

A method for generating coherent spatially explicit maps of seasonal palaeoclimates from site-based reconstructions

Article

Published Version

Creative Commons: Attribution 4.0 (CC-BY)

Open Access

Cleator, S. F., Harrison, S. P. ORCID: <https://orcid.org/0000-0001-5687-1903>, Nichols, N. K. ORCID: <https://orcid.org/0000-0003-1133-5220>, Prentice, I. C. and Roulstone, I. (2020) A method for generating coherent spatially explicit maps of seasonal palaeoclimates from site-based reconstructions. *Journal of Advances in Modeling Earth Systems*, 12 (1). e2019MS001630. ISSN 1942-2466 doi: 10.1029/2019MS001630 Available at <https://centaur.reading.ac.uk/86621/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1029/2019MS001630>

Publisher: American Geophysical Union

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online



RESEARCH ARTICLE

10.1029/2019MS001630

Key Points:

- A 3-D-variational technique combining site-based estimates with a model prior is used to reconstruct seasonal paleoclimates
- The technique produces spatially and seasonally coherent temperature and moisture changes relative to present
- Climate reconstructions are provided for southern Europe at the Last Glacial Maximum, ca. 21000 years ago

Supporting Information:

- Data Set S1
- Data Set S2

Correspondence to:

S. F. Cleator,
s.cleator@surrey.ac.uk

Citation:

Cleator, S. F., Harrison, S. P., Nichols, N. K., Prentice, I. C., & Roulstone, I. (2020). A method for generating coherent spatially explicit maps of seasonal paleoclimates from site-based reconstructions. *Journal of Advances in Modeling Earth Systems*, 12, e2019MS001630. <https://doi.org/10.1029/2019MS001630>

Received 31 JAN 2019

Accepted 26 SEP 2019

Accepted article online 14 OCT 2019

©2019. The Authors.

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

A Method for Generating Coherent Spatially Explicit Maps of Seasonal Paleoclimates From Site-Based Reconstructions

S. F. Cleator¹, S. P. Harrison², N. K. Nichols³, I. C. Prentice⁴, and I. Roulstone¹

¹Department of Mathematics, University of Surrey, Guildford, UK, ²School of Archaeology, Geography and Environmental Sciences (SAGES), University of Reading, Reading, UK, ³Department of Mathematics & Statistics, University of Reading, Reading, UK, ⁴AXA Chair of Biosphere and Climate Impacts, Department of Life Sciences, Imperial College London, London, UK

Abstract We describe a method for reconstructing spatially explicit maps of seasonal paleoclimate variables from site-based reconstructions. Using a 3-D-Variational technique, the method finds the best statistically unbiased, and spatially continuous, estimate of the paleoclimate anomalies through combining the site-based reconstructions and a prior estimate of the paleoclimate state. By assuming a set of correlations in the error of the prior, the resulting climate is smoothed both from month to month and from grid cell to grid cell. The amount of smoothing can be controlled through the choice of two length-scale values. The method is applied to a set of reconstructions of the climate of the Last Glacial Maximum (ca. 21,000 years ago) for southern Europe derived from pollen data with a prior derived from results from the third phase of the Palaeoclimate Modelling Intercomparison Project. We demonstrate how to choose suitable values for the smoothing length scales for the data sets used in the reconstruction.

1. Introduction

Past climates provide useful examples of how the climate system has responded to changes in external forcing, such as orbitally induced changes in incoming solar radiation, and internal Earth system feedbacks, such as changes in atmospheric CO₂ concentration ([CO₂]) or ice sheet extent (Harrison & Bartlein, 2012). Reconstructions of past climate states are now routinely used to evaluate the performance of the climate models that are used to project the trajectory of future climate changes (Harrison et al., 2014, 2015; Kageyama et al., 2018; Schmidt et al., 2014). The Last Glacial Maximum (LGM; ca. 21,000 years ago) has been a major focus for these evaluations because the change in climate forcing was as large (albeit different in type) as “high-end” changes projected for the end of the 21st century (Braconnot et al., 2012; Kageyama et al., 2018). These evaluations obviously depend on the availability of quantitative reconstructions of key climate variables, and this has led to the creation of benchmark data sets documenting climate conditions over land (e.g., Bartlein et al., 2011) and ocean (e.g., MARGO Project Members, 2009).

Past climate conditions can be inferred from environmental records which respond to climate, including sedimentological, geomorphological, chemical, isotopic, and biological records (Bradley, 1999; Gornitz, 2008). Quantitative reconstructions of climate variables can be obtained from these records either using statistical techniques based on modern day climate-response relationships (e.g., Ter Braak & Juggins, 1993; see also discussion in Bartlein et al., 2011) or by inversion of a model that simulates the response of a particular type of environmental record to climate (e.g., Garreta et al., 2010; Steiger et al., 2017). Pollen preserved in anoxic lake and bog sediments through time is the most widespread source of data for the reconstruction of terrestrial climates (Bartlein et al., 2011; Marsicek et al., 2018), because pollen abundance reflects the distribution of different plant taxa that have highly specific climatic requirements (Harrison et al., 2010; Woodward, 1987) and the pollen-preserving sediments can be accurately dated using radiometric techniques. One important characteristic of all of the environmental records that are used for climate reconstruction, including pollen, is that both the primary data and the climate reconstructions are generated for individual sites. Geological and climatic factors mean that the distribution of potential sites is spatially nonuniform: Speleothem records, for example, are confined to karst areas; pollen preservation requires anoxic conditions, and thus, pollen records are not common in arid regions. Furthermore, issues of accessibility and scientific interests

means that the actual sampling of potential sites is nonuniform, so there are often large geographic gaps in the data coverage (Harrison et al., 2016). While pollen records, for example, are well sampled across Europe and North America, there are far fewer records from central Eurasia or the tropics. Furthermore, geological preservation issues mean that the number of sites available tends to decrease through time: There is an order of magnitude more pollen data available for the middle Holocene (ca. 6,000 yr BP) than for the LGM (e.g., Harrison et al., 2016). Ideally, a benchmark data set for model evaluation would provide continuous climate fields. However, while gridding the data sets at a scale comparable to that of the climate models (see, e.g., Bartlein et al., 2011) can improve the situation, this still does not solve the problem of significant gaps in site-based data coverage.

Alternative approaches to generating spatially continuous paleoclimate reconstructions have been developed that involve combining observations with model simulations of paleoclimates. Goosse et al. (2006), for example, used observations to select the most realistic member from an ensemble of climate-model simulations. They ran a relatively large ensemble of simulations using a range of different climate forcings and/or model parametrizations to encompass uncertainties and then selected the members of the ensemble that best matched the observations at each time step before running these simulations for longer to gain a new estimate of the climate. In this approach, the most realistic climate is taken to be the simulated climate(s) that best matched observations after multiple simulations. Although this approach provides continuous and self-consistent fields of climate variables, the reconstructions cannot deviate fundamentally from the model predictions and thus could still be influenced by systematic errors inherent in the model construction. Annan and Hargreaves (2013) also used an ensemble of model simulations, but in this case they used multiple models. The ultimate climate reconstruction was assumed to be a weighted average of those climate models, where the weighting was determined by the goodness of fit to the observations. They applied a global weighting to each model rather than allowing the goodness of fit to vary regionally. As a result, there are regions where the reconstructed paleoclimate is far from the observations, producing a paleoclimate reanalysis that is highly influenced by systematic errors in the models.

Variational data assimilation techniques provide a way of combining observations and model outputs to produce climate reconstructions that are not explicitly constrained to a given source (Lahoz & Schneider, 2014; Nichols, 2010). Variational techniques are widely used by the weather forecasting community (e.g., Daley, 1994) and have also been used to reconstruct paleoclimate. Gebhardt et al. (2008), for example, used this approach to reconstruct European climates during the Last Interglacial. Simonis et al. (2012) applied the same basic approach to reconstruct January and July temperatures across Europe during the late Glacial (13,000 yr BP) and early Holocene (8,000 yr BP). The method involves applying a spatial constraint, based on a two-dimensional advection-diffusion equation of atmospheric dynamics, to upscale climate variables derived from statistical transfer functions relating the abundance of plant taxa with January and July temperature. In both examples, modern-day wind fields were used as the prior to determine the spatial scale and assumed to be the same in the past.

Tardif et al. (2018) also use variational techniques to create paleoclimate reconstructions for the Last Millennium, using an ensemble of transient paleoclimate simulations. They first determine the relationship between paleoclimate reconstructions and the model-derived prior using linear regression and then determine the best linear unbiased estimate using the Kalman formulation, to create the analytical reconstructions. Thus, temporal relationships are not based on an explicit analytical function designed to preserve structures (autocorrelations and/or discontinuities) in the prior and/or observations. Spatial correlation is generated from the prior ensemble with a covariance localization applied to prevent spurious correlations. This, and the necessity to define scaling parameters, involves a number of arbitrary choices which influence the final reconstructions and make it difficult for these reconstructions to deviate substantially from the prior.

The 3-D-Variational method finds the maximum a posteriori Bayesian estimate of the paleoclimate given the site-based reconstructions and a prior estimate. While this could lead to the generation of reconstructions with sharp changes in time and/or space, it is possible to incorporate additional assumptions about the error of the prior estimate (the difference between the true climate and the prior) to prevent this by ensuring continuity of the solution. The degree of continuity in the change of the reconstructed climate field can be controlled by adjusting two length scales: a spatial length scale that determines how smooth the spatial

correlation in the prior is between different geographical areas and a temporal length scale that determines how smooth it is through the seasonal cycle.

Here we apply this method to reconstruct six paleoclimate variables across southern Europe at the LGM. The six climate variables are those provided in the Bartlein et al. (2011) data set, namely, mean annual temperature (MAT, °C), mean temperature of the coldest month (MTCO, °C), mean temperature of the warmest month (MTWA, °C), growing degree days above a baseline of above 5 °C (GDD5, d°C), mean annual precipitation (MAP, mm), and an index of plant-available moisture (the ratio of actual to equilibrium evapotranspiration or α in Bartlein et al., 2011, reexpressed as a moisture index [MI, unitless] defined as the ratio of MAP to equilibrium evapotranspiration in our analyses. The conversion was made using the Zhang et al., 2004, formulation of the Budyko relationship). We use pollen-based reconstructions of climatic variables for the region of southern Europe (defined here as south of 50° N and extending eastward to 50° E) from Bartlein et al. (2011) as our observations. Although the sites from Europe were used to produce a gridded map in Bartlein et al. (2011), here we used the underlying individual site reconstructions. Some of the reconstructions used in Bartlein et al. (2011) were derived by model inversion, and these were excluded from our data set. Bartlein et al. (2011) gives mean values as anomalies from the modern climate, as well as standard errors. We use eight LGM climate simulations (CCSM4, CNRM-CM5, MPI-ESM-P, MRI-CGCM3, FGOALS-g2, COSMOS-ASO, IPSL-CM5A-LR, and MIROC-ESM) from the third phase of the Palaeoclimate Modelling Intercomparison Project (PMIP3: Braconnot et al., 2012) to create a prior. These simulations were forced by changes in incoming solar radiation, changes in land-sea geography, and the size and extent of ice sheets, and a reduction in atmospheric [CO₂] (see Braconnot et al., 2012, for details of the modeling protocol).

Our approach introduces features novel to the field of paleoclimate data assimilation, explicitly designed to maximize the usefulness of the reconstructions for climate model evaluation. Specifically, by solving the full variational problem, we take into account nonlinearities in the system. Furthermore, we minimize the dependency of the final analytical reconstructions on the prior generated from the climate models by using a prescribed correlation function for the error of the prior and by using a resolution matrix (Delahaies et al., 2017; Menke, 2012) to determine the temporal correlation length scale. The resolution matrix provides a particularly useful way to overcome problems caused by the sparsity of site-based paleoclimate reconstructions at the LGM. In addition to investigating methods to determine appropriate spatial and temporal length scales, we provide a way of calculating the error in the final reconstructions.

2. Data Assimilation With Spatial and Temporal Correlations in the Prior

In this section we describe the underlying method used in this paper. Section 2.1 describes the inverse problem solved by the method and the types of data used. Section 2.2 shows how the different climate variables can be related to one another by specifying correlations from our prior estimate of the system. Finally section 2.3 describes how the problem is preconditioned in order to reduce the computation cost.

2.1. The Inverse Problem

Our problem is to determine the paleoclimate that existed from a particular set of site-based reconstructions. We label the reconstructions as the column vector $\mathbf{y}_i \in \mathbb{R}^6$ for site i . For each reconstruction, $\mathbf{y}_i$, there are a total of six variables that may have been reconstructed, namely, α , MAP, MAT, MTCO, MTWA, and GDD5. All these reconstructions together make the observations labeled $\mathbf{y} \in \mathbb{R}^{6N}$ such that

$$\mathbf{y} = (\mathbf{y}_1^T | \mathbf{y}_2^T | \dots | \mathbf{y}_N^T)^T, \quad (1)$$

where N is the number of reconstructions. The reconstruction technique gives the variances for each reconstruction that we label as the column vector $\mathbf{v}_y \in \mathbb{R}^{6N}$ in the same order as $\mathbf{y}$. Not all variables are reconstructed at every site; for these variables, the variance tends to infinity; this is achieved by setting their inverse to 0.

From these reconstructions, we want to produce a gridded climate, the state vector, $\mathbf{x} \in \mathbb{R}^{13M}$ where there are M grid cells. The j th grid cell of the state is labeled $\mathbf{x}_j \in \mathbb{R}^{13}$ where

$$\mathbf{x} = (\mathbf{x}_1^T | \mathbf{x}_2^T | \dots | \mathbf{x}_M^T)^T. \quad (2)$$

For each grid cell, the method determines a set of 13 variables: the MAP (P) and the 12 average temperatures for each month, $\mathbf{T}$ where

$$\mathbf{T} = (T_1 \ T_2 \ \dots \ T_{12})^T,$$

where T_m is the temperature at month m .

For a general function $\mathbf{h}$ that maps a gridded climate $\mathbf{x}$ to the site-based observations, we state the problem as trying to find an $\mathbf{x}$ such that

$$\mathbf{h}(\mathbf{x}) = \mathbf{y}. \quad (3)$$

Solving equation (3) for $\mathbf{x}$ is ill-posed as there are several $\mathbf{x}$ that are possible solutions. A prior estimate of the state called the background or prior ($\mathbf{x}_b$) allows us to find the best $\mathbf{x}$ that solves equation (3) and remains close to the prior. The standard deviations of the prior are labeled as the vector $\mathbf{v}_b \in \mathbb{R}^{13M}$ in the same order as $\mathbf{x}_b$.

It can be shown (Nichols, 2010) that the optimal solution of equation (3) with a prior estimate of the state is defined as the analysis, $\mathbf{x}_a$, where

$$\mathbf{x}_a = \arg \min_{\mathbf{x}} J(\mathbf{x}), \quad (4)$$

with the cost function J as

$$J(\mathbf{x}) = \frac{1}{2}(\mathbf{x} - \mathbf{x}_b)^T \mathbf{B}^{-1}(\mathbf{x} - \mathbf{x}_b) + \frac{1}{2}(\mathbf{y} - \mathbf{h}(\mathbf{x}))^T \mathbf{R}^{-1}(\mathbf{y} - \mathbf{h}(\mathbf{x})). \quad (5)$$

Here $\mathbf{B}$ is the covariance of the uncertainties in the prior (conventionally denoted $\mathbf{B}$, for background), and $\mathbf{R}$ is the covariance of the uncertainties in the site-based reconstructions. Equations (4) and (5) ensure that the solution is the optimal distance from the observations subject to ensuring that the solution is not too far from the prior estimate, weighted by the error statistics in each. We assume that there are no correlations in the errors of the observations so we set

$$\mathbf{R} = \text{diag}(\mathbf{v}_y).$$

The prior error covariance matrix can be represented as the product of the standard deviations of the prior and the correlations between the errors in the variables in the prior. Hence, we write

$$\mathbf{B} = \Sigma \mathbf{C} \Sigma, \quad (6)$$

where

$$\Sigma = \text{diag}(\mathbf{v}_b^{\frac{1}{2}}) \quad (7)$$

is the diagonal matrix formed of the standard deviations of the prior error and $\mathbf{C}$ is the prior error correlation matrix.

2.2. Prior Error Correlation

The difference between the true $\mathbf{x}$ and the prior, the error in the prior, is expected to be smooth between adjacent grid cells and also from month to month since it would be unlikely that the observations would contain sharp jumps in climate that are not present in the prior. It would be unusual, for example, to have very high temperature in March if the temperatures in February and April are very low, if this behavior is not seen in the prior. To achieve this, we impose a structure on the prior error correlation matrix, $\mathbf{C}$, that weighs the cost function so that its minimum is smooth. This allows the prior error to be smooth but still allows nonsmooth areas if there is significant evidence to support it in the prior and/or the observations.

We assume there are two independent sets of correlations in the prior. The first correlation is spatially between the different grid cells. We also assume that the spatial correlation between the grid cells is homogeneous and valid on a sphere, so that for an angle θ_{ij} on a great circle of the Earth between each cell i and j , the correlation is given by

$$c_L(\theta_{ij}) = \left(\frac{a}{L} \sin \left(\frac{\theta_{ij}}{2} \right) \right) \mathcal{K}_1 \left(\frac{a}{L} \sin \left(\frac{\theta_{ij}}{2} \right) \right), \quad (8)$$

where c_L is a case of a Matérn function (Handcock & Wallis, 1994; Matérn, 1986) with order 1 and $\mathcal{K}$ is the modified Bessel function of the second kind, evaluated using the boost C++ library (Maddock et al., 2018). Here the correlation length scale is $L = L_s$ and $a = 6,371$ km is the radius of the Earth. The correlation matrix between all grid cells, $\mathbf{C}_{L_s}$, is given as

$$(\mathbf{C}_{L_s})_{ij} = c_{L_s}(\theta_{ij}).$$

The choice of L_s is dependent on the data sets used in $\mathbf{y}$ and $\mathbf{x}_b$ and so is specific to each problem. In section 3.2, a method of finding L_s is shown for a particular experiment.

The second assumed correlation is between the error in the average temperatures of the prior. We assume that there is a correlation between the average temperatures of a month and the surrounding months given by equation (8). Here $\theta_{ij} = \text{mod}_{12}(|i - j|)$ between months i and j . The correlation length scale is $L = L_t$ and $a = 6/\pi$. The appropriate value of L_t again depends on the data sets used and is shown for a particular experiment in section 3.2. For each grid cell, the correlation between the different climate variables is given by $\mathbf{C}_{L_t}$ where

$$\mathbf{C}_{L_t} = \left(\begin{array}{c|ccc} 1 & 0 & \dots & 0 \\ \hline 0 & & & \\ \vdots & & \{c_{L_t}(\theta_{ij})\}_{ij} & \\ 0 & & & \end{array} \right). \quad (9)$$

Note how $\{c_{L_t}(\theta_{ij})\}_{ij}$ is offset by the first row and column due to the presence of the precipitation term which is assumed to be uncorrelated to the temperature terms.

These two sets of correlations imply that all the variables in the error of the prior are correlated. For instance, the grid cells i and j are correlated by $(\mathbf{C}_{L_s})_{ij}$ and the temperatures in month l and k are correlated by $(\mathbf{C}_{L_t})_{lk}$. This means that the temperatures in month l in grid cell i and month k in grid cell j are correlated by the product $(\mathbf{C}_{L_s})_{ij}(\mathbf{C}_{L_t})_{lk}$. Repeating this for every variable gives an overall correlation for the prior ($\mathbf{C}$ from equation (6)) as

$$\mathbf{C} = \mathbf{C}_{L_s} \otimes \mathbf{C}_{L_t}, \quad (10)$$

where $\otimes$ is the Kronecker product of matrices.

The incorporation of correlation structures is due to the fact that the state space covers space and time. We introduce the $\mathbf{C}_{L_s}$ and $\mathbf{C}_{L_t}$ correlations to make the prior error smooth in space and time respectively. The presence of the scales L_s and L_t allows the adjustment of the smoothing in both dimensions and should depend, at least in part, on the spatial and temporal distribution of the prior and site-based reconstructions. In section 3.2, we discuss methods for choosing these scales.

2.3. Preconditioning and the Condition Number

The minimum of the cost function is sensitive to change in the input data used and to computational errors. This sensitivity reflects the difficulty in solving the problem and is measured by the condition number of the Hessian of the cost function (Golub & Loan, 1996). We define the condition number κ of a symmetric positive definite matrix $\mathbf{M}$ to be

$$\kappa(\mathbf{M}) = \frac{\lambda_{\max}(\mathbf{M})}{\lambda_{\min}(\mathbf{M})}, \quad (11)$$

where $\lambda_{\max}(\mathbf{M})$ and $\lambda_{\min}(\mathbf{M})$ are the maximum and minimum eigenvalues of $\mathbf{M}$. Here, $\mathbf{M}$ is the Hessian of the cost function, given by its (first order) second derivative $\mathbf{S} = \mathbf{H}\mathbf{B}\mathbf{H}^T + \mathbf{R}$ where $\mathbf{H}$ is the Jacobian of $\mathbf{h}$. This condition number indicates the computational work needed to minimize the cost function. Equation (11) shows how the condition number of $\mathbf{S}$ represents the disparity in scales of the problem. As the eigenvalues represent the sizes of the scales of $\mathbf{S}$, their ratio represents the largest scale that will be encountered when inverting $\mathbf{S}$. Since large-scale differences create more numerical inaccuracy, a large condition number will increase the computational cost and lead to an inaccurate solution.

Haben et al. (2011) shows that the bounds on the condition number can be reduced by minimizing the cost function around $\mathbf{w}$ instead of $\mathbf{x}$ where

$$\mathbf{B}^{\frac{1}{2}} \mathbf{w} = \mathbf{x} - \mathbf{x}_b, \quad (12)$$

where $\mathbf{B}^{\frac{1}{2}}$ is the symmetric square root of the matrix $\mathbf{B}$ such that

$$\mathbf{B} = \mathbf{B}^{\frac{1}{2}} \mathbf{B}^{\frac{1}{2}}.$$

The use of this linear transformation can be thought of as a z score to work with uncorrelated states.

Equation (12) transforms the inverse problem from equation (4) into finding

$$\mathbf{w}_a = \operatorname{argmin}_{\mathbf{w}} J(\mathbf{w}), \quad (13)$$

where $J(\mathbf{w})$ is

$$J(\mathbf{w}) = \frac{1}{2} \mathbf{w}^T \mathbf{w} + \frac{1}{2} (\mathbf{y} - \mathbf{h}(\mathbf{x}_b + \mathbf{B}^{\frac{1}{2}} \mathbf{w}))^T \mathbf{R}^{-1} (\mathbf{y} - \mathbf{h}(\mathbf{x}_b + \mathbf{B}^{\frac{1}{2}} \mathbf{w})). \quad (14)$$

We use the limited memory Broyden-Fletcher-Goldfarb-Shanno method to find the state, $\mathbf{w}_a$, which has the minimum J . The limited memory Broyden-Fletcher-Goldfarb-Shanno is a quasi-Newton method that maintains a limited memory version of an approximated Hessian as described in Liu and Nocedal (1989). At each evaluation step, we calculate the gradient of J as

$$\nabla J(\mathbf{w}) = \mathbf{w} - \mathbf{B}^{\frac{1}{2}} \mathbf{H}_x^T \mathbf{R}^{-1} (\mathbf{y} - \mathbf{h}_u(\mathbf{x}_b + \mathbf{B}^{\frac{1}{2}} \mathbf{w})), \quad (15)$$

where $\mathbf{H}_x$ is the Jacobian of $\mathbf{h}$ evaluated at $\mathbf{x}$. Once $\mathbf{w}_a$ is found we use equation (12) to find $\mathbf{x}_a$, the solution.

The error covariance of the analysis, $\mathbf{x}_a$, is denoted by $\mathbf{A}$ and is given (to first order) following Nichols (2010) as

$$\mathbf{A} = (\mathbf{I} - \mathbf{K} \mathbf{H}_{x_b}) \mathbf{B}, \quad (16)$$

where the gain matrix $\mathbf{K}$ is

$$\mathbf{K} = \mathbf{B} \mathbf{H}_{x_b}^T (\mathbf{H}_{x_b} \mathbf{B} \mathbf{H}_{x_b}^T + \mathbf{R})^{-1}. \quad (17)$$

3. Experimental Design

We use our method to reconstruct the paleoclimate of southern Europe during the LGM. The LGM had insolation forcing relatively similar to the present day, but Northern Hemisphere ice sheets were more extensive, sea level was lower and the area of the continents therefore larger, and the atmospheric $[\text{CO}_2]$ was less than half of the concentration today. In this section we describe the choices of $\mathbf{h}$, $\mathbf{y}$, and $\mathbf{x}_b$ used to make this reconstruction and our choices for L_t and L_s , the correlation length scales.

3.1. Experiment Setup

We use pollen-based reconstructions of climatic variables from Bartlein et al. (2011) as our observations. Bartlein et al. (2011) gives means as anomalies from the modern climate as well as standard errors. We add the anomalies to the CRU CL v2.0 data set (New et al., 2002) to derive absolute climate reconstructions. We nondimensionalize the climate variables in order to avoid computational issues in the calculation of the cost function because they are on different scales. After solving for the nondimensionalized case, we redimensionalize each of the variables to be on the original scale of the observations and the prior. Details of the dimensionalization and nondimensionalization of variables can be found in Appendix A. We use the nondimensionalized variables as our $\mathbf{y}$ and their nondimensional standard errors, formed from the product of the standard errors and the derivative of D_y (equation (A1)), as $\mathbf{v}_y^{\frac{1}{2}}$.

We use the LGM outputs from PMIP3 as our prior. We use the variables of monthly precipitation (that are summed to annual precipitation), monthly temperature, and monthly total cloud fraction. For each of the selected PMIP models that ran an LGM experiment, we interpolate the output to a $2^\circ \times 2^\circ$ grid producing a set of maps all at the same resolution. In order to minimize the impact of potential individual systematic model

biases in the simulated climate at the LGM, experiments are generally expressed relative to that specific model's preindustrial control (PI) experiment. We therefore interpolate each of the PI experiments to the same grid and take the difference between the LGM and PI experiments of each model as the anomaly to the modern day for each model. We then add each model's anomalies to values from the modern day (from CRU CL v2.0, as above, bilinearly interpolated to the $2^\circ \times 2^\circ$ grid) in order to produce absolute values for each model. For each variable in the set, we take the mean and variance across the set of all models to produce a gridded map. As for the observation space, we nondimensionalize the state space to remove any dimensional effects using equation (A3). The nondimensional variables form the prior $\mathbf{x}_b$ and their nondimensional variances, formed from taking the product of the variances and the derivative of equation (A3), form $\mathbf{v}_{x_b}$.

The observation function, $\mathbf{h}$, links together the variables from both data sets. At each site, i , we define the observation function as

$$\hat{\mathbf{h}}(\mathbf{x}_i) = \hat{\mathbf{h}} \left(\begin{array}{c} \bar{P} \\ \bar{\mathbf{T}} \end{array} \right) = \begin{pmatrix} \mu(\mathbf{x}_i) \\ \bar{P} \\ \text{mean}(\bar{\mathbf{T}}) \\ \text{max}(\bar{\mathbf{T}}) \\ \text{min}(\bar{\mathbf{T}}) \\ G(\bar{\mathbf{T}}) \end{pmatrix}. \quad (18)$$

The derivatives, $\frac{\partial \text{max}(T)}{\partial T_m}$ and $\frac{\partial \text{min}(T)}{\partial T_m}$ are taken to be 1 if T_m is the maximum or minimum of T and 0 elsewhere. The MI function μ is

$$\mu(\mathbf{x}_i) = 1 + m(\mathbf{x}_i) - (1 + m(\mathbf{x}_i)^\omega)^{\frac{1}{\omega}} \quad (19)$$

as given by the Budyko curve with $\omega = 3$ as described in Zhang et al. (2004). The MI m is calculated as

$$m(\mathbf{x}_i) = P\lambda \left[\sum_k^{12} l_k \frac{R(T_k, S_k) \frac{\partial e_s}{\partial T} |_{T_k}}{\frac{\partial e_s}{\partial T} |_{T_k} + \gamma} \right]^{-1} \quad (20)$$

where γ (0.067 kPa K^{-1}) is the psychrometer constant at sea level, l_j is the length of month j in days and where

$$\frac{\partial e_s}{\partial T} = \frac{10.5485}{(237.3 + T)^2} \exp \left(\frac{17.27T}{237.3 + T} \right)$$

is the differentiated Roche-Magnus formula from Allen et al. (1998). The function $R(T_k, S_k)$ is the daily net radiation at the vegetated surface defined in Davis et al. (2017) for the middle day in month k . The variable S_k is the total cloud fraction for month j which is taken from the PMIP3 average described above. We define

$$G(\bar{\mathbf{T}}) = \frac{1}{N_y} \sum_k^{12} \begin{cases} l_k \left(\bar{T}_k - \frac{5}{T_s} \right) & \bar{T}_k > \frac{5}{T_s} \\ 0 & \text{else} \end{cases},$$

and the mean function to be $\text{mean}(\bar{\mathbf{T}}) = \frac{1}{N_y} \sum_k^{12} l_k \bar{T}_k$ and $\text{max}(\bar{\mathbf{T}})$ and $\text{min}(\bar{\mathbf{T}})$ to be the maximum and minimum temperature in $\bar{\mathbf{T}}$, respectively. The full observation function, $\mathbf{h}$, is formed by applying $\hat{\mathbf{h}}$ at each grid cell where there is an observation and defining

$$\mathbf{h}(\mathbf{x}) = \left(\hat{\mathbf{h}}(\mathbf{x}_1)^T | \hat{\mathbf{h}}(\mathbf{x}_2)^T | \dots \right),$$

and so $\mathbf{h}$ will have the dimension $6N$, and hence $\mathbf{H}$ the Jacobian of $\mathbf{h}$, will have dimension $(6N)^2$.




Figure 1. Yearly temperature for the assimilation performed on a single simulated site at N37.50° and E33.73° with varying values of L_t . The different colored dots are the results of the assimilation for different values of L_t . The black dots in the center are the prior for the grid cell that contains this site with error bars of 1 standard deviation. The B-spline interpolation of the dots is shown as the curved lines. The observations of MTWA and MTCO are represented by the higher and lower solid black lines, respectively, with the dotted lines showing 1 standard deviation around the mean.

3.2. Determining L_t and L_s

The two correlation length scales, L_t and L_s , in **C** (section 2.2) determine the strength of the correlation in the errors of the prior. By varying the length scales, we can vary how smooth the error of the prior is and hence how smooth the solution is. If the length scale is too large, then the error will be oversmoothed and the solution will miss smaller-scale features such as interannual temperature changes or spatially small features caused by topography. A length scale too small will mean that the solution will be too coarse and contain unrealistic jumps.

In order to determine a suitable value for L_t , we consider a single grid cell with a single simulated observation at 37.50°N and 33.73°E, which allows us to ignore the effects of $\mathbf{C}_{L_s}$. The example only has observations of MTCO and MTWA (−15 and 30 °C, respectively), allowing us to ignore the nonlinear effects of calculating α . Figure 1 shows the prior and observations for the sample as well as the estimated states after assimilation for different values of L_t . For all values of L_t , the analysis does not match the observed MTCO since the prior temperature for January has low uncertainty. Low values of L_t create an analysis that swaps between the prior and the observations. Although the solution always matches either the reconstructions or the prior, the jumps between them are unrealistic. On the other hand, high values of L_t create an analysis that follows the prior too closely and is unable to create high and low temperatures. The value of $L_t = 1$ produces an assimilation that follows the shape of the prior but lies between the values of the prior and the observations.

We can further understand L_t by seeing how information is changed by the method. If we consider the hypothetical, true solution to the inverse problem, $\mathbf{w}_t$, then by equation (3) we have that

$$\mathbf{H}_{\mathbf{x}_b} \mathbf{B}^{\frac{1}{2}} \mathbf{w}_t \approx \mathbf{y} - \mathbf{h}(\mathbf{x}_b),$$

since, up to first order,

$$\mathbf{H}_{\mathbf{x}_b}(\mathbf{x} - \mathbf{x}_b) \approx \mathbf{h}(\mathbf{x}) - \mathbf{h}(\mathbf{x}_b). \quad (21)$$

Further, Nichols (2010) shows how

$$\mathbf{x}_a - \mathbf{x}_b \approx \mathbf{K}(\mathbf{y} - \mathbf{h}(\mathbf{x}_b)),$$




Figure 2. The resolution matrices for the assimilation method with a sample single grid cell and a simulated observation at N37.50° and E33.73°. The color is the log value of the resolution matrix N for values of $L_t = 0.1, 1$, and 2 .

where $\mathbf{K}$ is the gain matrix defined in equation (17). Hence, we can consider the change from the true solution to our computed one ($\mathbf{w}_a$) as being given by

$$\mathbf{w}_a \approx \mathbf{N}\mathbf{w}_t,$$

where

$$\mathbf{N} = \mathbf{B}^{-\frac{1}{2}} \mathbf{K} \mathbf{H}_{\mathbf{x}_b} \mathbf{B}^{\frac{1}{2}}$$

is the resolution matrix as described in Menke (2012) and Delahaies et al. (2017).

Resolution matrices where the diagonal elements are close to 0 describe a situation where if perfect information is input, then the solution would only contain part of this information. In situations where the resolution matrix has large off-diagonal terms, the solution is degraded by interference between variables. If the opposite is true, the resolution matrix is close to the identity matrix. The best method will have a resolution matrix that resolves as many variables as possible while having few variables interfering with each other.

Figure 2 shows how the resolution matrix changes with respect to L_t for the same test grid cell as in Figure 1. The simulated prior temperatures are closest to the observations in January and July such that for small




Figure 3. The condition number of our example problem as a function of L_s , the spatial length scaling.




Figure 4. The result, $\mathbf{h}(\mathbf{x}_a)$, is dimensionalized and represented by the color field with the dots representing the observations ($\mathbf{y}$). Observations of α have been translated to moisture index through equation (19).

values of L_t , the method resolves temperatures in these months well. However, for large L_t , the method improves the patterns away from these months while degrading reconstructions of January and July. Values of L_t in between the large and small values show a mixture of both high resolution and low interference. These results together with the results from Figure 1 suggest a value of $L_t = 1$ is suitable for this problem.

The choice of the other scale, L_s , is especially relevant for the relatively sparse data set used here. A higher L_s represents errors in the prior being correlated even though they are far away, whereas a low L_s represents errors not being highly correlated even though they are close together. A large L_s means that information from the reconstructions could be propagated over a large distance. While this is useful in maximizing the use of a geographically sparse data set, it could be unrealistic if this propagation extends too far beyond the source area for the pollen on which the site reconstructions are based (which is generally, though not always, of the order of 20–100 km around the site). In order to obtain a realistic solution while maximizing the use of the data, we choose L_s such that the assumed average source area of the different sites does not overlap.

L_s corresponds to the area that each observation impacts, so an increase in L_s makes greater use of the observations. Haben et al. (2011) show that the condition number of the inverse problem is proportional to the distance between the reconstruction sites which, in this case, is proportional to L_s . However, the condition number corresponds to the sensitivity of inverting the Hessian to inputs and so is inversely proportional to the computational accuracy of the problem, up to first order. Hence, it is important to check that a choice of large L_s does not lead to a condition number for the problem that is too large to give an accurate result.

Figure 3 plots $\kappa(\mathbf{S})$ against L_s and shows how $\kappa(\mathbf{S})$ begins to increase with higher L_s . Also, Figure 3 shows several inflection points which could indicate values of L_s that allow multiple observations to interact. For this paper, we pick a value of 400 km for L_s as this is large enough to propagate information sufficiently far from the different reconstructions. As seen in Figure 3, $L_s = 400$ km still has a relatively low condition number, and hence, the solution will be relatively accurate.




Figure 5. The color field is the difference between the reconstructed climate field and the prior, $\mathbf{h}(\mathbf{x}_a) - \mathbf{h}(\mathbf{x}_b)$, dimensionalized. The dots are the differences between the site-based observations, $\mathbf{y}$, and the reconstructed climate of the grid cell they are in. Observations of α have been translated to moisture index through equation (19).

4. Results

The solution using scaling values of $L_t = 1$ and $L_a = 400$ (Figure 4) produces climates at 50 sites and surrounding grid cells that are close to the reconstructions, as expected, over much of the region. However, this is not the case for the MI values of the three sites at the eastern tip of the Black sea (Apiancha, Kobuleti, and Sukhumi). These discrepant cases occur either where there is significant disagreement between different reconstructions and/or disagreement between the reconstructions and the prior with at least one of the reconstructions having relatively low variance. This reconstruction is weighted highly in the cost function, and the solution does not meet the other reconstructed variables or the prior. This creates a situation in which the best possible solution differs from both the reconstructions and prior.

The difference between the solution and the prior, transformed by equation (18) at each grid cell and dimensionalised via equation (A1), shows that the climate is much drier than the prior in the western part of the area, as shown by MI and precipitation (Figure 5). MAT has increased in some regions but decreased in others; this suggests that the inclusion of $\mathbf{C}_{L_s}$ is working as intended, since although there are varied changes in MAT, the changes occur in a spatially smooth way. Furthermore, there has been an increase in temperature seasonality as MTCO has become colder at all sites and MTWA has become warmer at most sites. This, together with the changes to MAT and GDD5, suggests that $\mathbf{C}_{L_t}$ is having the desired effect, as the changes to MTCO and MTWA are impacting the whole of the seasonal cycle of the climate and giving reasonable and smooth changes in both MAT and GDD5.

In general (Figure 6), grid cells near reconstruction sites have less error, because the solution is using information from both the prior and the reconstructions, while grid cells further away from reconstruction sites have higher error by defaulting to the error in the prior. There are some areas near reconstruction sites with high errors in MTCO, particularly in the northeast. This could reflect the fact that vegetation toward the cold and dry end of the winter temperature gradient is less sensitive to temperature change than in the




Figure 6. The standard deviation of the result, given by the dimensionalized square root of the main diagonal of $\mathbf{H}_a \mathbf{A} \mathbf{H}_a^T$ (the analysis error covariance in observation space), is represented by the color field where the dots represent sites of observations. Observations of α have been translated to moisture index through equation (19). For areas with very low temperature, it is almost certain that GDD5 is zero, and so these areas have been left blank.

Mediterranean region. However, the high median error for MTCO overall shows that there need to be large changes in MTCO from the prior to match the observations.

5. Discussion

Our final temperature reconstructions show good coherence spatially, plausible seasonal relationships and no systematic discrepancies from pollen-based reconstructions at individual sites. However, the reconstructions of moisture variables, MAP and MI, are wetter than indicated by the pollen-based reconstructions. This was expected and is realistic. The atmospheric CO_2 concentration, $[\text{CO}_2]$, was considerably lower during the LGM than it is today (180 ppm compared to 280 ppm in the PI simulations and ~ 400 ppm today). Low $[\text{CO}_2]$ decreases the water-use efficiency of plants and favors drought-adapted plants at the expense of trees, even without a change in climate (Jolly & Haxeltine, 1997; Prentice & Harrison, 2009). Although there are methods of accounting for this direct $[\text{CO}_2]$ effect (Prentice et al., 2017), statistical techniques based strictly on the application of modern analogs do not account for this impact. All of the pollen-based reconstructions for southern Europe from the Bartlein et al. (2011) data set are based on statistical reconstruction techniques. Application of the theoretically based correction factor derived by Prentice et al. (2017) to the reconstructed moisture variables would be a useful next step to improve their realism.

Sites suitable for obtaining pollen records are not uniformly distributed geographically, and in any case, the actual sampling of potential environments is extremely uneven in many regions of the world (Figure 4; Bartlein et al., 2011; Harrison et al., 2016). We have shown that the condition number can be used to identify an appropriate scale for interpolating the site-based data spatially and that a scale of 400–500 km appears to be appropriate for southern Europe at the LGM given the data currently available. This spatial scale is not uniformly appropriate, however. The standard deviation of the reconstructions (Figure 6) provides a measure of how reliable the interpolation is. More importantly, the standard deviation of the reconstruction could be used to determine when the interpolated values provide a realistic measure of the actual climate

and when they do not. Establishing an acceptable threshold value for reliability would be a useful step in the creation of the kind of paleoclimate reanalysis we are proposing here.

While the values of both scales, L_s and L_t , have been shown to be appropriate for the example shown in this paper, they are somewhat subjective. The spatial scale, L_s , is chosen to give high utilization of sparse observation data and is shown, by the condition number in Figure 3, not to lead to a numerically inaccurate solution. A value for L_t is determined by plotting the resolution matrix for multiple L_t , as shown in Figure 2; however, this only provides a range of possible values. A more objective method for selecting L_t could be developed by selecting the L_t which gives the resolution matrix closest to the identity.

6. Conclusions

In this paper we have demonstrated a novel method for reconstructing spatially explicit paleoclimate reconstructions from site-based data. The method allows the effects of each site in the data set to be tuned by imposing a structure on the error of the prior that creates reconstructions that are spatially smooth and hence more realistic. By assuming that the error in the prior with respect to temperature has a given correlation month by month, it also allows the generation of a solution that is temporally smooth. We show that a length scale L_t of 1 provides a smooth solution for the seasonal cycle, both using single sites and over multiple grid cells. Our analyses suggest that a spatial length scale (L_s) of 400 km is reasonable for southern Europe at the LGM; although this is larger than the assumed source area of most of the reconstruction sites, it reflects the large-scale coherence of the regional climate change between LGM and present. Additional work could help to determine a more objective way to determine these length scales, but nevertheless, the final climate maps appear plausible and suggest that the application of this new method should yield more robust data sets for climate-model evaluation.

Appendix A: Nondimensionalization

Most of the variables from the site-based reconstructions and PMIP3 have a dimension. This can cause a problem when computing the cost function as different variables can be at different scales and it is difficult to compare different scales together computationally. To avoid this problem, we nondimensionalize all the variables involved before computing the cost function and then redimensionalize the variables when the analytical result has been found.

We nondimensionalize the observation space using

$$D_y(\mathbf{y}_i) = \begin{pmatrix} \alpha \\ D_P(P) \\ \frac{MAT}{T_s} \\ \frac{MTWA}{T_s} \\ \frac{MTCO}{T_s} \\ \frac{GDD5}{N_y T_s} \end{pmatrix}, \quad (A1)$$

where N_y is the number of days in a year, T_s is a temperature scaling value (5 °C). The function D_P is defined as

$$D_P(P) = \begin{cases} \ln\left(\frac{P\lambda}{I_{sc}}\right) + 1 & P < \frac{I_{sc}}{\lambda} \\ \frac{P\lambda}{I_{sc}} & \text{else} \end{cases}, \quad (A2)$$

where I_{sc} is the solar constant (1,360.8 W m⁻²) and λ is the latent heat of vaporization of water (2.45 MJ kg⁻¹). D_P ensures that the method never creates a situation where $P < 0$. Similar to the observation space, we also nondimensionalize the state space using

$$D_x(\mathbf{x}_j) = \begin{pmatrix} D_P(P) \\ \frac{1}{T_s} \mathbf{T} \end{pmatrix}. \quad (A3)$$

Code and Data Availability

The reconstructions of southern Europe in this paper are available from Zenodo (<https://dx.doi.org/10.5281/zenodo.3567150>), as well as the code used to perform the data assimilation (<http://doi.org/10.5281/zenodo.3445166>). The CRU CL v2.0 data set was downloaded from the University of East Anglia Climatic Research Unit and was published as New et al. (2002). The PMIP3 LGM simulations (Braconnot et al., 2012) are available at CMIP5 archives and, for this paper, were downloaded from the Earth System Grid Federation at the Pierre Simon Laplace Institute (<https://esgf-node.ipsl.upmc.fr/projects/esgf-ipsl/>). The pollen reconstructions used from Bartlein et al. (2011) are available at the National Oceanic and Atmospheric Administration's National Centers for Environmental Information (<https://www.ncdc.noaa.gov/paleo-search/study/9897>).

References

Acknowledgments

S. F. C. was supported by a UK Natural Environment Research Programme (NERC) scholarship as part of the SCENARIO Doctoral Training Partnership at the University of Reading. S. P. H. acknowledges support from the ERC-funded project GC 2.0 (Global Change 2.0: Unlocking the past for a clearer future, Grant 694481). I. C. P. acknowledges support from the ERC under the European Union's Horizon 2020 research and innovation programme (Grant Agreement 787203 REALM). This research is a contribution to the AXA Chair Programme in Biosphere and Climate Impacts and the Imperial College initiative on Grand Challenges in Ecosystems and the Environment (ICP). N. K. N. is supported in part by the NERC National Centre for Earth Observation (NCEO). We thank PMIP colleagues who contributed to the production of the paleoclimate reconstructions. We also acknowledge the World Climate Research Programme's Working Group on Coupled Modeling, which is responsible for CMIP, and the climate modeling groups in the Palaeoclimate Modelling Intercomparison Project (PMIP) for producing and making available their model output. For CMIP, the U.S. Department of Energy's Program for Climate Model Diagnosis and Intercomparison provides coordinating support and led development of software infrastructure in partnership with the Global Organization for Earth System Science Portals. The analyses and figures are based on data archived at CMIP on 12 September 2018. We thank the Next-Generation Vegetation Modeling group for providing model code for the calculation of bioclimatic variables and for discussion of the results.

- Allen, R. G., Pereira, L. S., Raes, D., & Smith, M. (1998). Crop evapotranspiration—Guidelines for computing crop water requirements—FAO Irrigation and drainage paper 56. *Food and Agriculture Organization*, 300, D05109.
- Annan, J. D., & Hargreaves, J. C. (2013). A new global reconstruction of temperature changes at the Last Glacial Maximum. *Climate of the Past*, 9(1), 367–376.
- Bartlein, P. J., Harrison, S. P., Brewer, S., Connor, S., Davis, B. A. S., Gajewski, K., et al. (2011). Pollen-based continental climate reconstructions at 6 and 21 ka: A global synthesis. *Climate Dynamics*, 37, 775–802.
- Braconnot, P., Harrison, S. P., Kageyama, M., Bartlein, P. J., Masson-Delmotte, V., Abe-Ouchi, A., et al. (2012). Evaluation of climate models using palaeoclimatic data. *Nature Climate Change*, 2, 417–424.
- Bradley, R. S. (1999). *Paleoclimatology: Reconstructing climates of the Quaternary* (Vol. 68). San Diego, CA: Academic Press.
- Daley, R. (1994). *Atmospheric data analysis*. Cambridge: Cambridge University Press.
- Davis, T. W., Prentice, I. C., Stocker, B. D., Thomas, R. T., Whitley, R. J., Wang, H., et al. (2017). Simple process-led algorithms for simulating habitats (SPLASH v1.0): Robust indices of radiation, evapotranspiration and plant-available moisture. *Geoscientific Model Development*, 10(2), 689–708.
- Delahaies, S., Roulstone, I., & Nichols, N. (2017). Constraining DALECV2 using multiple data streams and ecological constraints: Analysis and application. *Geoscientific Model Development*, 10(7), 2635–2650.
- Garreta, V., Miller, P. A., Guiot, J., Hély, C., Brewer, S., Sykes, M. T., & Litt, T. (2010). A method for climate and vegetation reconstruction through the inversion of a dynamic vegetation model. *Climate Dynamics*, 35(2), 371–389. <https://doi.org/10.1007/s00382-009-0629-1>
- Gebhardt, C., Kühl, N., Hense, A., & Litt, T. (2008). Reconstruction of quaternary temperature fields by dynamically consistent smoothing. *Climate Dynamics*, 30(4), 421–437. <https://doi.org/10.1007/s00382-007-0299-9>
- Golub, G. H., & Loan, C. F. V. (1996). *Matrix computations* (3rd ed.). Baltimore & London: Johns Hopkins University Press.
- Goosse, H., Renssen, H., Timmermann, A., Bradley, R. S., & Mann, M. E. (2006). Using palaeoclimate proxy-data to select optimal realisations in an ensemble of simulations of the climate of the past millennium. *Climate Dynamics*, 27(2), 165–184.
- Gornitz, V. (2008). *Encyclopedia of paleoclimatology and ancient environments*. New York: Springer Science & Business Media.
- Haben, S. A., Lawless, A. S., & Nichols, N. K. (2011). Conditioning of incremental variational data assimilation, with application to the Met Office system. *Tellus A: Dynamic Meteorology and Oceanography*, 63(4), 782–792.
- Handcock, M. S., & Wallis, J. R. (1994). An approach to statistical spatial-temporal modeling of meteorological fields. *Journal of the American Statistical Association*, 89(426), 368–378. <https://doi.org/10.1080/01621459.1994.10476754>
- Harrison, S. P., & Bartlein, P. J. (2012). Records from the past, lessons for the future: What the palaeo-record implies about mechanisms of global change. In A. Henderson-Sellers & K. McGuffie (Eds.), *The future of the world's climate* (pp. 403–436). Amsterdam: Elsevier.
- Harrison, S. P., Bartlein, P. J., Brewer, S., Prentice, I. C., Boyd, M., Hessler, I., et al. (2014). Climate model benchmarking with glacial and mid-Holocene climates. *Climate Dynamics*, 43, 671–688.
- Harrison, S. P., Bartlein, P. J., Izumi, K., Li, G., Annan, J., Hargreaves, J., et al. (2015). Evaluation of CMIP5 palaeo-simulations to improve climate projections. *Nature Climate Change*, 5, 735–743.
- Harrison, S. P., Bartlein, P. J., & Prentice, I. C. (2016). What have we learnt from palaeoclimate simulations? *Journal of Quaternary Science*, 31(4), 363–385.
- Harrison, S. P., Prentice, I. Colin, Barboni, D., Kohfeld, K. E., Ni, J., & Sutra, J.-P. (2010). Ecophysiological and bioclimatic foundations for a global plant functional classification. *Journal of Vegetation Science*, 21, 300–317.
- Jolly, D., & Haxeltine, A. (1997). Effect of low glacial atmospheric CO₂ on tropical African montane vegetation. *Science*, 276(5313), 786–788.
- Kageyama, M., Braconnot, P., Harrison, S. P., Haywood, A. M., Jungclauss, J. H., Otto-Bliesner, B. L., et al. (2018). The PMIP4 contribution to CMIP6—Part 1: Overview and over-arching analysis plan. *Geoscientific Model Development*, 11(3), 1033–1057.
- Lahoz, W. A., & Schneider, P. (2014). Data assimilation: Making sense of earth observation. *Frontiers in Environmental Science*, 2, 16. <https://doi.org/10.3389/fenvs.2014.00016>
- Liu, D. C., & Nocedal, J. (1989). On the limited memory BFGS method for large scale optimization. *Mathematical Programming*, 45(1), 503–528.
- MARGO Project Members (2009). Constraints on the magnitude and patterns of ocean cooling at the Last Glacial Maximum. *Nature Geoscience*, 2, 127–132.
- Maddock, J., Bristow, P., Holin, H., & Zhang, X. (2018). Boost c++/math/special functions. <https://www.boost.org/>, (ver. 1.69).
- Marsicek, J., Shuman, B. N., Bartlein, P. J., Shafer, S. L., & Brewer, S. (2018). Reconciling divergent trends and millennial variations in Holocene temperatures. *Nature*, 554(7690), 92.
- Matérn, B. (1986). *Spatial variation* (2nd ed., Vol. 36). New York, NY: Springer New York.
- Menke, W. (2012). *Geophysical data analysis: Discrete inverse theory* (3rd ed.). Cambridge, Massachusetts: Academic Press.
- New, M., Lister, D., Hulme, M., & Makin, I. (2002). A high-resolution data set for surface climate over global land areas. *Climate Research*, 21, 1–25.
- Nichols, N. K. (2010). Mathematical concepts of data assimilation. In W. Lahoz, B. Khattatov, & R. Menard (Eds.), *Data assimilation* (pp. 718). Verlag Berlin Heidelberg: Springer.
- Prentice, I. C., Cleator, S. F., Huang, Y. H., Harrison, S. P., & Roulstone, I. (2017). Reconstructing ice-age palaeoclimates: Quantifying low-CO₂ effects on plants. *Global and Planetary Change*, 149, 166–176.

- Prentice, I. C., & Harrison, S. P. (2009). Ecosystem effects of CO₂ concentration: Evidence from past climates. *Climate of the Past*, 5(3), 297–307.
- Schmidt, G. A., Annan, J. D., Bartlein, P. J., Cook, B. I., Guilyardi, E., Hargreaves, J. C., et al. (2014). Using palaeo-climate comparisons to constrain future projections in CMIP5. *Climate of the Past*, 10, 221–250.
- Simonis, D., Hense, A., & Litt, T. (2012). Reconstruction of late Glacial and Early Holocene near surface temperature anomalies in Europe and their statistical interpretation. *Quaternary International*, 274, 233–250. <https://doi.org/10.1016/j.quaint.2012.02.050>
- Steiger, N. J., Steig, E. J., Dee, S. G., Roe, G. H., & Hakim, G. J. (2017). Climate reconstruction using data assimilation of water isotope ratios from ice cores. *Journal of Geophysical Research: Atmospheres*, 122, 1545–1568. <https://doi.org/10.1002/2016JD026011>
- Tardif, R., Hakim, G. J., Perkins, W. A., Horlick, K. A., Erb, M. P., Emile-Geay, J., et al. (2018). Last millennium reanalysis with an expanded proxy database and seasonal proxy modeling. *Climate of the Past Discussions*, 2018, 1–37. <https://doi.org/10.5194/cp-2018-120>
- Ter Braak, C. J. F., & Juggins, S. (1993). Weighted averaging partial least squares regression (WA-PLS): An improved method for reconstructing environmental variables from species assemblages. In H. van Dam (Ed.), *Twelfth international diatom symposium* (pp. 485–502). Dordrecht: Springer Netherlands.
- Woodward, F. I. (1987). *Climate and plant distribution*. Cambridge: Cambridge University Press.
- Zhang, L., Hickel, K., Dawes, W. R., Chiew, F. H. S., Western, A. W., & Briggs, P. R. (2004). A rational function approach for estimating mean annual evapotranspiration. *Water Resources Research*, 40, W02502. <https://doi.org/10.1029/2003WR002710>