. <https://doi.org/10.1038/sdata.2017.88>

Paillard, D. (2001). Glacial cycles: Toward a new paradigm. *Reviews of Geophysics*, 39(3), 325–346. <https://doi.org/10.1029/2000RG000091>

Peñil, P., Domínguez, A., Buson, S., Ajello, M., Otero-Santos, J., Barrio, J. A., et al. (2020). Systematic search for γ -ray periodicity in active galactic nuclei detected by the fermi large area telescope. *The Astrophysical Journal*, 896(2), 134. <https://doi.org/10.3847/1538-4357/ab910d>

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12(85), 2825–2830. Retrieved from <http://jmlr.org/papers/v12/pedregosa11a.html>

Perkel, J. M. (2015). Programming: Pick up Python. *Nature*, 518(7537), 125–126. <https://doi.org/10.1038/518125a>

Predybaylo, E., Torrence, C., & Compo, G. (2022). Python wavelet software [software]. ATOC. Retrieved from <http://atoc.colorado.edu/research/wavelets/>

Rose, B. (2018). Climlab: A python toolkit for interactive, process-oriented climate modeling. *Journal of Open Source Software*, 3(24), 659. <https://doi.org/10.21105/joss.00659>

Runge, J., Nowack, P., Kretschmer, M., Flaxman, S., & Sejdinovic, D. (2019). Detecting and quantifying causal associations in large nonlinear time series datasets. *Science Advances*, 5(11), eaau4996. <https://doi.org/10.1126/sciadv.aau4996>

Samorodnitsky, G. (2007). Long range dependence. *Foundations and Trends in Stochastic Systems*, 1(3), 163–257. <https://doi.org/10.1561/0900000004>

Savitzky, A., & Golay, M. J. (1964). Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry*, 36(8), 1627–1639. <https://doi.org/10.1021/ac60214a047>

Scargle, J. (1982). Studies in astronomical time series analysis. ii. Statistical aspects of spectral analysis of unevenly spaced data. *The Astrophysical Journal*, 263(2), 835–853. <https://doi.org/10.1086/160554>

Scargle, J. (1989). Studies in astronomical time series analysis. III. Fourier transforms, autocorrelation functions, and cross-correlation functions of unevenly-spaced data. *The Astrophysical Journal*, 343(2), 874–887. <https://doi.org/10.1086/167757>

Schoellhamer, D. H. (2001). Singular spectrum analysis for time series with missing data. *Geophysical Research Letters*, 28(16), 3187–3190. <https://doi.org/10.1029/2000GL012698>

Schulz, M., & Mudelsee, M. (2002). Redfit: Estimating red-noise spectra directly from unevenly spaced paleoclimatic time series. *Computers & Geosciences*, 28(3), 421–426. [https://doi.org/10.1016/s0098-3004\(01\)00044-9](https://doi.org/10.1016/s0098-3004(01)00044-9)

Schulz, M., & Stattegger, K. (1997). Spectrum: Spectral analysis of unevenly spaced time series. *Computers & Geosciences*, 23(9), 929–945. [https://doi.org/10.1016/s0098-3004\(97\)00087-3](https://doi.org/10.1016/s0098-3004(97)00087-3)

Seabold, S., & Perktold, J. (2010). Statsmodels: Econometric and statistical modeling with python. In *9th python in science conference*.

Shackleton, N. J., Hall, M., & Pate, D. (1995). Pliocene stable isotope stratigraphy of ODP site 846. *Proceedings of the Ocean Drilling Program, Scientific Results*, 138, 337–356.

Snyder, C. W. (2016). Evolution of global temperature over the past two million years. *Nature*, 538(7624), 226–228. <https://doi.org/10.1038/nature19798>

Sugihara, G., May, R., Ye, H., Hsieh, C.-h., Deyle, E., Fogarty, M., & Munch, S. (2012). Detecting causality in complex ecosystems. *Science*, 338(6106), 496–500. <https://doi.org/10.1126/science.1227079>

Tardif, R., Hakim, G. J., Perkins, W. A., Horlick, K. A., Erb, M. P., Emile-Geay, J., et al. (2019). Last millennium reanalysis with an expanded proxy database and seasonal proxy modeling. *Climate of the Past*, 15(4), 1251–1273. <https://doi.org/10.5194/cp-15-1251-2019>

Thomson, D. J. (1982). Spectrum estimation and harmonic analysis. *Proceedings of the IEEE*, 70(9), 1055–1096. <https://doi.org/10.1109/proc.1982.12433>

Timm, O., & Timmermann, A. (2007). Simulation of the last 21,000 years using accelerated transient boundary conditions. *Journal of Climate*, 20(17), 4377–4401. <https://doi.org/10.1175/JCLI4237.1>

Torrence, C., & Compo, G. (1998). A practical guide to wavelet analysis. *Bulletin of the American Meteorological Society*, 79(1), 61–78. [https://doi.org/10.1175/1520-0477\(1998\)079<0061:apgtwa>2.0.co;2](https://doi.org/10.1175/1520-0477(1998)079<0061:apgtwa>2.0.co;2)

VanderPlas, J. T. (2018). Understanding the lomb–scargle periodogram. *The Astrophysical Journal - Supplement Series*, 236(1), 16. <https://doi.org/10.3847/1538-4365/aab766>

Vautard, R., & Ghil, M. (1989). Singular spectrum analysis in nonlinear dynamics, with applications to paleoclimatic time series. *Physica D*, 35(3), 395–424. [https://doi.org/10.1016/0167-2789\(89\)90077-8](https://doi.org/10.1016/0167-2789(89)90077-8)

Vautard, R., Yiou, P., & Ghil, M. (1992). Singular-spectrum analysis: A toolkit for short, noisy chaotic signals. *Physica D: Nonlinear Phenomena*, 58(1), 95–126. [https://doi.org/10.1016/0167-2789\(92\)90103-T](https://doi.org/10.1016/0167-2789(92)90103-T)

Vinod, H. D., & de Lacalle, J. L. (2009). Maximum entropy bootstrap for time series: The meboot r package. *Journal of Statistical Software*, 29(5), 1–19. <http://www.jstatsoft.org/v29/i05>

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., et al. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods*, 17(3), 261–272. <https://doi.org/10.1038/s41592-019-0686-2>

Waskom, M. L. (2021). Seaborn: Statistical data visualization. *Journal of Open Source Software*, 6(60), 3021. <https://doi.org/10.21105/joss.03021>

Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., Appleton, G., Axton, M., Baak, A., et al. (2016). The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data*, 3(1), 160018. <https://doi.org/10.1038/sdata.2016.18>

Zhu, F., Emile-Geay, J., Anchukaitis, K. J., Hakim, G. J., Wittenberg, A. T., Morales, M. S., et al. (2022). A re-appraisal of the ENSO response to volcanism with paleoclimate data assimilation. *Nature Communications*, 13(1), 747. <https://doi.org/10.1038/s41467-022-28210-1>

Zhu, F., Emile-Geay, J., Hakim, G. J., King, J., & Anchukaitis, K. J. (2020). Resolving the differences in the simulated and reconstructed temperature response to volcanism. *Geophysical Research Letters*, 47(8), e2019GL086908. <https://doi.org/10.1029/2019GL086908>

Zhu, F., Emile-Geay, J., McKay, N. P., Hakim, G. J., Khider, D., Ault, T. R., et al. (2019). Climate models can correctly simulate the continuum of global-average temperature variability. *Proceedings of the National Academy of Sciences*, 116(18), 8728–8733. <https://doi.org/10.1073/pnas.1809959116>

Zou, Y., Donner, R. V., Marwan, N., Donges, J. F., & Kurths, J. (2019). Complex network approaches to nonlinear time series analysis. *Physics Reports*, 787, 1–97. <https://doi.org/10.1016/j.physrep.2018.10.005>