

# cfr (v2024.1.26): a Python package for climate field reconstruction

Feng Zhu<sup>1</sup>, Julien Emile-Geay<sup>2</sup>, Gregory J. Hakim<sup>3</sup>, Dominique Guillot<sup>4</sup>, Deborah Khider<sup>5</sup>, Robert Tardif<sup>3</sup>, and Walter A. Perkins<sup>6</sup>

<sup>1</sup>Climate and Global Dynamics Laboratory, National Center for Atmospheric Research, Boulder, CO, USA

<sup>2</sup>Department of Earth Sciences, University of Southern California, Los Angeles, CA USA

<sup>3</sup>Department of Atmospheric Sciences, University of Washington, Seattle, WA USA

<sup>5</sup>University of Southern California, Information Sciences Institute, Marina Del Rey, CA

<sup>4</sup>Department of Mathematical Sciences, University of Delaware, Newark, DE, USA

<sup>6</sup>Allen Institute for Artificial Intelligence, Seattle, WA USA

**Correspondence:** Feng Zhu (fengzhu@ucar.edu)

**Abstract.** Climate field reconstruction (CFR) refers to the estimation of spatiotemporal climate fields (such as surface temperature) from a collection of pointwise paleoclimate proxy datasets. Such reconstructions can provide rich information on climate dynamics and provide an out-of-sample validation of climate models. However, most CFR workflows are complex and time-consuming, as they involve: (i) preprocessing of the proxy records, climate model simulations, and instrumental observations, (ii) application of one or more statistical methods, and (iii) analysis and visualization of the reconstruction results. Historically, this process has lacked transparency and accessibility, limiting reproducibility and experimentation by non-specialists. This article presents an open-source and object-oriented Python package called `cfr` that aims to make CFR workflows easy to understand and conduct, saving climatologists from technical details and facilitating efficient and reproducible research. `cfr` provides user-friendly utilities for common CFR tasks such as proxy and climate data analysis and visualization, proxy system modeling, and modularized workflows for multiple reconstruction methods, enabling methodological intercomparisons within the same framework. The package is supported with an extensive documentation of the application programming interface (API) and a growing number of tutorial notebooks illustrating its usage. As an example, we present two `cfr`-driven reconstruction experiments using the PAGES 2k temperature database: applying the last millennium reanalysis (LMR) paleoclimate data assimilation (PDA) framework and the Graphical Expectation-Maximization (GraphEM) algorithm, respectively.

## 1 Introduction

Paleoclimate reconstructions provide critical context for recent changes and out-of-sample validation of climate models (Masson-Delmotte et al., 2013). Site-based reconstructions have been used to infer climate change with the assumption that site-based paleoclimate records (i.e., proxies) can represent the regional or even global climate variations to some extent (e.g., Shakun et al., 2012; Marcott et al., 2013), yet such approaches may lead to conclusions biased to specific sites (Bova et al., 2021; Osman et al., 2021). Climate field reconstruction (CFR), in contrast, aims to optimally combine the available proxy data and infer multivariate climate variability over an entire domain of interest, hence alleviating the biases in site-based reconstruction methods. It has become the emerging approach to studying spatiotemporal climate history and model-data comparisons over

the past 2,000 years (e.g. Tingley et al., 2012; Neukom et al., 2019a; Tierney et al., 2020; Zhu et al., 2020; Osman et al., 2021; King et al., 2021; Zhu et al., 2022).

25 Existing CFR methods can be classified into three main categories. The first category is based on regression models, which fit climate variables (e.g., temperature, pressure, precipitation, and so on) onto proxy observations using some variant of least squares (e.g., Mann et al., 1998, 1999; Evans et al., 2002; Cook et al., 2004, 2010; Tingley et al., 2012; Smerdon and Pollack, 2016). A more formal approach is based on Bayesian hierarchical models, which explicitly model the various levels of the relationship linking climate processes to proxy observations, and invert the proxy-climate relation via Bayes' theorem (e.g., 30 Tingley and Huybers, 2010a, b, 2013; Christiansen and Ljungqvist, 2017; Shi et al., 2017). Both approaches rely on the covariance structure between proxies and climate variables that is estimated from observational data over a calibration period. The third is based on paleoclimate data assimilation (PDA, e.g., Dirren and Hakim, 2005; Goosse et al., 2006b; Ridgwell et al., 2007; Steiger and Hakim, 2016; Hakim et al., 2016; Tardif et al., 2019), which follows a similar idea to the previous approaches, but with a critical difference that the proxy-climate covariance matrix is estimated from climate model simulations 35 instead of observational data, as such spatial, multivariate relationships are available and subject to dynamical constraints. As a result, PDA methods can, in principle, estimate any field simulated in the model prior, though the reconstruction quality will be a function of how strongly those variables can be constrained by paleo observations.

Since some CFR methods have been shown to depend sensitively on the input data (Wang et al., 2015), it is imperative to apply more than one method to the same problem to establish a result's robustness. This has heretofore been difficult, as 40 most CFR workflows are complex, possibly involving the selection and processing of proxy records, the processing of gridded climate model simulation and instrumental observation data, the calibration of proxy system models (PSMs, Evans et al., 2013), the setup and execution of the specific reconstruction algorithms, as well as the validation and visualization of the reconstruction results. Each of these steps can lead to a long decision tree (Bürger et al., 2006; Büntgen et al., 2021), complicating comparisons. Furthermore, these steps require a comprehensive knowledge of proxies, data analysis and modeling, 45 reconstruction methods, as well as scientific programming and visualization, and can be obscure and error-prone for climatologists not familiar with all of them. Here we present a Python package called `cfr` that is designed to make CFR workflows easy to understand and perform, with the aim of improving the efficiency and reproducibility of CFR-related research.

The paper is organized as follows: Sect. 2 presents the design philosophy of the package. Sect. 3 provides a primer for the two classes of reconstruction methods presently included in the package. Sect. 4 introduces the architecture and major 50 modules. Sect. 5 and 6 describe the reconstruction workflows of two methods: paleoclimate data assimilation and graphical expectation-maximization, respectively. A summary and an outlook are provided in Sect. 7.

## 2 The `cfr` design philosophy

`cfr` aims to provide intuitive and effective workflows for climate field reconstruction with the following features:

**Reproducible Computational Narratives.** Although a traditional script-based style is also supported, `cfr` is natively de- 55 signed to be used in the context of reproducible computational narratives known as Jupyter notebooks (Kluyver et al., 2016),

which provides an interactive programming laboratory for data analysis and visualization, and has become the new standard for analysis-driven scientific research.

**Intuitive.** `cfr` is coded in the object-oriented programming (OOP) style, which provides intuitive ways to manipulate data objects. For instance, each proxy record object (`ProxyRecord`) supports a collection of methods for analysis and visualization; proxy record objects can be added together (by the operator “+”) to form a proxy database object (`ProxyDatabase`) that supports another collection of methods specific to a proxy database. So-called “method cascading” is also supported to allow for smooth processing of the data objects in one line of code combining multiple processing steps.

**Flexible.** `cfr` provides multiple levels of modularization. It supports object-specific workflows, which provide in-depth operations of the data objects, as well as workflows specific to reconstruction tasks, enabling macroscopic manipulations of the reconstruction-related tasks.

**Community-based.** `cfr` is built upon community efforts on (paleoclimate) data analysis, modeling, and visualization, including but not limited to NumPy (van der Walt et al., 2011), SciPy (Virtanen et al., 2020), Pandas (McKinney, 2010), Xarray (Hoyer and Hamman, 2017), Matplotlib (Hunter, 2007), Cartopy (Met Office, 2010), Seaborn (Waskom, 2021), Plotly (Plotly Technologies Inc., 2015), Statsmodels (Seabold and Perktold, 2010), Pyleoclim (Khider et al., 2022), and PRYSM (Dee et al., 2015). This makes the codebase of `cfr` not only concise and efficient, but also ready for grassroots open development.

**User-friendly.** `cfr` is designed to be easy to install and use. It is supported with an extensive documentation on the installation and the essential application programming interface (API), as well as a growing number of tutorial notebooks illustrating its usage. In addition, the commonly used proxy databases and gridded climate data for CFR applications can be remotely loaded from the cloud with `cfr`’s data fetching functionality.

## 75 3 Reconstruction Methods

At the moment, `cfr` supports two classes of reconstruction methods, though it is designed to accommodate many more.

### 3.1 Offline Paleoclimate Data Assimilation (PDA)

In recent years, paleoclimate data assimilation (PDA) has provided a novel way to reconstruct past climate variations (Jones and Widmann, 2004; Goosse et al., 2006a; Gebhardt et al., 2008; Widmann et al., 2010; Goosse et al., 2010; Annan and Hargreaves, 2013; Steiger et al., 2014; Hakim et al., 2016; Franke et al., 2017; Acevedo et al., 2017; Steiger et al., 2018; Shi et al., 2019; Tierney et al., 2020; Osman et al., 2021; King et al., 2021; Lyu et al., 2021; Zhu et al., 2022; Valler et al., 2022; Annan et al., 2022; Fang et al., 2022). PDA proceeds by drawing from a prior distribution of climate states, which it updates by comparison with observations (Wikle and Berliner, 2007). The prior comes from physics-based model simulations, typically from general circulation models experiments covering past intervals (e.g. Otto-Bliesner et al., 2016).

85 The observations, in this case, are the values of climate proxies, which indirectly sense the climate field(s) of interest (e.g. surface temperature, precipitation, wind speed). The link between latent climate states and paleo proxy observations is provided by observation operators (one for each proxy type), which in this case are called proxy system models (PSMs, Evans

et al., 2013). PSMs are low-order representations of the processes that translate climate conditions to the physical or chemical observations made on various proxy archives, such as tree rings, corals, ice cores, lake and marine sediments, or cave deposits (speleothems).

`cfr` implements a version of a data assimilation algorithm known as an offline ensemble Kalman filter (EnKF), popularized in the paleoclimate context by Steiger et al. (2013); Hakim et al. (2016); Tardif et al. (2019). The code for the EnKF solver is derived from the Last Millennium Reanalysis (LMR) codebase (<https://github.com/modons/LMR>) and was streamlined with utilities from the `cfr` package. The output of this algorithm is a time-evolving distribution (the “posterior”) quantifying the probability of particular climate states over time.

PDA can reconstruct as many climate fields as are present in the prior, though they will not be equally well constrained by the observations. Prior values are usually obtained by drawing  $n$  (100 to 200) samples at random from a model simulation at the outset. This precludes sampling variations, ensuring that all variability in the posterior is driven by the observations.

Uncertainty quantification is carried out by random sampling of a subset of the proxy set (typically 75%) for separate training and validation. This process is typically repeated MC (20-50) times, yielding so many Monte-Carlo “iterations” of the reconstruction, each of which contains a unique set of  $n$  samples from the prior. The posterior is therefore composed of contains  $n \times \text{MC}$  trajectories, typically numbering in the thousands. For scalar indices like the global mean surface temperature or the North Atlantic Oscillation index, the entire ensemble is exported. However, for climate fields, this added dimension would yield unacceptably large files. As a compromise, for each field variable, `cfr` by default exports the ensemble mean of the  $n$  reconstructions during each Monte-Carlo iteration, and the final output thus contains MC ensemble members. In case the full ensemble is needed, `cfr` also provides a switch to export the entire  $n \times \text{MC}$  members of the reconstructed fields.

The framework treats each target year for reconstruction in a standalone manner, i.e. the reconstruction of a certain year will not affect that of the others. This design is a key feature of offline DA. It alleviates the issue of non-uniform temporal data availability and facilitates the parallel reconstruction of multiple years.

### 110 **3.2 Graphical Expectation Maximization (GraphEM)**

GraphEM (Guillot et al., 2015, hereafter G15) is a variant of the popular Regularized Expectation-Maximization (RegEM) algorithm introduced by Schneider (2001), and popularized in the paleoclimate reconstruction context by Michael Mann and co-authors (Rutherford et al., 2005; Mann et al., 2007; Mann et al., 2008, 2009; Emile-Geay et al., 2013a, b). Like RegEM, GraphEM mitigates the problems posed to the traditional EM algorithm (Dempster et al., 1977) by the presence of rank-deficient covariance matrices, due the “large  $p$ , small  $n$ ” problem: climate fields are observed over large grids (e.g. a  $5^\circ \times 5^\circ$  global grid contains nearly 2500 grid points) but a comparatively short time (say 170 annual samples for the period 1851-2020). GraphEM addresses this issue by exploiting the conditional independence relations inherent to a climate field (Vaccaro et al., 2021): namely, surface temperature at a given gridpoint is conditionally independent of the vast majority of the rest of the grid, except for a handful of neighbors, typically a fraction of a percent of the total grid size. Effectively, this reduces the dimensionality of the problem, to the point that the estimation is well-posed, and the various submatrices of the covariance matrix  $\Sigma$  are invertible and well estimated.

The degree of regularization is determined by the density of non-zero entries in the inverse covariance matrix  $\Omega$ , which can be used to construct a graph adjacency matrix (Lauritzen, 1996). The more zero entries in this matrix, the sparser the graph, and the more damped the estimation. The fuller the graph, the more connections are captured between distant locales, but the more potential there is for a non-invertible covariance matrix, resulting in unphysical values of the reconstructed field. Therefore, the quality of the GraphEM estimation hinges on an appropriate choice of graph, which is a compromise between the two poles described above. In practice, it is helpful to split  $\Omega$  into three parts:

$\Omega_{FF}$  stands for “field-field”, and encodes covariances intrinsic to a given field, for instance covariance between temperature at different grid points.

$\Omega_{FP}$  stands for “field-proxy”, and encodes the way the proxies covary with the target field (here, temperature).

$\Omega_{PP}$  stands for “proxy-proxy”, and encodes how the values of one proxy may be predictive of another.

This last part is the identity matrix, because a proxy cannot provide any more information about another proxy than the underlying climate field can (Guillot et al., 2015). Therefore, the only parts of  $\Omega$  that need be specified are  $\Omega_{FF}$  and  $\Omega_{FP}$ . `cfr` supports two ways to do so:

**Neighborhood graphs** are the simplest method, identifying as neighbors all gridpoints that lie within a certain radius (say, 1000 km) of a given location. The smaller the radius, the sparser the graph; the larger the radius, the fuller the graph. Neighborhood graphs take no time to compute, but their structure is fairly rigid, as its sole dependence on distance means that it cannot exploit anisotropic features like land/ocean boundaries, orography, or teleconnection patterns.

**Graphical LASSO** (GLASSO) is an  $\ell_1$  penalized likelihood method introduced by Friedman et al. (2008) in the context of high-dimensional covariance estimation. A major advantage of GLASSO is that it can extract non-isotropic dependencies from a climate field (Vaccaro et al., 2021). In `cfr`, the level of regularization is controlled by two sparsity parameters, which explicitly target the proportion of non-zero entries in  $\Omega$  (see G15 for details): `sp_FF` controls the sparsity of submatrix  $\Omega_{FF}$ , and `sp_FP` controls the sparsity of submatrix  $\Omega_{FP}$ . The smaller the sparsity parameters, the sparser the graph, and the more damped the estimate.

G15 found that for suitable choices of the sparsity parameters, the GLASSO approach yielded much better estimates of the temperature field than the neighborhood graph method. However, the GLASSO approach has two main challenges: (1) the graph optimization can be computationally intensive and (2) it requires a complete data matrix to operate. In the case of paleoclimate reconstructions, this means that proxy series and instrumental observations of the climate field of interest must not contain any missing values over the entire calibration period (e.g. 1850-2000).

To address the first challenge, `cfr` uses a greedy algorithm to find the optimal `sp_FP` and `sp_FF`, as proposed by G15. To address the second challenge, we use GraphEM with neighborhood graphs to obtain a first guess for the climate field, use it to run GLASSO and obtain a more flexible graph, which is then used within GraphEM on the original (incomplete) data matrix (Section 6). Another advantage of this method is that it provides GraphEM (an iterative method) with a reasonable first

guess for the parameters of the distribution  $(\mu_0, \Sigma_0)$ , which can considerably speed up convergence. We call this approach the  
155 "hybrid" method in `cfr`.

### 3.3 Generalization

The `cfr` framework is designed to be quite general, and can in principle accommodate any CFR method. PDA and GraphEM  
are only two of the methods used, for instance, in Neukom et al. (2019b), which are themselves a subset of all possible CFR  
methods. The two options included at this stage, based on completely different assumptions and methodologies, provide a way  
160 to test the method-dependence of a given result. This is critical, as some CFR methods have been shown to depend sensitively  
on the input data (Wang et al., 2015). More methods will be added into the package in subsequent versions. We also welcome  
of contributions from the community for the inclusion of more other Python-based CFR methods. A contributing guide can be  
accessed at: <https://fzhu2e.github.io/cfr> (Zhu et al., 2024).

## 4 The `cfr` architecture

165 To make CFR workflows intuitive and flexible, `cfr` takes an object-oriented programming (OOP) approach and defines  
Python classes describing and organizing the essential data objects along with their connections, such as proxy records, proxy  
databases, gridded climate fields, as well as the PSMs and reconstruction method solvers, under different modules, which are  
described in the following sections.

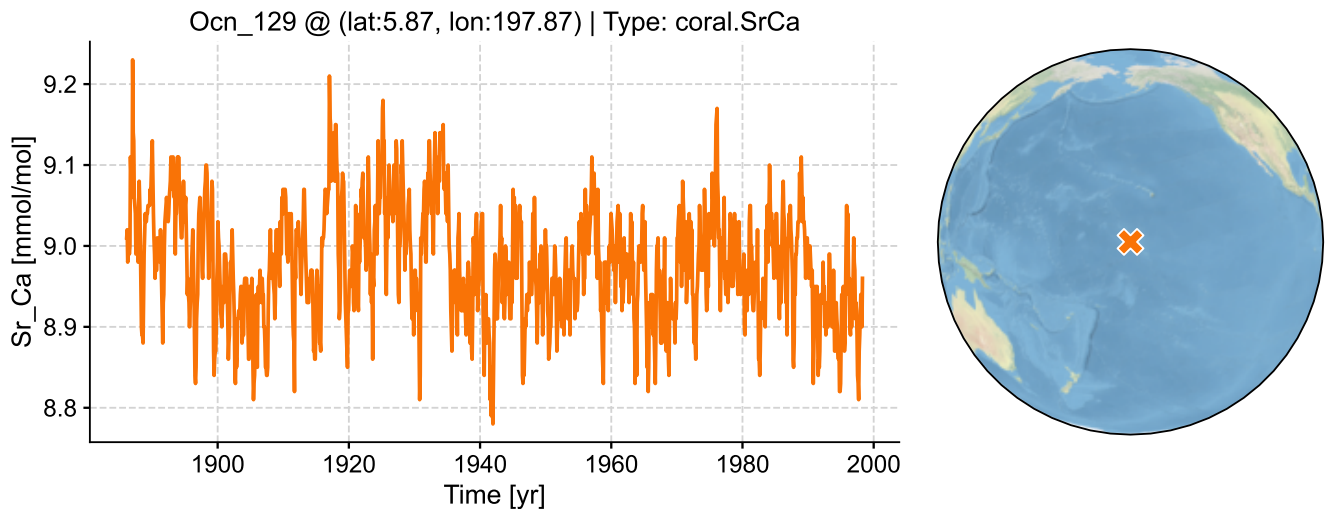
### 4.1 proxy: proxy data processing and visualization

170 The proxy module contains a `ProxyRecord` class and a `ProxyDatabase` class, providing a collection of operation methods  
for proxy data processing and visualization, which is fundamental for paleoclimate data analysis.

The `ProxyRecord` class provides a structure to store, process, and visualize a proxy record. The essential attributes and  
methods are listed in Table 1. For instance, suppose we have a `ProxyRecord` object named `pobj`, then calling `pobj.plot()`  
will visualize the proxy time series along with its location depicted on a map (Fig. 1). Basic metadata stored as its attributes  
175 will be displayed on the plot, including the proxy ID, the location in latitude and longitude, the proxy type, and the proxy  
variable with the unit, if available. There are methods for basic processing: calling `pobj.slice()` will slice the record based  
on the given timespan (while “index slicing” is also supported, please refer to the notebook tutorial: <https://fzhu2e.github.io/cfr/notebooks/proxy-ops.html#Slice-a-ProxyRecord> (Zhu et al., 2024) for the details), calling `pobj.annualize()` will annu-  
annualize/seasonalize the record based on the specified list of months, and calling `pobj.center()` or `pobj.standardize()` will  
180 center or standardize the value axis of the record. There are also methods related to proxy system modeling: the `get_clim()`  
method helps to get the grid point value from a gridded climate field nearest to the record, after which the `get_pseudo()`  
method can be called to generate the pseudoproxy estimate. To support intuitive conversions between data objects, multiple  
`ProxyRecord` objects can be added together with the “+” operator to form a `ProxyDatabase` that we introduce in the next  
section (i.e., `pdb = pobj1 + pobj2 + ...`).

**Table 1.** The essential attributes and methods of `cfr.ProxyRecord`.

Attributes	Description	Methods	Description
<code>pid</code>	The unique proxy ID.	<code>plot()</code>	Plot the record time series.
<code>time</code>	The time axis.	<code>slice()</code>	Slice the record based on a give timespan.
<code>value</code>	The value axis.	<code>+</code>	Combine multiple <code>ProxyRecord</code> 's as a <code>ProxyDatabase</code> .
<code>ptype</code>	The proxy type.	<code>annualize()</code>	Annualize/seasonalize the proxy record.
<code>lat, lon, elev</code>	The latitude, longitude, and elevation.	<code>get_clim()</code>	Get climate data nearest to the record.
<code>seasonality</code>	The seasonality.	<code>get_pseudo()</code>	Generate the pseudoproxy estimate.
<code>dt</code>	The median of time differences.	<code>center()</code>	Center the record against a reference timespan.
<code>tags</code>	A tag set for filtering purposes.	<code>standardize()</code>	Standardize the value axis of the record.



**Figure 1.** A visualization example of a coral Sr/Ca record using `cfr.ProxyRecord.plot()`.

185 The `ProxyDatabase` class provides a structure to organize multiple `ProxyRecord` objects at once. The essential attributes and methods are listed in Table 2. For example, suppose we have a `ProxyDatabase` object named `pdb`, then calling `pdb.plot()` will visualize the proxy database on a static map (Fig. 2 top), along with the count of the records by proxy type if the argument `plot_count=True` is set (Fig. 2 middle). In contrast, calling `pdb.plotly()` will visualize the proxy database on an interactive map that allows query of the location and proxy type of each record by hovering a mouse pointer (Fig. 2 bottom). The

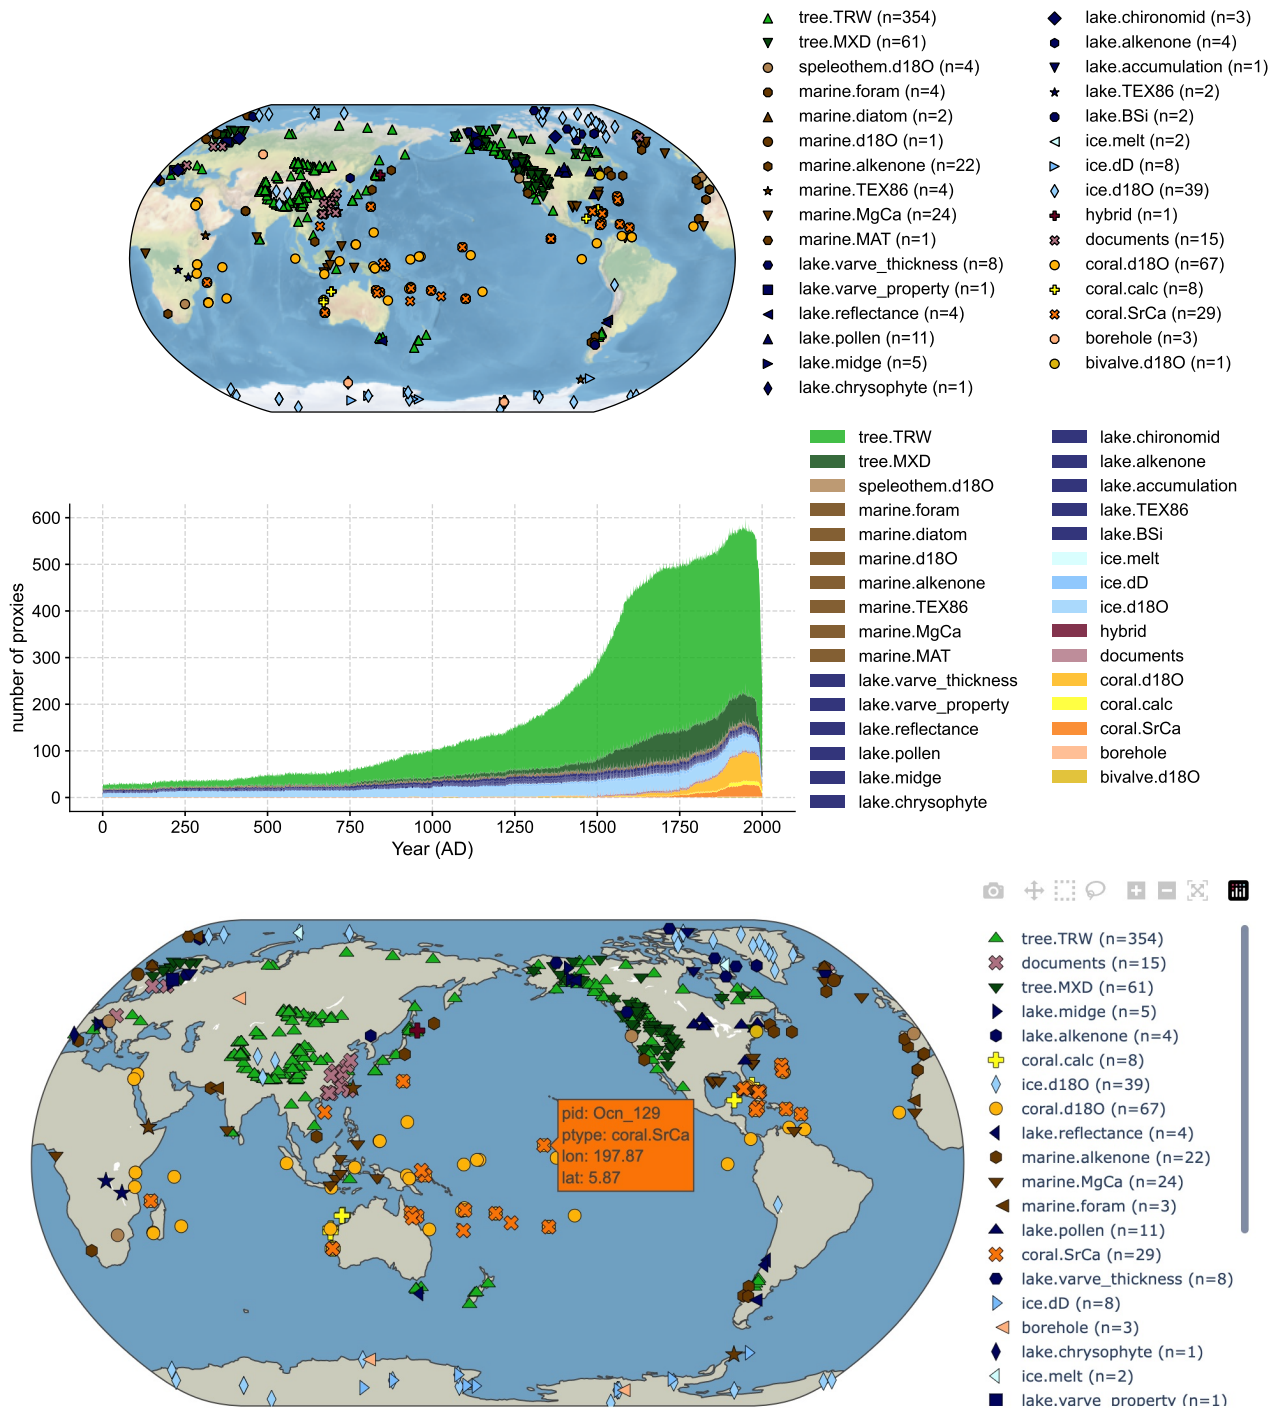
190 `from_df()` and `to_df()` methods allow data input and output in the form of a `pandas.DataFrame`, which is a more common format for tabular data. The `filter()` method offers a handy tool to filter the proxy database in flexible ways, such as by proxy types, latitude/longitude ranges, center location and radius distance, and arbitrary tag combinations. For more details, please refer to the notebook tutorial on this topic: <https://fzhu2e.github.io/cfr/notebooks/pp2k-pdb-filter.html> (Zhu et al., 2024). To

support intuitive conversions between data objects, multiple ProxyRecord objects and/or ProxyDatabase objects can be added together with the “+” operator to form a ProxyDatabase (i.e., `pdb = pdb1 + pobj1 + ...`), and multiple ProxyRecord objects can be removed from a ProxyDatabase with the “-” operator (i.e., `pdb = pdb1 - pobj1 - ...`). ProxyDatabase also comes with `find_duplicates()` and `squeeze_dups()` to locate and remove duplicated records when adding multiple databases together. In addition, with the `fetch` method, the commonly used proxy databases can be conveniently fetched from the cloud. For instance, `pdb = cfr.ProxyDatabase.fetch('PAGES2kv2')` will remotely load the PAGES 2k phase 2 database (PAGES 2k Consortium, 2017), and `pdb = cfr.ProxyDatabase.fetch('pseudoPAGES2k/ppwn_SNRinf_rta')` will fetch and load the base version of the “pseudoPAGES2k” database (Zhu et al., 2023a, b). Its argument can also be an arbitrary URL to a supported file stored in the cloud to support flexible loading of the data. If the `fetch` method is called without an argument, the supported predefined entries will be printed out to give a hint to the users.

**Table 2.** The essential attributes and methods of `cfr.ProxyDatabase`.

Attributes	Description	Methods	Description
<code>records</code>	The dictionary of the records.	<code>plot()</code>	Plot the database on a static map.
<code>pids</code>	The list of proxy IDs.	<code>plotly()</code>	Plot the proxy database on an interactive map.
<code>nrec</code>	The number of the proxy records.	<code>from_df()</code>	Load the database from a <code>pandas.DataFrame</code> .
<code>type_dict</code>	The count of each proxy type.	<code>to_df()</code>	Convert the database to a <code>pandas.DataFrame</code> .
<code>lats, lons</code>	Location information in lists.	<code>filter()</code>	Filter the database by supported conditions.
		<code>+</code>	Combine multiple ProxyRecords’s/ProxyDatabase’s together.
		<code>-</code>	Removing multiple ProxyRecord’s from a ProxyDatabase.
		<code>find_duplicates()</code>	Find duplicated records.
		<code>squeeze_dups()</code>	Remove duplicated records and keep one.
		<code>fetch()</code>	Fetch and load the database from cloud.





**Figure 2.** Visualization examples of the PAGES 2k version 2 multiproxy database (PAGES 2k Consortium, 2017) using `cfr.ProxyDatabase.plot()` and `cfr.ProxyDatabase.plotly()`.


## 4.2 climate: gridded climate data processing and visualization

205 The `climate` module comes with a `ClimateField` class that is essentially an extension of `xarray.DataArray` to better fit CFR applications. It contains a collection of methods for processing and visualizing the gridded climate model simulation and instrumental observation data.

The essential attributes and methods are listed in Table 3. For instance, suppose we have a `ClimateField` object named `fd`, then calling `fd.get_anom()` will return the anomaly field against a reference time period, and calling `fd.get_anom().plot()` (an example of method cascading) will plot the anomaly field on a map (Fig. 3). The spatial resolution and coverage of the field can be altered by calling `fd.regrid()` and `fd.crop()`. Since `ClimateField` is based on `xarray.DataArray`, the slice method applies similarly; however, we augmented the original method to more robustly account for time and calendars in the paleoclimate context. For more details on time slicing, please refer to the notebook tutorial: <https://fzhu2e.github.io/cfr/notebooks/climate-ops.html#Time-slicing> (Zhu et al., 2024). Similar to a `ProxyRecord`, a `ClimateField` with monthly time resolution can be annualized/seasonalized by calling the `annualize()` method with a specified list of months. To support convenient calculation of many popular climate indices such as global mean surface temperature (GMST) and NINO3.4, the `geo_mean()` method calculates the “cos-lat” weighted geospatial mean over a region defined by the latitude and longitude range. To support quick validation of a reconstructed field, the `compare()` method facilitates the comparison against a reference field in terms of a metric such as the gridpoint-wise linear correlation ( $r$ ), coefficient of determination ( $R^2$ ), or coefficient of efficiency ( $CE$ ). Lastly, the `load_nc()` and `to_nc()` methods allow loading and writing gridded datasets in the netCDF format (Rew and Davis, 1990), while the `fetch` method can be used to conveniently load gridded datasets from the cloud. For instance, `fd = cfr.ClimateField.fetch('iCESM_past1000historical/tas')` will fetch and load the iCESM simulated surface temperature field since 850 CE (Brady et al., 2019). Similar to that of `ProxyDatabase`, the argument can also be an arbitrary URL to a netCDF file hosted in the cloud, and calling without an argument will print out supported predefined entries.

**Table 3.** The essential attributes and methods of `cfr.ClimateField`.

Attributes	Description	Methods	Description
<code>da</code>	The data in <code>xarray.DataArray</code> .	<code>plot()</code>	Plot the climate field on a map.
		<code>get_anom()</code>	Calculate the anomaly field against a reference time period.
		<code>regrid()</code>	Regrid the field.
		<code>crop()</code>	Crop the field.
		<code>annualize()</code>	Annualize/seasonalize the field.
		<code>geo_mean()</code>	Calculate the geospatial mean over a region.
		<code>compare()</code>	Compare against a reference field.
		<code>load_nc()</code>	Load the field from a netCDF file.
		<code>to_nc()</code>	Output the field to a netCDF file.
		<code>fetch()</code>	Fetch and load the gridded climate field from cloud.



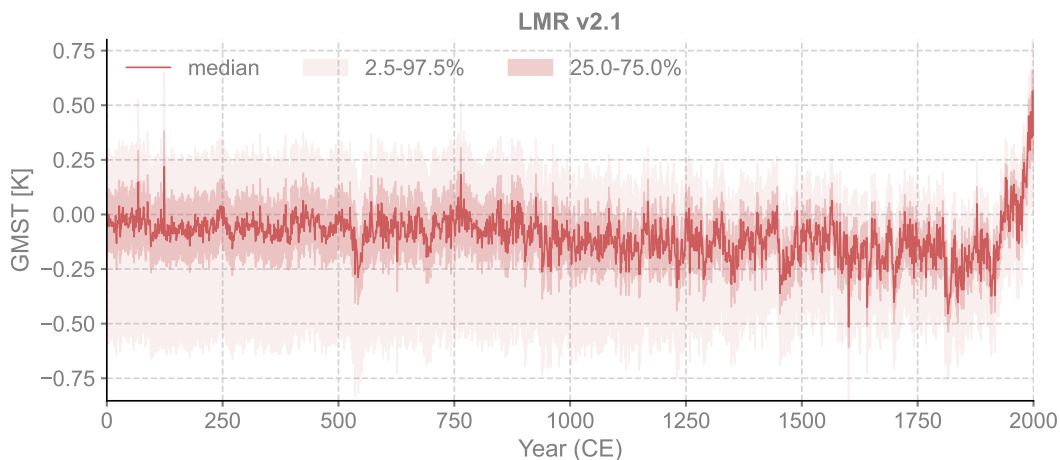
**Figure 3.** A visualization example of an iCESM simulated surface temperature anomaly field using `cfr.ClimateField.plot()`.

### 225 4.3 ts: time series processing and visualization

The `ts` module is for time series processing and visualization in general, and it comes with an `EnsTS` class to handle ensemble time series. In specific, the `EnsTS` class is mainly designed to visualize and validate the reconstructed time series such as GMST and NINO3.4, and its essential attributes and methods are listed in Table 4. The `plot()` method visualizes each member of the ensemble time series, while the `plot_qs()` method plots only the quantile envelopes (Fig. 4). Similar to a `ClimateField`, the `compare()` method is useful for a quick comparison of the ensemble median against a reference time series. The `fetch` method is supported as well. For instance, `bc09 = cfr.ProxyDatabase.fetch('BC09_NIN034')` will remotely load the NINO3.4 estimate by Bunge and Clarke (2009).

**Table 4.** The essential attributes and methods of `cfr.EnsTS`.

Attributes	Description	Methods	Description
<code>time</code>	The time axis.	<code>plot()</code>	Plot each of the ensemble time series.
<code>value</code>	The ensemble value in matrix.	<code>plot_qs()</code>	Plot the quantile envelopes.
<code>nt</code>	The temporal length.	<code>compare()</code>	Compare against a reference time series.
<code>nEns</code>	The ensemble size.	<code>fetch()</code>	Fetch and load the time series data from cloud.



**Figure 4.** A visualization example of the LMR v2.1 ensemble global mean surface temperature (GMST) (Tardif et al., 2019) using `cfr.EnsTS.plot_qs()`.

#### 4.4 psm: proxy system modeling

The `psm` module incorporates classes for multiple popular proxy system models (PSMs, Evans et al., 2013), including a univariate linear regression model (Hakim et al., 2016; Tardif et al., 2019) (`Linear`), a bivariate linear regression model (Hakim et al., 2016; Tardif et al., 2019) (`Bilinear`), a tree-ring width model VS-Lite (Tolwinski-Ward et al., 2011, 2013; Zhu and Tolwinski-Ward, 2023) (`VSLite`), a simple lake varve thickness model (Zhu et al., 2023a) (`Lake_VarveThickness`), as well as the ice core  $\delta^{18}\text{O}$  model (`Ice_d180`), coral  $\delta^{18}\text{O}$  model (`Coral_d180`), coral Sr/Ca ratio model (`Coral_SrCa`) adopted from PRYSM (Dee et al., 2015), among which `Linear` and `Bilinear` are commonly used in paleoclimate data assimilation (PDA), and others are more useful to generate pseudoproxy emulations (Zhu et al., 2023a, b). A summary of available PSMs in `cfr` is listed in Table 5.

Taking the default univariate linear regression model (`Linear`) as a typical example, the essential attributes and methods are listed in Table 6. As illustrated in the notebook tutorial <https://fzhu2e.github.io/cfr/notebooks/psm-linear.html> (Zhu et al., 2024), a `Linear` PSM object can be initialized with a `ProxyRecord` object, and by calling the `calibrate()` method, the PSM utilizes the proxy measurement and the nearest instrumental observation to calibrate the regression coefficients over the instrumental period, with the option to test multiple seasonality candidates, yielding an optimal regression model. Then by calling the `forward()` method, the calibrated model will forward the nearest model simulated climate and generate the pseudoproxy estimate, translating the climate signal to the proxy space. For more illustrations of the other available PSMs, please refer to <https://fzhu2e.github.io/cfr/ug-psm.html> (Zhu et al., 2024).

**Table 5.** A summary of the available proxy system models (PSMs) in `cfr`.

PSM name	Class	References
Univariate Linear	<code>cfr.psm.Linear</code>	Hakim et al. (2016); Tardif et al. (2019)
Multivariate Linear	<code>cfr.psm.BiLinear</code>	Hakim et al. (2016); Tardif et al. (2019)
VS-Lite	<code>cfr.psm.VSLite</code>	Tolwinski-Ward et al. (2011, 2013)
Lake Varve Thickness	<code>cfr.psm.Lake_VarveThickness</code>	Zhu et al. (2023a)
Ice Core $\delta^{18}\text{O}$	<code>cfr.psm.Ice_d180</code>	Dee et al. (2015)
Coral $\delta^{18}\text{O}$	<code>cfr.psm.Coral_d180</code>	Dee et al. (2015)
Coral Sr/Ca	<code>cfr.psm.Coral_SrCa</code>	Dee et al. (2015)

**Table 6.** The essential attributes and methods of the `cfr.psm.Linear` proxy system model (PSM).

Attributes	Description	Methods	Description
<code>pobj</code>	The <code>ProxyRecord</code> object.	<code>calibrate()</code>	Calibrate the PSM.
<code>climate_required</code>	The list of variable names of the required climate input.	<code>forward()</code>	Forward the calibrated model.

## 250 4.5 da: data assimilation

The `da` module implements the data assimilation algorithms, currently including the ensemble Kalman filter (Kalman, 1960; Evensen, 2009) used in the last millenium reanalysis (LMR, Hakim et al., 2016; Tardif et al., 2019) framework. The class is named `EnKF` and Table 7 lists its essential attributes and methods. Direct operations on this class are not recommended unless for debugging purposes. It is designed to be called internally by the `ReconJob` class that we introduce later for a high-level control of the workflow.

255

**Table 7.** The essential attributes and methods of `cfr.da.EnKF`.

Attributes	Description	Methods	Description
<code>prior</code>	The dictionary of prior fields.	<code>gen_prior_samples()</code>	Generate prior samples.
<code>pdb_assim</code>	The assimilated <code>ProxyDatabase</code> .	<code>gen_Ye()</code>	Generate the forward estimates.
<code>seed</code>	The seed for randomization.	<code>gen_Xb()</code>	Generate the background matrix.
<code>nens</code>	The ensemble size.	<code>update_yr()</code>	Update a specific year utilizing the <code>EnKF</code> solver.
<code>recon_vars</code>	The names of the reconstructed variables.	<code>run()</code>	Run the <code>EnKF</code> solver through multiple years.

## 4.6 graphem: Graphical Expectation Maximization

The `graphem` module implements the GraphEM algorithm (Guillot et al., 2015). The class is named `GraphEM` and Table 8 lists its essential attributes and methods. Direct operations on this class are not recommended unless for debugging purposes. Similar to `da`, `graphem` is also designed to be called internally by the `ReconJob` class.

**Table 8.** The essential attributes and methods of `cf r . graphem . GraphEM`.

Attributes	Description	Methods	Description
<code>field_r</code>	The reconstructed field.	<code>fit()</code>	Estimate the parameters and reconstruct the climate field.
<code>proxy_r</code>	The reconstructed proxy matrix.		
<code>calib</code>	Indices of the calibration period.		

## 260 4.7 reconjob: reconstruction workflow management

The `reconjob` module provides pre-defined reconstruction workflows attached to the class `ReconJob`, as listed in Table 9. We will illustrate the usage of these methods in detail in the sections on PDA and GraphEM workflows.

**Table 9.** The essential attributes and methods of `cf r . ReconJob`.

Attributes	Description	Methods	Description
<code>configs</code>	The dictionary of configurations.	<code>load_proxydb()</code>	Load the proxy database.
<code>proxydb</code>	The loaded <code>ProxyDatabase</code> .	<code>filter_proxydb()</code>	Filter the proxy database.
<code>prior</code>	The dictionary of prior fields.	<code>annualize_proxydb()</code>	Annualize the proxy database.
<code>obs</code>	The dictionary of instrumental observations.	<code>split_proxydb()</code>	Split the proxy database.
		<code>load_clim()</code>	Load the simulated or observed climate data.
		<code>annualize_clim()</code>	Annualize the climate data.
		<code>regrid_clim()</code>	Regrid the climate data.
		<code>crop_clim()</code>	Crop the climate data.
		<code>calib_psms()</code>	Calibrate the PSMs for each proxy record.
		<code>forward_psms()</code>	Forward the PSMs for each proxy record.
		<code>run_da()</code>	Run the DA solver.
		<code>run_da_mc()</code>	Run DA with Monte-Carlo iterations.
		<code>prep_graphem()</code>	Prepare data for the GraphEM solver.
		<code>run_graphem()</code>	Run the GraphEM solver.

## 4.8 reconres: result analysis and visualization

The reconres module focuses on postprocessing and visualization of the reconstruction results. It contains a ReconRes class that helps to load the reconstruction outputs stored in netCDF files from a given directory path, and organize the data in the form of either ClimateField or EnsTS. The essential attributes and methods of ReconRes are listed in Table 10. We will illustrate the usage in detail in the sections on PDA and GraphEM workflows.

**Table 10.** The essential attributes and methods of `cfr.ReconRes`.

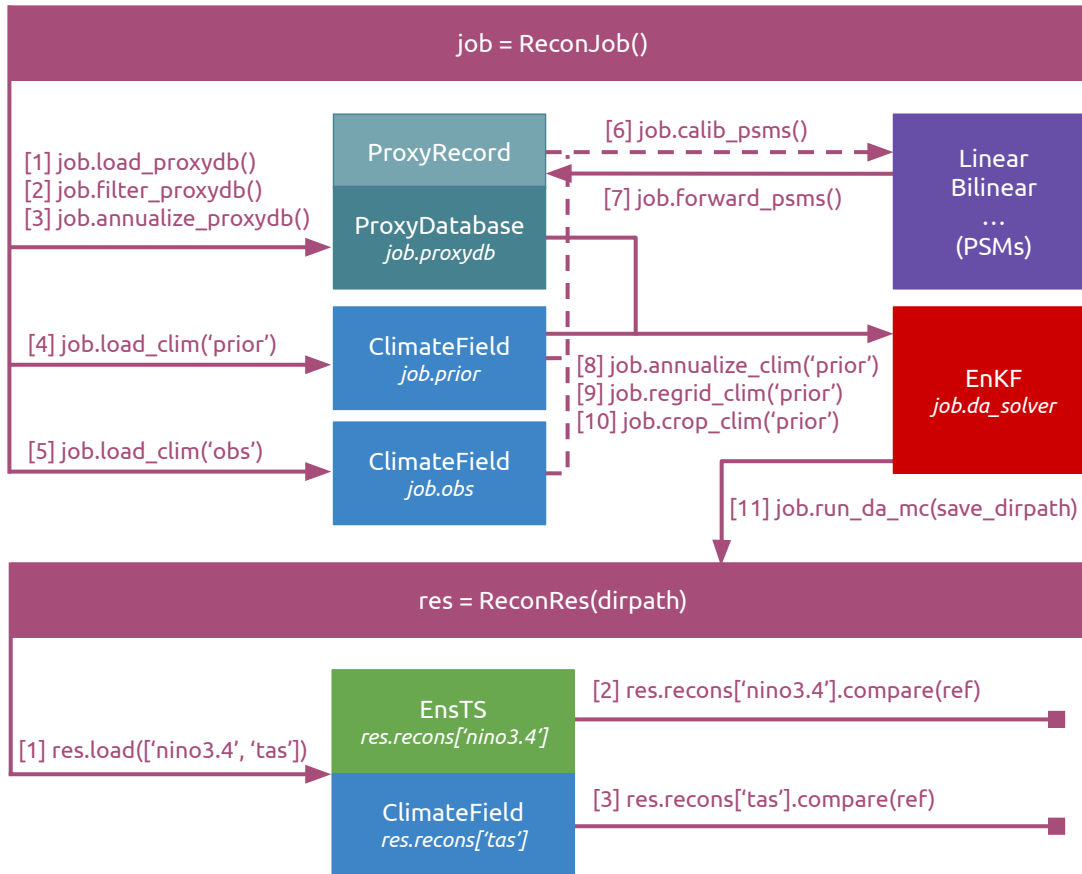
Attributes	Description	Methods	Description
<code>paths</code>	The paths of the reconstruction netCDF files.	<code>load()</code>	Load the reconstruction files.
<code>recons</code>	The dictionary of reconstruction objects.	<code>valid()</code>	Perform validation against given targets.
<code>da</code>	The dictionary of reconstructions in <code>xarray.DataArray</code> .	<code>plot_valid()</code>	Visualize the validation.

## 5 cfr’s PDA workflow

In this section, we illustrate the `cfr` workflow (Fig. 5) with a reconstruction experiment taking the last millennium reanalysis (LMR, Hakim et al., 2016; Tardif et al., 2019) paleoclimate data assimilation (PDA) approach. A similar pseudoproxy reconstruction experiment can be accessed at <https://fzhu2e.github.io/cfr/notebooks/pp2k-ppe-pda.html> (Zhu et al., 2024).

The task here is to assimilate tropical coral records and reconstruct the boreal winter (December-February, DJF) surface temperature field, which can be used to calculate the El Niño/Southern Oscillation (ENSO) indices (NINO3.4 in this example). The “iCESM1” last millennium simulation (Brady et al., 2019) is utilized as the model prior, and the coral records from the PAGES 2k Phase 2 database (PAGES 2k Consortium, 2017) are used as the observations to update the prior. NASA Goddard’s Global Surface Temperature Analysis (GISTEMP) (Lenssen et al., 2019) combining land surface air temperatures primarily from the GHCN-M version 4 (Menne et al., 2018) with the ERSSTv5 sea surface temperature analysis (Huang et al., 2017) is used as the calibration target for the PSMs. Finally, a spatially completed version of the near-surface air temperature and sea-surface temperature analyses product HadCRUT4.6 (Morice et al., 2012) leveraging the GraphEM algorithm (Guillot et al., 2015; Vaccaro et al., 2021) and the “BC09” NINO3.4 reanalysis (Bunge and Clarke, 2009) are used as the validation target for the posterior, i.e., the reconstruction.

# CFR/PDA Workflow



**Figure 5.** The cfr workflow for the last millennium reanalysis (LMR, Hakim et al., 2016; Tardif et al., 2019) paleoclimate data assimilation (PDA) framework. Here “nino3.4” and “tas” stand for two examples of the reconstructed variables for validation, referring to the El Niño/Southern Oscillation index series calculated as the spatial average of surface temperature within the tropical Pacific region (5°N-5°S, 170°W-120°W), and the surface temperature field, respectively. PSMs: proxy system models; EnKF: ensemble Kalman filter.



## 5.1 Load the proxy records, model simulations, and instrumental observations.

First of all, we import the `cfr` package and create a `ReconJob` object named `job`.

```
import cfr
job = cfr.ReconJob()
```

Next, we load the PAGES 2k proxy database (PAGES 2k Consortium, 2017) from the cloud and pick coral records by calling  
285 the `filter_proxydb()` method, then annualize them to be boreal winter average by calling `annualize_proxydb()`:

```
job.load_proxydb('PAGES2kv2')
job.filter_proxydb(by='ptype', keys=['coral'])
job.annualize_proxydb(months=[12, 1, 2], ptypes=['coral'])
```

Now we load the model simulated prior “iCESM1” (Brady et al., 2019) and the instrumentally observed climate data GIS-TEMP (Lenssen et al., 2019) by calling the `load_clim()` method.

```
job.load_clim(tag='prior', path_dict={'tas': 'iCESM_past1000historical/tas'}, anom_period=(1951, 1980))
job.load_clim(tag='obs', path_dict={'tas': 'gistemp1200_GHCnv4_ERSSTv5'}, anom_period=(1951, 1980),
              rename_dict={'tas': 'tempanomaly'})
```

Here, the `tag` argument is used to specify whether the loaded climate data will be used as prior or observation. The  
290 `rename_dict` argument is used to map variable names if the netCDF file does not name a variable as we assumed internally in  
`cfr`. For instance, we assume “lat” for latitude, “lon” for longitude, and “time” for the temporal dimension. The `anom_period`  
argument specifies against which time period we will calculate the anomaly.

Note that here we assume both the prior and observation fields being regularly gridded (lat/lon) datasets, which ensures a fast  
nearest-neighbor search in later processing steps. Any irregularly gridded datasets should be regridded beforehand using the  
`cfr.utils.regrid_field_curv_rect()` function or other third-party tools such as `scipy.interpolate.griddata` and  
295 `ESMPy` (<https://earthsystemmodeling.org/esmpy>).

## 5.2 Calibrate and forward the PSMs.

With the loaded proxy data, climate simulation, and instrumental observation, we are ready to calibrate the proxy system models  
for each proxy record by calling the `calib_psms()` method. For coral records, we use the univariate linear regression model  
Linear as defined by the `ptype_psm_dict` argument. The `ptype_season_dict`, `calib_period`, and `ptype_clim_dict`  
300 arguments specify the season, timespan (in year), and climate variables involved for calibration, respectively.

```
job.calib_psms(
    ptype_psm_dict={'coral.d180': 'Linear', 'coral.calc': 'Linear', 'coral.SrCa': 'Linear'},
    ptype_season_dict={'coral.d180': [12, 1, 2], 'coral.calc': [12, 1, 2], 'coral.SrCa': [12, 1, 2]},
    ptype_clim_dict={'coral.d180': ['tas'], 'coral.calc': ['tas'], 'coral.SrCa': ['tas']},
    calib_period=(1850, 2015))
```

In this example, we have specified a fixed seasonality for the whole proxy database. Users may also perform calibration of each proxy record through a for loop of the proxy database, specifying various seasonality manually or based on the metadata if available. For records whose corresponding PSM is successfully calibrated, a “calibrated” tag is added to the record to facilitate filtering later. Proxy records with low calibration skill will be assigned a small weight during the assimilation process as per the EnKF algorithm, so there is no need to set a calibration threshold to manually filter the records. However, with the flexibility of the package, users have the freedom to easily filter the proxy database, using the `cfr.ProxyDatabase.filter()` method, based on any calibration threshold they like.

Once the PSMs are calibrated (and filtered), we can launch (i.e., forward) them by simply calling the `forward_psms()` method.

```
job.forward_psms()
```

### 310 5.3 Define the season, resolution, and domain of the reconstruction.

Now we annualize the prior by calling the `annualize_clim()` method, and specify on which resolution and domain we would like to conduct the reconstruction by calling the `regrid_clim()` and `crop_clim()` method.

```
job.annualize_clim(tag='prior', months=[12, 1, 2])
job.regrid_clim(tag='prior', nlat=42, nlon=63)
job.crop_clim(tag='prior', lat_min=-35, lat_max=35)
```

Here we annualize the prior to boreal winter, and regrid the field to a global grid with 42 latitudes and 63 longitudes, and then crop the domain to be within the 35°S to 35°N band where the corals are located.

### 315 5.4 Conduct the Monte-Carlo iterations of the data assimilation steps.

Once the pre-processing is complete, we are ready to conduct the data assimilation step. We call the `run_da_mc()` method to perform Monte-Carlo iterations of the EnKF assimilation steps. The argument `save_dirpath` specifies the directory path where we store the reconstruction results, and the argument `recon_seeds` specify the seed for randomization for each Monte-Carlo iteration.

```
job.run_da_mc(save_dirpath='./recons/lmr-real-pages2k', recon_seeds=list(range(1, 11)))
```

### 320 5.5 Validate the reconstruction.

Once the Monte-Carlo iterations are done, we should have several netCDF files stored in the specified directory, and we can initiate a `ReconRes` object by specifying the directory path, and load the results by calling the `load()` method.

```
res = cfr.ReconRes('./recons/lmr-real-pages2k')
res.load(['nino3.4', 'tas'])
```

Note that the variable names listed in the `load()` method are predefined. 'nino3.4' refers to the NINO3.4 index and the reconstructed ensemble time series will be formed as a `EnSTS` object. 'tas' refers to the surface temperature field and the ensemble mean will be formed as a `ClimateField` object. These two objects will be organized in a dictionary named `res.recons`, and their original `xarray.DataArray` forms will be stored as `res.recons['nino3.4'].da` and `res.recons['tas'].da` for more universal purposes.

To evaluate the reconstruction skill of the surface temperature field, we load a reference target named “HadCRUT4.6\_GraphEM” (Vaccaro et al., 2021), which is a spatially completed version of the near-surface air temperature and sea-surface temperature analyses product HadCRUT4.6 (Morice et al., 2012) leveraging the GraphEM algorithm (Guillot et al., 2015; Vaccaro et al., 2021), and we validate both the prior and the posterior (`res.recons['tas']`) by calling the `compare()` method over the 1874-2000 CE timespan, which returns another `ClimateField` object, and we are able to visualize it by calling its `plot()` method.

```
target = cfr.ClimateField().fetch('HadCRUT4.6_GraphEM', vn='tas').get_anom((1951, 1980))
target = target.annualize(months=[12, 1, 2]).crop(lat_min=-35, lat_max=35)

# validate the prior
stat = 'corr'
valid_fd = job.prior['tas'].compare(target, stat=stat, timespan=(1874, 2000))
fig, ax = valid_fd.plot(
    title=f'{stat}(prior, obs), mean={valid_fd.geo_mean().value[0,0]:.2f}',
    projection='PlateCarree', latlon_range=(-32, 32, 0, 360), plot_cbar=False)

# validate the reconstruction
valid_fd = res.recons['tas'].compare(target, stat=stat, timespan=(1874, 2000))
valid_fd.plot_kwargs.update({'cbar_orientation': 'horizontal', 'cbar_pad': 0.1})

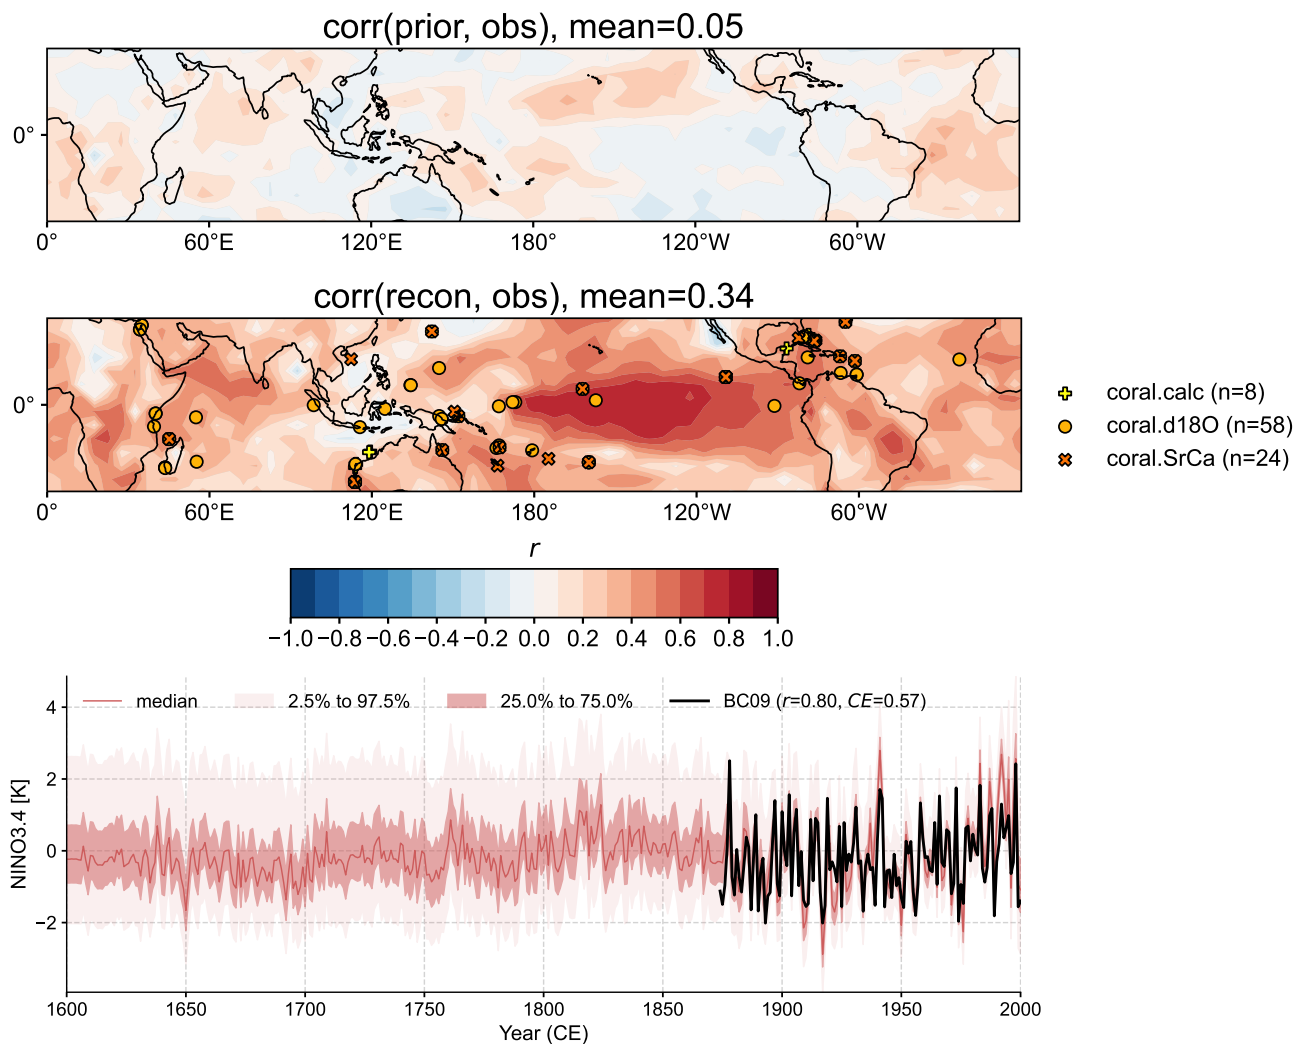
fig, ax = valid_fd.plot(
    title=f'{stat}(prior, obs), mean={valid_fd.geo_mean().value[0,0]:.2f}',
    projection='PlateCarree', latlon_range=(-32, 32, 0, 360),
    plot_proxydb=True, proxydb=job.proxydb.filter(by='tag', keys=['calibrated']), plot_proxydb_lgd=True,
    proxydb_lgd_kws={'loc': 'lower left', 'bbox_to_anchor': (1, 0)})
```

With the above code lines, we get Fig. 6 top and middle. The top shows the map of Pearson correlation between the prior and the reference target (“HadCRUT4.6\_GraphEM”), while the middle shows that between the posterior (i.e., reconstruction) and the reference target, indicating that the reconstruction is working as expected, boosting the mean correlation from 0.05 to 0.34. This is one measure of the information added by the proxies.

To evaluate the reconstruction skill of the NINO3.4 index, we load a reference target named “BC09” (Bunge and Clarke, 2009), and then validate `res.recons['nino3.4']` against BC09 by calling the `compare()` method, which returns another  
340 `EnsTS` object that can be visualized by calling its `plot_qs()` method.

```
bc09 = cfr.EnsTS().fetch('BC09_NINO34').annualize(months=[12, 1, 2])
fig, ax = res.recons['nino3.4'].compare(bc09, timespan=(1874, 2000), ref_name='BC09').plot_qs()
ax.set_xlim(1600, 2000)
ax.set_ylabel('NINO3.4 [K]')
```

The above code lines yield Fig. 6 bottom, and it indicates the reconstruction skill over the instrumental period is remarkably high, with correlation coefficient  $r = 0.80$  and the coefficient of efficiency  $CE = 0.57$ .




**Figure 6.** Validation of the reconstruction taking the paleoclimate data assimilation approach. (top) The correlation map between the prior filed and the observation target (Vaccaro et al., 2021). (middle) The correlation map between the reconstruction median field and the observation target. (bottom) The correlation ( $r$ ) and coefficient of efficiency ( $CE$ ) between the reconstructed NINO3.4 median series and the BC09 reanalysis (Bunge and Clarke, 2009). Note that the proxy calibration is performed over 1850-2015 CE, while the validation is performed over 1874-2000 CE, indicating an overlap of calibration and validation periods. The risk of overfitting is mitigated by randomly picking model prior and proxies during each Monte-Carlo iteration (Hakim et al., 2016; Tardif et al., 2019).

## 6 cfr's GraphEM workflow

345 In this section, we illustrate the `cfr` workflow (Fig. 7) for GraphEM (Section 3.2) with a similar reconstruction experiment setup that we presented above for the illustration of the PDA workflow. Another similar pseudoproxy reconstruction experiment can be accessed at <https://fzhu2e.github.io/cfr/notebooks/pp2k-ppe-graphem.html> (Zhu et al., 2024).

### CFR/GraphEM Workflow



**Figure 7.** The `cfr` workflow for the Graphical Expectation-Maximization (GraphEM, Guillot et al., 2015) algorithm. Here “nino3.4” and “tas” stand for two examples of the reconstructed variables for validation, referring to the El Niño/Southern Oscillation index series calculated as the spatial average of the surface temperature within the tropical Pacific region (5°N-5°S, 170°W-120°W), and the surface temperature field, respectively.

## 6.1 Load the proxy records and instrumental observations.

The cfr workflow for GraphEM is very similar to that for PDA. The first part is loading and preprocessing the required data, including the PAGES 2k proxy records (PAGES 2k Consortium, 2017) and the GISTEMP instrumental observation (Lenssen et al., 2019). The only difference compared to the PDA workflow is that no model prior is needed. Due to the data sparsity requirement of the GraphEM method, we set a relatively small domain around the NINO3.4 region for the reconstruction (20°S-20°N, 150°E-100°W) in this specific experiment.

```
import cfr
job = cfr.ReconJob()

# load and preprocess the proxy database
job.load_proxydb('PAGES2kv2')
job.filter_proxydb(by='ptype', keys=['coral'])
job.annualize_proxydb(months=[12, 1, 2], ptypes=['coral'])

# load the instrumental observation dataset
job.load_clim(tag='obs', path_dict={'tas': 'gistemp1200_GHCnv4_ERSSTv5'},
             anom_period=(1951, 1980), rename_dict={'tas': 'tempanomaly'})

# define the season, resolution, and domain of the reconstruction
job.annualize_clim(tag='obs', months=[12, 1, 2])
job.regrid_clim(tag='obs', nlat=42, nlon=63)
job.crop_clim(tag='obs', lat_min=-20, lat_max=20, lon_min=150, lon_max=260)
```

## 6.2 Prepare the GraphEM solver

Compared to PDA, the preparation of the solver is much easier with GraphEM. Here, we set the calibration period to be 1901-2000 CE, while the reconstruction period is 1871-2000 CE; this reserves 1871-1900 for validation. Such dates can be easily changed with the arguments recon\_period and calib\_period. We also process the proxy database to be more uniform by keeping only the records that span the full reconstruction time period (uniform\_pdb=True) to make the GraphEM algorithm more efficient.

```
job.prep_graphem(recon_period=(1871, 2000), calib_period=(1901, 2000), uniform_pdb=True)
```

## 6.3 Run the GraphEM solver.

We run the GraphEM solver with the “hybrid” approach (Section 3.2) with a cutoff radius (cutoff\_radius) of 5000 km for the neighborhood graph, and 2% target sparsity of the in-field part (sp\_FF) and 2% target sparsity for the climate field/proxy

part (sp\_FP) of the inverse covariance matrix for the GLASSO method. For more details about these parameters, please refer to Guillot et al. (2015).

```
job.run_graphem(save_dirpath='./recons/graphem-real-pages2k',
               graph_method='hybrid', cutoff_radius=5000, sp_FF=2, sp_FP=2)
```

#### 6.4 Validate the reconstruction.

365 Similar to the PDA reconstruction experiment, we validate the reconstructed surface temperature field against the “Had-CRUT4.6\_GraphEM” dataset (Vaccaro et al., 2021) and the reconstructed NINO3.4 index against the “BC09” reanalysis (Bunge and Clarke, 2009), but only over the 1874-1900 CE timespan as the reconstruction over the calibration period (1901-2000 CE) is identical to the GISTEMP observation (Fig. 8).

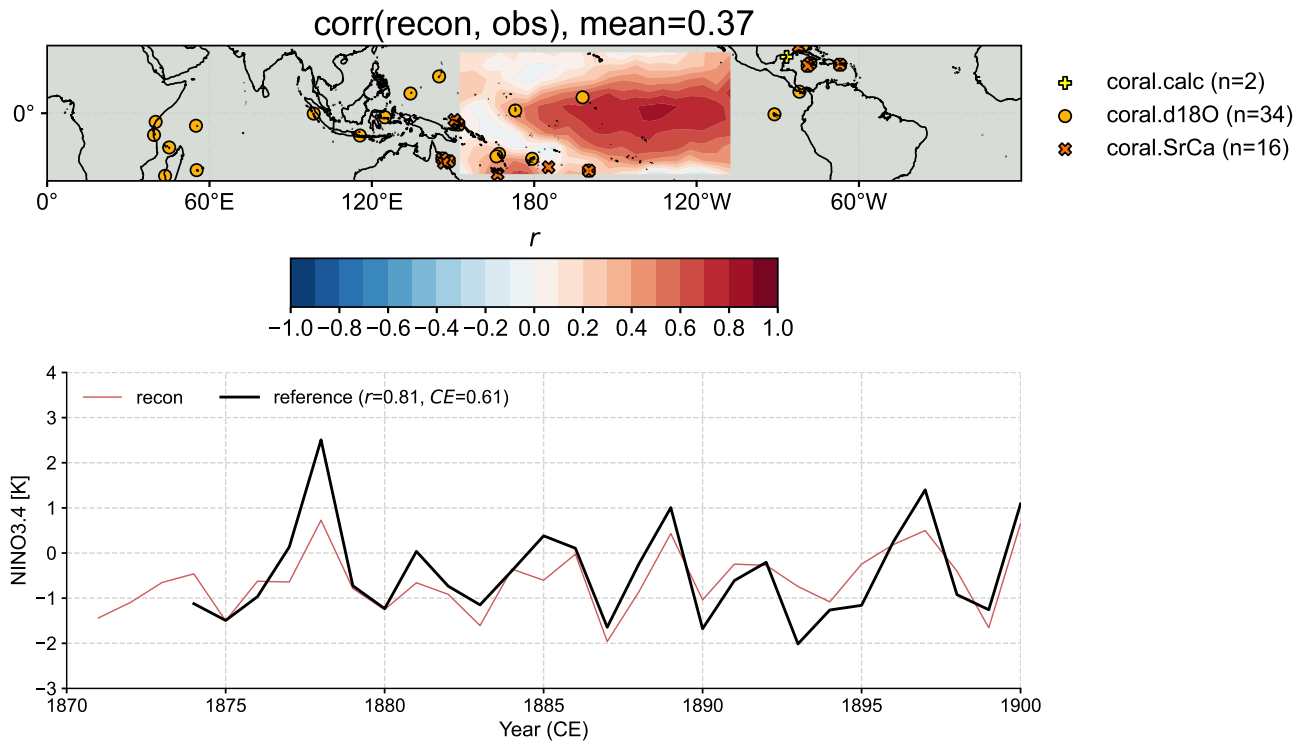
```
res = cfr.ReconRes('./recons/graphem-real-pages2k')
res.load(['nino3.4', 'tas'], verbose=True)

# validate the reconstructed surface temperature field
target = cfr.ClimateField().fetch('HadCRUT4.6_GraphEM', vn='tas').get_anom((1951, 1980))
target = target.annualize(months=[12, 1, 2]).crop(lat_min=-25, lat_max=25, lon_min=120, lon_max=280)

stat = 'corr'
valid_fd = res.recons['tas'].compare(target, stat=stat, timespan=(1874, 1900))
valid_fd.plot_kwargs.update({'cbar_orientation': 'horizontal', 'cbar_pad': 0.1})
fig, ax = valid_fd.plot(
    title=f'{stat}(recon, obs), mean={valid_fd.geo_mean().value[0,0]:.2f}',
    projection='PlateCarree', latlon_range=(-25, 25, 0, 360),
    plot_cbar=True, plot_proxydb=True, proxydb=job.proxydb, plot_proxydb_lgd=True,
    proxydb_lgd_kws={'loc': 'lower left', 'bbox_to_anchor': (1, 0)})

# validate the reconstructed NINO3.4
bc09 = cfr.EnsTS().fetch('BC09_NINO34').annualize(months=[12, 1, 2])
fig, ax = res.recons['nino3.4'].compare(bc09, timespan=(1874, 1900)).plot(label='recon')
ax.set_xlim(1870, 1900)
ax.set_ylim(-3, 4)
ax.set_ylabel('NINO3.4 [K]')
```





**Figure 8.** Validation of the reconstruction taking the Graphical Expectation Maximization (GraphEM) approach. (top) The correlation map between the reconstruction median field and the observation target (Vaccaro et al., 2021). (bottom) The correlation ( $r$ ) and coefficient of efficiency ( $CE$ ) between the reconstructed NINO3.4 median series and the BC09 reanalysis (Bunge and Clarke, 2009). Note that the calibration is performed over 1901-2000 CE, while the validation is performed over 1874-1900 CE.

The result indicates that GraphEM yields a seemingly slightly better skill compared to PDA, although we note that this skill is sensitive to the parameters (`cutoff_radius`, `sp_FF`, and `sp_FP`), which in practice requires some tuning work in order to be optimized. In addition, we note that the current comparison is not “apples-to-apples” due to the different domains and timespans for validation. A more systematic comparison with a consistent setup is presented in the next section.

## 6.5 Comparison to PDA

A main value proposition of `cfr` is to enable convenient inter-methodological comparisons of multiple reconstructions. Here, we illustrate how to conduct an “apples-to-apples” comparison of the reconstructions generated by the GraphEM and PDA approaches.

First, we create `ReconRes` objects for the two reconstructions and load the validation targets “HadCRUT4.6\_GraphEM” (Vaccaro et al., 2021) and “BC09” (Bunge and Clarke, 2009).

```

res_graphem = cfr.ReconRes('./recons/graphem-real-pages2k')
res_lmr = cfr.ReconRes('./recons/lmr-real-pages2k')

tas_HadCRUT = cfr.ClimateField().fetch('HadCRUT4.6_GraphEM', vn='tas').get_anom((1951, 1980))
tas_HadCRUT = tas_HadCRUT.annualize(months=[12, 1, 2])
nino34_bc09 = cfr.EnsTS().fetch('BC09_NINO34').annualize(months=[12, 1, 2])

```

Then we call the `valid()` method for the reconstructions to compute the validation statistics, including Pearson correlation and coefficient of efficiency ( $CE$ ) against the same set of validation targets over the same domain (17°S-17°N, 153°E-108°W) and the same time span 1874-1900 CE.

```

res_graphem.valid(target_dict={'tas': tas_HadCRUT, 'nino3.4': nino34_bc09},
                  timespan=(1874, 1900), stat=['corr', 'CE'])

res_lmr.valid(target_dict={'tas': tas_HadCRUT, 'nino3.4': nino34_bc09},
              timespan=(1874, 1900), stat=['corr', 'CE'])

```

Finally, we call the `plot_valid()` method to visualize the validation results (Fig. 9).

```

fig, ax = res_graphem.plot_valid(
    target_name_dict={'tas': 'HadCRUT4.6', 'nino3.4': 'BC09'},
    recon_name_dict={'tas': 'GraphEM/tas', 'nino3.4': 'NINO3.4 [K]'},
    valid_fd_kws=dict(projection='PlateCarree', latlon_range=(-17, 17, 153, 252), plot_cbar=True),
    valid_ts_kws=dict(xlim=(1870, 1900), ylim=(-3, 4)))

fig, ax = res_lmr.plot_valid(
    target_name_dict={'tas': 'HadCRUT4.6', 'nino3.4': 'BC09'},
    recon_name_dict={'tas': 'PDA/tas', 'nino3.4': 'NINO3.4 [K]'},
    valid_fd_kws=dict(projection='PlateCarree', latlon_range=(-17, 17, 153, 252), plot_cbar=True),
    valid_ts_kws=dict(xlim=(1870, 1900), ylim=(-3, 4)))

```

It can be seen that the GraphEM and PDA approaches show overall comparable skill with a consistent spatiotemporal validation setup, with the mean correlation skill being around 0.5 in the specific domain (Fig. 9a, d). PDA shows fairly uniform reconstruction skill, while GraphEM's is more spatially variable, and is notably higher in the eastern equatorial Pacific. This leads to a slightly higher reconstruction skill for NINO3.4 (Fig. 9c, f), as measured by both  $r$  (0.82 vs 0.72) and  $CE$  (0.63 vs 0.51).

The weak and negative correlation shown in Fig. 9a is located far from proxy locales, and could be due to two main causes, which are not mutually exclusive:

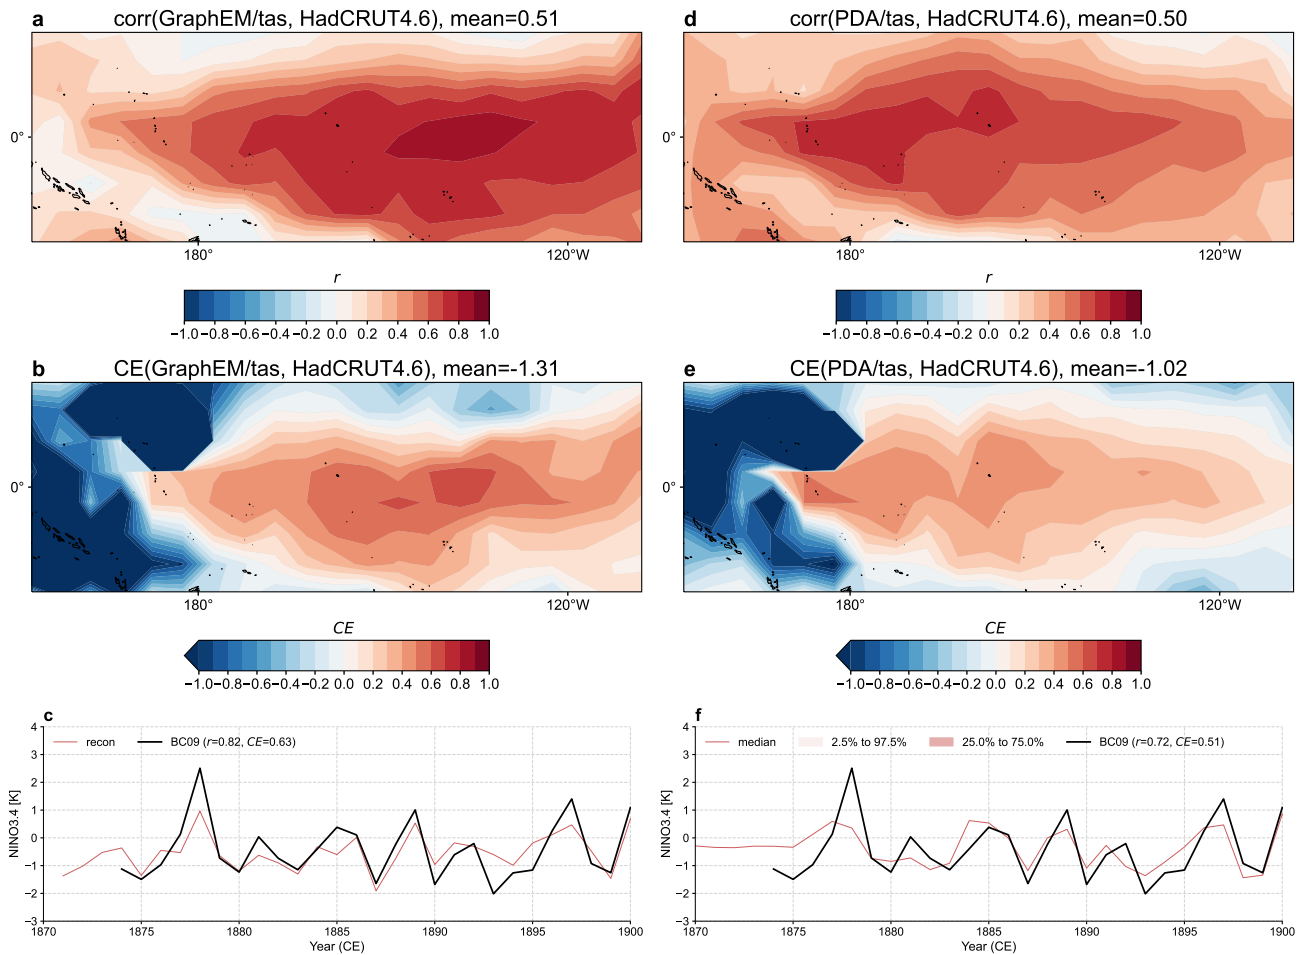
390 **Covariance Modeling** The graph used here has been minimally optimized, and it is possible that it is too sparse, which would  
understate the strength of teleconnections. This would result in under-constrained temperature far from proxy locales, as  
observed here.

**Proxy Modeling** GraphEM implicitly applies a so-called “direct regression” framework (Brown, 1994), whereby the climate  
variable (here,  $t_{as}$ ) is modeled as a linear combination of all the proxies within the neighborhood identified by the graph.  
395 This stands in contrast to the “indirect regression” framework (Brown, 1994) used as part of PDA, where proxies are  
modeled as linearly dependent on local temperature. In the climate context, it has been shown that direct regression is  
problematic and may result in high variance estimates, especially when underconstrained by observations (Tingley and  
Li, 2012). This high variance would manifest as temperature estimates with poor correlation to the instrumental target.

Both the GraphEM and PDA based reconstructions show poor  $CE$  skill over the western Pacific region (Fig. 9b, e), indicating  
400 large amplitude biases, which is likely due to issues in proxy modeling: here corals are considered temperature recorders, while  
in the western Pacific their signal is known to be largely driven by surface hydrology (e.g. Lough, 2010; Tierney et al., 2015).

By conducting more reconstruction experiments, we find that the PDA approach shows a robust ability to handle various  
spatiotemporal reconstruction setups, while the GraphEM approach is computationally limited to restricted reconstruction  
domains and time intervals so as to ensure a relatively dense data matrix, without which the EM algorithm can take a long time  
405 converging. This computational restriction also limits the degree of graph optimization that can be performed.

It is also worth noting that the GraphEM-based reconstruction produced by `cfr` does not, at this stage, produce ensembles. In  
future versions, approaches such as block bootstrapping will be implemented for uncertainty quantification, following Guillot  
et al. (2015).



**Figure 9.** Comparison of the reconstructions generated by the GraphEM and PDA approaches against the same validation targets “HadCRUT4.6\_GraphEM” (Vaccaro et al., 2021) and “BC09” (Bunge and Clarke, 2009) over the same domain (17°S-17°N, 153°E-108°W) and the same timespan 1874-1900 CE.

## 7 Summary

410 Climate field reconstruction provides spatial details about past climate, and the `cfr` package makes this process intuitive, modular, and efficient for Python users. In particular, `cfr` enables a sophisticated laboratory for comparisons of different climate field reconstruction methods, which is of critical importance to better estimate and interpret reconstructions, given their sensitivity to multiple error sources rooted in the climate model prior, proxy records, proxy system modeling, and the reconstruction algorithm itself. In addition, `cfr` is a handy toolbox for paleoclimate data analysis and visualization, lowering

415 the bar for processing and visualizing proxy records, climate model output, and instrumental observations.

The current version of the package provides native support for climate field reconstruction over the Common Era, and its framework can be readily expanded for deep-time reconstructions (e.g., Tierney et al., 2020; Osman et al., 2021) through porting PSMs of deep-time proxies such as planktic foraminiferal  $\delta^{18}\text{O}_c$ , Mg/Ca, TEX'86, and UK'37 (e.g., Malevich et al., 2019; Tierney and Tingley, 2014, 2018; Tierney et al., 2019; King et al., 2023). As an open-source and community-based code, future versions of the package will combine contributions from the core development team and the community to support more CFR methods such as BARCAST (Tingley and Huybers, 2013), an online ensemble Kalman filter (Perkins and Hakim, 2017, 2021), particle filters (Goosse et al., 2010; Dubinkina et al., 2011; Dubinkina and Goosse, 2013; Liu et al., 2017; Shi et al., 2019; Lyu et al., 2021), a wider array of PSMs for deep-time reconstructions, along with more postprocessing, visualization, and validation functionalities, catalyzing open and reproducible paleoclimate research. A contributing guide is available at: <https://fzhu2e.github.io/cfr>. We hope this ethos invites more replicability and reproducibility in paleoclimate reconstructions, thereby deepening confidence in our knowledge of past climates.

*Code and data availability.* The `cfr` codebase is available at: <https://github.com/fzhu2e/cfr> under the BSD-3 Clause licence. Its documentation can be accessed at: <https://fzhu2e.github.io/cfr>. The exact version used to produce the results shown in this paper is archived on Zenodo at: <https://zenodo.org/record/10575537> (Zhu et al., 2024), which has been tested with the current mainstream Python versions: 3.9, 3.10, and 3.11. Please follow the evolving installation guide at <https://fzhu2e.github.io/cfr/ug-installation.html> for more updated details, including the recommended Python version.

All datasets leveraged in the examples illustrated in this study can be automatically fetched from the cloud with `cfr`'s remote-data-loading feature, without the hassle of manually searching and downloading, including:

- The PAGES 2k phase 2 global multiproxy database (PAGES 2k Consortium, 2017) hosted on the National Center for Environmental Information's World Data Service for Paleoclimatology (<https://www.ncei.noaa.gov/access/paleo-search/study/21171>). A reorganized copy for the remote-data-loading purpose is hosted on Zenodo (<https://zenodo.org/record/8367746>) (Zhu, 2023) and Github (<https://github.com/fzhu2e/cfr-data/raw/main/pages2kv2.json>).
- The “pseudoPAGES2k” pseudoproxy dataset (Zhu et al., 2023a) hosted on Zenodo (<https://doi.org/10.5281/zenodo.8173256>) (Zhu et al., 2023b) and Github (<https://github.com/fzhu2e/paper-pseudoPAGES2k>).
- The “iCESM1” last millennium simulation (Brady et al., 2019) hosted on a data server at University of Washington by Rorbert Tardif (<https://atmos.washington.edu/~rtardif/LMR/prior>).
- The NASA Goddard's Global Surface Temperature Analysis (GISTEMP) (Lenssen et al., 2019) combining land surface air temperatures primarily from the GHCN-M version 4 (Menne et al., 2018) with the ERSSTv5 sea surface temperature analysis (Huang et al., 2017) hosted on NASA's website ([https://data.giss.nasa.gov/pub/gistemp/gistemp1200\\_GHCNv4\\_ERSSTv5.nc.gz](https://data.giss.nasa.gov/pub/gistemp/gistemp1200_GHCNv4_ERSSTv5.nc.gz)).
- The spatially completed version of the near-surface air temperature and sea-surface temperature analyses product HadCRUT4.6 (Morice et al., 2012) leveraging the GraphEM algorithm (Vaccaro et al., 2021) hosted on Zenodo (<https://zenodo.org/records/4601616>) and Github ([https://github.com/fzhu2e/cfr-data/raw/main/HadCRUT4.6\\_GraphEM\\_median.nc](https://github.com/fzhu2e/cfr-data/raw/main/HadCRUT4.6_GraphEM_median.nc)).
- The “BC09” NINO3.4 reanalysis (Bunge and Clarke, 2009) hosted on Zenodo (<https://zenodo.org/record/8367746>) (Zhu, 2023) and Github ([https://github.com/fzhu2e/cfr-data/raw/main/BC09\\_NINO34.csv](https://github.com/fzhu2e/cfr-data/raw/main/BC09_NINO34.csv)) for the remote-data-loading purpose.

450 *Author contributions.* F.Z. and J.E.G. designed the research and drafted the manuscript. All authors reviewed and revised the manuscript, and contributed to the development of the package.

*Competing interests.* The authors declare no competing interests.

*Acknowledgements.* The authors acknowledge support from the National Science Foundation (grant 1948822 to F.Z., J.E.G., D.K., and grants 2202777, 2303530 to F.Z.), the Climate Program Office of the National Oceanographic and Atmospheric Administration (grants 455 NA18OAR4310426 to J.E.G. and NA18OAR4310422 to G.J.H.), the Office of Naval Research (grant N00014-21-1-2437 to D.K.), and the Simons Foundation Collaboration Grant for Mathematicians (grant 526851 to D.G.).

## References

- Acevedo, W., Fallah, B., Reich, S., and Cubasch, U.: Assimilation of pseudo-tree-ring-width observations into an atmospheric general circulation model, *Climate of the Past*, 13, 545–557, <https://doi.org/https://doi.org/10.5194/cp-13-545-2017>, 2017.
- 460 Annan, J. D. and Hargreaves, J. C.: A New Global Reconstruction of Temperature Changes at the Last Glacial Maximum, *Climate of the Past*, 9, 367–376, <https://doi.org/10.5194/cp-9-367-2013>, 2013.
- Annan, J. D., Hargreaves, J. C., and Mauritsen, T.: A New Global Surface Temperature Reconstruction for the Last Glacial Maximum, *Climate of the Past*, 18, 1883–1896, <https://doi.org/10.5194/cp-18-1883-2022>, 2022.
- Bova, S., Rosenthal, Y., Liu, Z., Godad, S. P., and Yan, M.: Seasonal origin of the thermal maxima at the Holocene and the last interglacial, 465 *Nature*, 589, 548–553, <https://doi.org/10.1038/s41586-020-03155-x>, number: 7843 Publisher: Nature Publishing Group, 2021.
- Brady, E., Stevenson, S., Bailey, D., Liu, Z., Noone, D., Nusbaumer, J., Otto-Bliesner, B. L., Tabor, C., Tomas, R., Wong, T., Zhang, J., and Zhu, J.: The Connected Isotopic Water Cycle in the Community Earth System Model Version 1, *Journal of Advances in Modeling Earth Systems*, 11, 2547–2566, <https://doi.org/10.1029/2019MS001663>, 2019.
- Brown, P. J.: Measurement, Regression, and Calibration, vol. 12 of *Oxford Statistical Science Series*, Oxford University Press, USA, s edn., 470 1994.
- Bunge, L. and Clarke, A. J.: A Verified Estimation of the El Niño Index Niño-3.4 since 1877, *Journal of Climate*, 22, 3979–3992, <https://doi.org/10.1175/2009JCLI2724.1>, 2009.
- Büntgen, U., Allen, K., Anchukaitis, K. J., Arseneault, D., Boucher, É., Bräuning, A., Chatterjee, S., Cherubini, P., Churakova (Sidorova), O. V., Corona, C., Gennaretti, F., GieBinger, J., Guillet, S., Guiot, J., Gunnarson, B., Helama, S., Hochreuther, P., Hughes, M. K., 475 Huybers, P., Kirilyanov, A. V., Krusic, P. J., Ludescher, J., Meier, W. J. H., Myglan, V. S., Nicolussi, K., Oppenheimer, C., Reinig, F., Salzer, M. W., Seftigen, K., Stine, A. R., Stoffel, M., St. George, S., Tejedor, E., Trevino, A., Trouet, V., Wang, J., Wilson, R., Yang, B., Xu, G., and Esper, J.: The influence of decision-making in tree ring-based climate reconstructions, *Nature Communications*, 12, 3411, <https://doi.org/10.1038/s41467-021-23627-6>, 2021.
- Bürger, G., Fast, I., and Cubasch, U.: Climate reconstruction by regression –32 variations on a theme, *Tellus A: Dynamic Meteorology and Oceanography*, 58, 227–235, <https://doi.org/10.1111/j.1600-0870.2006.00164.x>, 2006.
- 480 Christiansen, B. and Ljungqvist, F. C.: Challenges and Perspectives for Large-Scale Temperature Reconstructions of the Past Two Millennia, *Reviews of Geophysics*, 55, 40–96, <https://doi.org/10.1002/2016RG000521>, 2017.
- Cook, E. R., Woodhouse, C. A., Eakin, C. M., Meko, D. M., and Stahle, D. W.: Long-Term Aridity Changes in the Western United States, *Science*, 306, 1015–1018, <https://doi.org/10.1126/science.1102586>, 2004.
- 485 Cook, E. R., Anchukaitis, K. J., Buckley, B. M., D’Arrigo, R. D., Jacoby, G. C., and Wright, W. E.: Asian Monsoon Failure and Megadrought During the Last Millennium, *Science*, 328, 486–489, <https://doi.org/10.1126/science.1185188>, 2010.
- Dee, S., Emile-Geay, J., Evans, M. N., Allam, A., Steig, E. J., and Thompson, D.: PRYSM: An open-source framework for PProX System Modeling, with applications to oxygen-isotope systems, *J. Adv. Model. Earth Syst.*, 7, 1220–1247, <https://doi.org/10.1002/2015MS000447>, 2015.
- 490 Dempster, A. P., Laird, N. M., and Rubin, D. B.: Maximum likelihood estimation from incomplete data via the EM algorithm (with discussion), *J. Roy. Stat. Soc. B.*, 39, 1–38, 1977.
- Dirren, S. and Hakim, G. J.: Toward the assimilation of time-averaged observations, *Geophysical Research Letters*, 32, <https://doi.org/https://doi.org/10.1029/2004GL021444>, 2005.

- Dubinkina, S. and Goosse, H.: An assessment of particle filtering methods and nudging for climate state reconstructions, *Climate of the Past*, 9, 1141–1152, <https://doi.org/10.5194/cp-9-1141-2013>, publisher: Copernicus GmbH, 2013.
- 495 Dubinkina, S., Goosse, H., Sallaz-Damaz, Y., Crespin, E., and Crucifix, M.: Testing a particle filter to reconstruct climate changes over the past centuries, *International Journal of Bifurcation and Chaos*, 21, 3611–3618, <https://doi.org/10.1142/S0218127411030763>, publisher: World Scientific Publishing Co., 2011.
- Emile-Geay, J., Cobb, K., Mann, M., and Wittenberg, A. T.: Estimating Central Equatorial Pacific SST variability over the Past Millennium. Part 1: Methodology and Validation, *J. Clim.*, 26, 2302–2328, <https://doi.org/10.1175/JCLI-D-11-00510.1>, 2013a.
- 500 Emile-Geay, J., Cobb, K., Mann, M., and Wittenberg, A. T.: Estimating Central Equatorial Pacific SST variability over the Past Millennium. Part 2: Reconstructions and Implications, *J. Clim.*, 26, 2329–2352, <https://doi.org/10.1175/JCLI-D-11-00511.1>, 2013b.
- Evans, M. N., Kaplan, A., and Cane, M. A.: Pacific sea surface temperature field reconstruction from coral  $\delta^{18}\text{O}$  data using reduced space objective analysis, *Paleoceanography*, 17, 7–17–13, <https://doi.org/https://doi.org/10.1029/2000PA000590>, 2002.
- 505 Evans, M. N., Tolwinski-Ward, S. E., Thompson, D. M., and Anchukaitis, K. J.: Applications of proxy system modeling in high resolution paleoclimatology, *Quaternary Science Reviews*, 76, 16–28, <https://doi.org/10.1016/j.quascirev.2013.05.024>, 2013.
- Evensen, G.: *Data Assimilation: The Ensemble Kalman Filter*, Springer, 2nd ed. 2009 edition edn., 2009.
- Fang, M., Li, X., Chen, H. W., and Chen, D.: Arctic Amplification Modulated by Atlantic Multidecadal Oscillation and Greenhouse Forcing on Multidecadal to Century Scales, *Nature Communications*, 13, 1865, <https://doi.org/10.1038/s41467-022-29523-x>, 2022.
- 510 Franke, J., Brönnimann, S., Bhend, J., and Brugnara, Y.: A Monthly Global Paleo-Reanalysis of the Atmosphere from 1600 to 2005 for Studying Past Climatic Variations, *Scientific Data*, 4, 170 076, <https://doi.org/10.1038/sdata.2017.76>, 2017.
- Friedman, J., Hastie, T., and Tibshirani, R.: Sparse inverse covariance estimation with the graphical lasso, *Biostat*, 9, 432–441, <https://doi.org/10.1093/biostatistics/kxm045>, 2008.
- Gebhardt, C., Kühl, N., Hense, A., and Litt, T.: Reconstruction of Quaternary Temperature Fields by Dynamically Consistent Smoothing, *Climate Dynamics*, 30, 421–437, <https://doi.org/10.1007/s00382-007-0299-9>, 2008.
- 515 Goosse, H., Renssen, H., Timmermann, A., Bradley, R. S., and Mann, M. E.: Using Paleoclimate Proxy-Data to Select Optimal Realisations in an Ensemble of Simulations of the Climate of the Past Millennium, *Climate Dynamics*, 27, 165–184, <https://doi.org/10.1007/s00382-006-0128-6>, 2006a.
- Goosse, H., Renssen, H., Timmermann, A., Bradley, R. S., and Mann, M. E.: Using paleoclimate proxy-data to select optimal realisations in an ensemble of simulations of the climate of the past millennium, *Climate Dynamics*, 27, 165–184, <https://doi.org/10.1007/s00382-006-0128-6>, 2006b.
- 520 Goosse, H., Crespin, E., de Montety, A., Mann, M. E., Renssen, H., and Timmermann, A.: Reconstructing Surface Temperature Changes over the Past 600 Years Using Climate Model Simulations with Data Assimilation, *Journal of Geophysical Research: Atmospheres*, 115, <https://doi.org/10.1029/2009JD012737>, 2010.
- Guillot, D., Rajaratnam, B., and Emile-Geay, J.: Statistical paleoclimate reconstructions via Markov random fields, *The Annals of Applied Statistics*, 9, 324–352, <https://doi.org/10.1214/14-AOAS794>, 2015.
- 525 Hakim, G. J., Emile-Geay, J., Steig, E. J., Noone, D., Anderson, D. M., Tardif, R., Steiger, N., and Perkins, W. A.: The Last Millennium Climate Reanalysis Project: Framework and First Results, *Journal of Geophysical Research: Atmospheres*, 121, 6745 – 6764, <https://doi.org/10.1002/2016JD024751>, 2016.
- 530 Hoyer, S. and Hamman, J.: xarray: N-D labeled arrays and datasets in Python, *Journal of Open Research Software*, 5, <https://doi.org/10.5334/jors.148>, 2017.



- Huang, B., Thorne, P. W., Banzon, V. F., Boyer, T., Chepurin, G., Lawrimore, J. H., Menne, M. J., Smith, T. M., Vose, R. S., and Zhang, H.-M.: Extended Reconstructed Sea Surface Temperature, Version 5 (ERSSTv5): Upgrades, Validations, and Intercomparisons, *Journal of Climate*, 30, 8179–8205, <https://doi.org/10.1175/JCLI-D-16-0836.1>, 2017.
- 535 Hunter, J. D.: Matplotlib: A 2D graphics environment, *Computing In Science & Engineering*, 9, 90–95, <https://doi.org/10.1109/MCSE.2007.55>, 2007.
- Jones, J. and Widmann, M.: Reconstructing Large-scale Variability from Palaeoclimatic Evidence by Means of Data Assimilation Through Upscaling and Nudging (DATUN), in: *The KIHZ project: Towards a Synthesis of Holocene Proxy Data and Climate Models*, edited by Fischer, H., Kumke, T., Lohmann, G., Flösser, G., Miller, H., von Storch, H., and Negendank, J., pp. 171–193, Springer, Heidelberg, Berlin, New York, 2004.
- 540 Kalman, R. E.: A New Approach to Linear Filtering and Prediction Problems, *Journal of Fluids Engineering*, 82, 35–45, <https://doi.org/10.1115/1.3662552>, 1960.
- Khider, D., Emile-Geay, J., Zhu, F., James, A., Landers, J., Ratnakar, V., and Gil, Y.: Pyleoclim: Paleoclimate Timeseries Analysis and Visualization With Python, *Paleoceanography and Paleoclimatology*, 37, e2022PA004 509, <https://doi.org/10.1029/2022PA004509>, 2022.
- 545 King, J., Tierney, J., Osman, M., Judd, E. J., and Anchukaitis, K. J.: DASH: A MATLAB Toolbox for Paleoclimate Data Assimilation, *Geoscientific Model Development*, 16, 5653–5683, <https://doi.org/10.5194/gmd-16-5653-2023>, 2023.
- King, J. M., Anchukaitis, K. J., Tierney, J. E., Hakim, G. J., Emile-Geay, J., Zhu, F., and Wilson, R.: A data assimilation approach to last millennium temperature field reconstruction using a limited high-sensitivity proxy network, *Journal of Climate*, -1, 1–64, <https://doi.org/10.1175/JCLI-D-20-0661.1>, publisher: American Meteorological Society Section: Journal of Climate, 2021.
- 550 Kluyver, T., Ragan-Kelley, B., Pérez, F., Granger, B., Bussonnier, M., Frederic, J., Kelley, K., Hamrick, J., Grout, J., Corlay, S., Ivanov, P., Avila, D., Abdalla, S., and Willing, C.: Jupyter Notebooks – a publishing format for reproducible computational workflows, in: *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, edited by Loizides, F. and Schmidt, B., pp. 87 – 90, IOS Press, 2016.
- Lauritzen, S. L.: *Graphical Models*, Oxford: Clarendon Press., 1996.
- Lenssen, N. J. L., Schmidt, G. A., Hansen, J. E., Menne, M. J., Persin, A., Ruedy, R., and Zyss, D.: Improvements in the GISTEMP Uncertainty Model, *Journal of Geophysical Research: Atmospheres*, 124, 6307–6326, <https://doi.org/10.1029/2018JD029522>, 2019.
- 555 Liu, H., Liu, Z., and Lu, F.: A Systematic Comparison of Particle Filter and EnKF in Assimilating Time-Averaged Observations, *Journal of Geophysical Research: Atmospheres*, p. 2017JD026798, <https://doi.org/10.1002/2017JD026798>, 2017.
- Lough, J. M.: Climate records from corals, *Wiley Interdisciplinary Reviews: Climate Change*, 1, 318–331, <https://doi.org/10.1002/wcc.39>, 2010.
- 560 Lyu, Z., Goosse, H., Dalaiden, Q., Klein, F., Shi, F., Wagner, S., and Braconnot, P.: Spatial Patterns of Multi-Centennial Surface Air Temperature Trends in Antarctica over 1–1000 CE: Insights from Ice Core Records and Modeling, *Quaternary Science Reviews*, 271, 107 205, <https://doi.org/10.1016/j.quascirev.2021.107205>, 2021.
- Malevich, S. B., Vetter, L., and Tierney, J. E.: Global Core Top Calibration of  $\delta^{18}\text{O}$  in Planktic Foraminifera to Sea Surface Temperature, *Paleoceanography and Paleoclimatology*, 34, 1292–1315, <https://doi.org/10.1029/2019PA003576>, 2019.
- 565 Mann, M. E., Bradley, R. S., and Hughes, M. K.: Global-scale temperature patterns and climate forcing over the past six centuries, *Nature*, 392, 779–787, <https://doi.org/10.1038/33859>, 1998.
- Mann, M. E., Bradley, R. S., and Hughes, M. K.: Northern hemisphere temperatures during the past millennium: Inferences, uncertainties, and limitations, *Geophysical Research Letters*, 26, 759–762, <https://doi.org/10.1029/1999GL900070>, 1999.

- Mann, M. E., Rutherford, S., Wahl, E., and Ammann, C.: Robustness of proxy-based climate field reconstruction methods, *Journal of Geophysical Research (Atmospheres)*, 112, 12 109–+, <https://doi.org/10.1029/2006JD008272>, 2007.
- Mann, M. E., Zhang, Z., Hughes, M. K., Bradley, R. S., Miller, S. K., Rutherford, S., and Ni, F.: Proxy-based reconstructions of hemispheric and global surface temperature variations over the past two millennia, *Proceedings of the National Academy of Sciences*, 105, 13 252–13 257, <https://doi.org/10.1073/pnas.0805721105>, 2008.
- Mann, M. E., Zhang, Z., Rutherford, S., Bradley, R. S., Hughes, M. K., Shindell, D., Ammann, C., Faluvegi, G., and Ni, F.: Global Signatures and Dynamical Origins of the Little Ice Age and Medieval Climate Anomaly, *Science*, 326, 1256–1260, <https://doi.org/10.1126/science.1177303>, 2009.
- Marcott, S. A., Shakun, J. D., Clark, P. U., and Mix, A. C.: A Reconstruction of Regional and Global Temperature for the Past 11,300 Years, *Science*, 339, 1198–1201, <https://doi.org/10.1126/science.1228026>, 2013.
- Masson-Delmotte, V., Schulz, M., Abe-Ouchi, A., Beer, J., Ganopolski, A., Rouco, J. G., Jansen, E., Lambeck, K., Luterbacher, J., Naish, T., Osborn, T., Otto-Bliesner, B., Quinn, T., Ramesh, R., Rojas, M., Shao, X., and Timmermann, A.: Information from Paleoclimate Archives, in: *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, edited by Stocker, T. F., Qin, D., Plattner, G.-K., Tignor, M., Allen, S., Boschung, J., Nauels, A., Xia, Y., Bex, V., and Midgley, P., pp. 383–464, Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, <https://doi.org/10.1017/CBO9781107415324.013>, 2013.
- McKinney, W.: Data Structures for Statistical Computing in Python, in: *Proceedings of the 9th Python in Science Conference*, edited by van der Walt, S. and Millman, J., pp. 51 – 56, 2010.
- Menne, M. J., Williams, C. N., Gleason, B. E., Rennie, J. J., and Lawrimore, J. H.: The Global Historical Climatology Network Monthly Temperature Dataset, Version 4, *Journal of Climate*, 31, 9835–9854, <https://doi.org/10.1175/JCLI-D-18-0094.1>, 2018.
- Met Office: Cartopy: a cartographic python library with a Matplotlib interface, Exeter, Devon, <https://scitools.org.uk/cartopy>, 2010.
- Morice, C. P., Kennedy, J. J., Rayner, N. A., and Jones, P. D.: Quantifying Uncertainties in Global and Regional Temperature Change Using an Ensemble of Observational Estimates: The HadCRUT4 Data Set, *Journal of Geophysical Research: Atmospheres*, 117, <https://doi.org/10.1029/2011JD017187>, 2012.
- Neukom, R., Barboza, L. A., Erb, M. P., Shi, F., Emile-Geay, J., Evans, M. N., Franke, J., Kaufman, D. S., Lücke, L., Rehfeld, K., Schurer, A., Zhu, F., Brönnimann, S., Hakim, G. J., Henley, B. J., Ljungqvist, F. C., McKay, N., Valler, V., and von Gunten, L.: Consistent multidecadal variability in global temperature reconstructions and simulations over the Common Era, *Nature Geoscience*, 12, 643–649, <https://doi.org/10.1038/s41561-019-0400-0>, 2019a.
- Neukom, R., Steiger, N., Gómez-Navarro, J. J., Wang, J., and Werner, J. P.: No evidence for globally coherent warm and cold periods over the preindustrial Common Era, *Nature*, 571, 550–554, <https://doi.org/10.1038/s41586-019-1401-2>, 2019b.
- Osman, M. B., Tierney, J. E., Zhu, J., Tardif, R., Hakim, G. J., King, J., and Poulsen, C. J.: Globally resolved surface temperatures since the Last Glacial Maximum, *Nature*, 599, 239–244, <https://doi.org/10.1038/s41586-021-03984-4>, 2021.
- Otto-Bliesner, B. L., Brady, E. C., Fasullo, J., Jahn, A., Landrum, L., Stevenson, S., Rosenbloom, N., Mai, A., and Strand, G.: Climate Variability and Change since 850 CE: An Ensemble Approach with the Community Earth System Model, *Bull. Am. Met. Soc.*, 97, 735–754, <https://doi.org/10.1175/BAMS-D-14-00233.1>, 2016.
- PAGES 2k Consortium: A global multiproxy database for temperature reconstructions of the Common Era, *Scientific Data*, 4, 170 088 EP, <https://doi.org/10.1038/sdata.2017.88>, 2017.

- Perkins, W. A. and Hakim, G. J.: Reconstructing paleoclimate fields using online data assimilation with a linear inverse model, *Climate of the Past*, 13, 421–436, <https://doi.org/10.5194/cp-13-421-2017>, 2017.
- Perkins, W. A. and Hakim, G. J.: Coupled Atmosphere–Ocean Reconstruction of the Last Millennium Using Online Data Assimilation, *Paleoceanography and Paleoclimatology*, 36, e2020PA003 959, <https://doi.org/10.1029/2020PA003959>, 2021.
- 610 Plotly Technologies Inc.: Collaborative data science, <https://plot.ly>, 2015.
- Rew, R. and Davis, G.: NetCDF: An Interface for Scientific Data Access, *IEEE Computer Graphics and Applications*, 10, 76–82, <https://doi.org/10.1109/38.56302>, 1990.
- Ridgwell, A., Hargreaves, J. C., Edwards, N. R., Annan, J. D., Lenton, T. M., Marsh, R., Yool, A., and Watson, A.: Marine geochemical data assimilation in an efficient Earth System Model of global biogeochemical cycling, *Biogeosciences*, 4, 87–104, <https://doi.org/https://doi.org/10.5194/bg-4-87-2007>, publisher: Copernicus GmbH, 2007.
- 615 Rutherford, S., Mann, M. E., Osborn, T. J., Bradley, R. S., Briffa, K. R., Hughes, M. K., and Jones, P. D.: Proxy-Based Northern Hemisphere Surface Temperature Reconstructions: Sensitivity to Method, Predictor Network, Target Season, and Target Domain., *J. Clim.*, 18, 2308–2329, 2005.
- Schneider, T.: Analysis of Incomplete Climate Data: Estimation of Mean Values and Covariance Matrices and Imputation of Missing Values., *J. Clim.*, 14, 853–871, [https://doi.org/10.1175/1520-0442\(2001\)014<0853:AOICDE>2.0.CO;2](https://doi.org/10.1175/1520-0442(2001)014<0853:AOICDE>2.0.CO;2), 2001.
- 620 Seabold, S. and Perktold, J.: statsmodels: Econometric and statistical modeling with python, in: 9th Python in Science Conference, 2010.
- Shakun, J. D., Clark, P. U., He, F., Marcott, S. A., Mix, A. C., Liu, Z., Otto-Bliesner, B., Schmittner, A., and Bard, E.: Global Warming Preceded by Increasing Carbon Dioxide Concentrations during the Last Deglaciation, *Nature*, 484, 49–54, <https://doi.org/10.1038/nature10915>, 2012.
- 625 Shi, F., Zhao, S., Guo, Z., Goosse, H., and Yin, Q.: Multi-Proxy Reconstructions of May–September Precipitation Field in China over the Past 500 Years, *Climate of the Past*, 13, 1919–1938, <https://doi.org/10.5194/cp-13-1919-2017>, 2017.
- Shi, F., Goosse, H., Klein, F., Zhao, S., Liu, T., and Guo, Z.: Monopole Mode of Precipitation in East Asia Modulated by the South China Sea Over the Last Four Centuries, *Geophysical Research Letters*, 46, 14 713–14 722, <https://doi.org/10.1029/2019GL085320>, 2019.
- Smerdon, J. E. and Pollack, H. N.: Reconstructing Earth’s surface temperature over the past 2000 years: the science behind the headlines, *WIREs Climate Change*, 7, 746–771, <https://doi.org/https://doi.org/10.1002/wcc.418>, 2016.
- 630 Steiger, N. and Hakim, G.: Multi-timescale data assimilation for atmosphere–ocean state estimates, *Climate of the Past*, 12, 1375–1388, <https://doi.org/https://doi.org/10.5194/cp-12-1375-2016>, publisher: Copernicus GmbH, 2016.
- Steiger, N. J., Hakim, G. J., Steig, E. J., Battisti, D. S., and Roe, G. H.: Assimilation of Time-Averaged Pseudoproxies for Climate Reconstruction, *Journal of Climate*, 27, 426–441, <https://doi.org/10.1175/JCLI-D-12-00693.1>, 2013.
- 635 Steiger, N. J., Hakim, G. J., Steig, E. J., Battisti, D. S., and Roe, G. H.: Assimilation of Time-Averaged Pseudoproxies for Climate Reconstruction, *Journal of Climate*, 27, 426–441, <https://doi.org/10.1175/JCLI-D-12-00693.1>, 2014.
- Steiger, N. J., Smerdon, J. E., Cook, E. R., and Cook, B. I.: A Reconstruction of Global Hydroclimate and Dynamical Variables over the Common Era, *Scientific Data*, 5, 1–15, <https://doi.org/10.1038/sdata.2018.86>, 2018.
- Tardif, R., Hakim, G. J., Perkins, W. A., Horlick, K. A., Erb, M. P., Emile-Geay, J., Anderson, D. M., Steig, E. J., and Noone, D.: Last Millennium Reanalysis with an expanded proxy database and seasonal proxy modeling, *Climate of the Past*, 15, 1251–1273, <https://doi.org/10.5194/cp-15-1251-2019>, 2019.
- Tierney, J. E. and Tingley, M. P.: A Bayesian, Spatially-Varying Calibration Model for the TEX86 Proxy, *Geochimica et Cosmochimica Acta*, 127, 83–106, <https://doi.org/10.1016/j.gca.2013.11.026>, 2014.

- Tierney, J. E. and Tingley, M. P.: BAYSPLINE: A New Calibration for the Alkenone Paleothermometer, *Paleoceanography and Paleoclimatology*, 33, 281–301, <https://doi.org/10.1002/2017PA003201>, 2018.
- 645 Tierney, J. E., Abram, N. J., Anchukaitis, K. J., Evans, M. N., Giry, C., Kilbourne, K. H., Saenger, C. P., Wu, H. C., and Zinke, J.: Tropical sea surface temperatures for the past four centuries reconstructed from coral archives, *Paleoceanography*, 30, 226–252, <https://doi.org/10.1002/2014PA002717>, 2015.
- Tierney, J. E., Malevich, S. B., Gray, W., Vetter, L., and Thirumalai, K.: Bayesian Calibration of the Mg/Ca Paleothermometer in Planktic Foraminifera, *Paleoceanography and Paleoclimatology*, 34, 2005–2030, <https://doi.org/10.1029/2019PA003744>, 2019.
- 650 Tierney, J. E., Zhu, J., King, J., Malevich, S. B., Hakim, G. J., and Poulsen, C. J.: Glacial Cooling and Climate Sensitivity Revisited, *Nature*, 584, 569–573, <https://doi.org/10.1038/s41586-020-2617-x>, 2020.
- Tingley, M. P. and Huybers, P.: Recent temperature extremes at high northern latitudes unprecedented in the past 600 years, *Nature*, 496, 201–205, <https://doi.org/10.1038/nature11969>, 2013.
- 655 Tingley, M. P. and Huybers, P. J.: A Bayesian Algorithm for Reconstructing Climate Anomalies in Space and Time. Part 1: Development and Applications to Paleoclimate Reconstruction Problems, *Journal of Climate*, <https://doi.org/10.1175/2009JCLI3015.1>, 2010a.
- Tingley, M. P. and Huybers, P. J.: A Bayesian Algorithm for Reconstructing Climate Anomalies in Space and Time. Part 2: Comparison with the Regularized Expectation-Maximization Algorithm, *Journal of Climate*, <https://doi.org/10.1175/2009JCLI3016.1>, 2010b.
- Tingley, M. P. and Li, B.: Comments on “Reconstructing the NH Mean Temperature: Can Underestimation of Trends and Variability Be Avoided?”, *Journal of Climate*, 25, 3441–3446, <https://doi.org/10.1175/JCLI-D-11-00005.1>, 2012.
- 660 Tingley, M. P., Craigmile, P. F., Haran, M., Li, B., Mannshardt, E., and Rajaratnam, B.: Piecing together the past: statistical insights into paleoclimatic reconstructions, *Quaternary Science Reviews*, 35, 1–22, <https://doi.org/10.1016/j.quascirev.2012.01.012>, 2012.
- Tolwinski-Ward, S. E., Evans, M. N., Hughes, M. K., and Anchukaitis, K. J.: An efficient forward model of the climate controls on interannual variation in tree-ring width, *Climate Dynamics*, 36, 2419–2439, <https://doi.org/10.1007/s00382-010-0945-5>, 2011.
- 665 Tolwinski-Ward, S. E., Anchukaitis, K. J., and Evans, M. N.: Bayesian parameter estimation and interpretation for an intermediate model of tree-ring width, *Climate of the Past*, 9, 1481–1493, <https://doi.org/10.5194/cp-9-1481-2013>, 2013.
- Vaccaro, A., Emile-Geay, J., Guillot, D., Verna, R., Morice, C., Kennedy, J., and Rajaratnam, B.: Climate field completion via Markov random fields – Application to the HadCRUT4.6 temperature dataset, *Journal of Climate*, -1, 1–66, <https://doi.org/10.1175/JCLI-D-19-0814.1>, publisher: American Meteorological Society Section: *Journal of Climate*, 2021.
- 670 Valler, V., Franke, J., Brugnara, Y., and Brönnimann, S.: An Updated Global Atmospheric Paleo-Reanalysis Covering the Last 400 Years, *Geoscience Data Journal*, 9, 89–107, <https://doi.org/10.1002/gdj3.121>, 2022.
- van der Walt, S., Colbert, S. C., and Varoquaux, G.: The NumPy Array: A Structure for Efficient Numerical Computation, *Computing in Science Engineering*, 13, 22–30, <https://doi.org/10.1109/MCSE.2011.37>, 2011.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., 675 van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, İ., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., and SciPy 1.0 Contributors: SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python, *Nature Methods*, 17, 261–272, <https://doi.org/10.1038/s41592-019-0686-2>, 2020.
- Wang, J., Emile-Geay, J., Guillot, D., McKay, N. P., and Rajaratnam, B.: Fragility of reconstructed temperature patterns over the Common 680 Era: Implications for model evaluation., *Geophysical Research Letters*, 42, 7162–7170, <https://doi.org/10.1002/2015GL065265>, 2015.
- Waskom, M. L.: seaborn: statistical data visualization, *Journal of Open Source Software*, 6, 3021, <https://doi.org/10.21105/joss.03021>, 2021.

- Widmann, M., Goosse, H., van der Schrier, G., Schnur, R., and Barkmeijer, J.: Using Data Assimilation to Study Extratropical Northern Hemisphere Climate over the Last Millennium, *Climate of the Past*, 6, 627–644, <https://doi.org/10.5194/cp-6-627-2010>, 2010.
- 685 Wikle, C. K. and Berliner, L. M.: A Bayesian tutorial for data assimilation, *Physica D: Nonlinear Phenomena*, 230, 1 – 16, <https://doi.org/10.1016/j.physd.2006.09.017>, data Assimilation, 2007.
- Zhu, F.: cfr-data v1.0, <https://doi.org/10.5281/zenodo.8367746>, 2023.
- Zhu, F. and Tolwinski-Ward, S.: PyVSL v0.4.0, <https://doi.org/10.5281/zenodo.7967103>, 2023.
- Zhu, F., Emile-Geay, J., Hakim, G. J., King, J., and Anchukaitis, K. J.: Resolving the Differences in the Simulated and Reconstructed Temperature Response to Volcanism, *Geophysical Research Letters*, 47, e2019GL086908, <https://doi.org/10.1029/2019GL086908>, 2020.
- 690 Zhu, F., Emile-Geay, J., Anchukaitis, K. J., Hakim, G. J., Wittenberg, A. T., Morales, M. S., Toohey, M., and King, J.: A re-appraisal of the ENSO response to volcanism with paleoclimate data assimilation, *Nature Communications*, 13, 747, <https://doi.org/10.1038/s41467-022-28210-1>, 2022.
- Zhu, F., Emile-Geay, J., Anchukaitis, K. J., McKay, N. P., Stevenson, S., and Meng, Z.: A Pseudoproxy Emulation of the PAGES 2k Database Using a Hierarchy of Proxy System Models, *Scientific Data*, 10, 624, <https://doi.org/10.1038/s41597-023-02489-1>, 2023a.
- 695 Zhu, F., Emile-Geay, J., Anchukaitis, K. J., McKay, N. P., Stevenson, S., and Meng, Z.: The pseudoPAGES2k Dataset v1.4, <https://doi.org/10.5281/zenodo.8173256>, 2023b.
- Zhu, F., Emile-Geay, J., Hakim, G. J., Guillot, D., Khider, D., Tardif, R., and Perkins, W. A.: cfr: a Python package for Climate Field Reconstruction, <https://doi.org/10.5281/zenodo.10575537>, 2024.