Journal of the Royal Statistical Society Series C: Applied Statistics, 2023, **72**, 1493–1511
https://doi.org/10.1093/jrsssc/qlad078
Advance access publication 18 September 2023
Original Article

ROYAL STATISTICAL SOCIETY DATA | EVIDENCE | DECISIONS
Journal of the Statistics Society
Series C
Applied Statistics

Reconstructing the Antarctic ice-sheet shape at the Last Glacial Maximum using ice-core data

Fiona E. Turner[1], Caitlin E. Buck[2], Julie M. Jones[3], Louise C. Sime[4], Irene Malmierca Vallet[4] and Richard D. Wilkinson[5]

[1]Department of Geography, King's College London, Bush House NE Wing, Aldwych, London, UK
[2]School of Mathematics and Statistics, University of Sheffield, UK
[3]Department of Geography, The University of Sheffield, Sheffield, UK
[4]British Antarctic Survey, Cambridge, UK
[5]School of Mathematical Sciences, University of Nottingham, University Park, Nottingham, UK

Address for correspondence: Fiona E. Turner, Department of Geography, King's College London, Bush House NE Wing, Aldwych, London WC2B 4BG, UK. Email: fiona.turner@kcl.ac.uk

Abstract

The Antarctic ice sheet (AIS) is the Earth's largest store of frozen water; understanding how it changed in the past allows us to improve projections of how it, and sea levels, may change. Here, we use previous AIS reconstructions, water isotope ratios from ice cores, and simulator predictions of the relationship between the ice-sheet shape and isotope ratios to create a model of the AIS at the Last Glacial Maximum. We develop a prior distribution that captures expert opinion about the AIS, generate a designed ensemble of potential shapes, run these through the climate model HadCM3, and train a Gaussian process emulator of the link between ice-sheet shape and isotope ratios. To make the analysis computationally tractable, we develop a preferential principal component method that allows us to reduce the dimension of the problem in a way that accounts for the differing importance we place in reconstructions, allowing us to create a basis that reflects prior uncertainty. We use Markov chain Monte Carlo to sample from the posterior distribution, finding shapes for which HadCM3 predicts isotope ratios closely matching observations from ice cores. The posterior distribution allows us to quantify the uncertainty in the reconstructed shape, a feature missing in other analyses.

Keywords: Antarctic ice sheet, Bayesian calibration, Gaussian process emulation, palaeo-climatology

1 Introduction

Understanding the history and evolution of the Earth's ice sheets is important for predicting the effects of future climate change. The Antarctic ice sheet (AIS) is the largest area of land ice globally, and looks set to contribute up to 42 cm of sea-level rise by 2,100 under current emission policies (Edwards et al., 2019). Predicting what changes we might see in the AIS over the next century requires us to understand what changes have occurred in the past. In this paper, we focus on inferring the AIS shape at the Last Glacial Maximum (LGM), the most recent time at which the ice sheets were at their greatest extent, which in the Antarctic continent occurred approximately 21 Ka BP (thousand years before present) (Clark et al., 2009). The LGM is of particular interest because of the uncertainty around the size of the AIS and its contribution to sea level change following the termination of the most recent glacial period (Blasco et al., 2021; Khan et al., 2019; Simms et al., 2019; Yokoyama & Purcell, 2021).

Whilst there are marine and terrestrial glacial geological datasets that can help determine the position of the LGM ice-sheet grounding-line (the region where ice transitions from grounded ice sheet to

Received: September 12, 2022. **Accepted:** August 4, 2023

freely floating ice shelf), the ice-sheet upper surface, and in some cases flow-directional features (e.g. Bentley et al., 2014), measurements from ice cores have rarely been used to attempt to ascertain the size and shape of the LGM ice sheet. It is therefore novel and of some interest to use proxy data on the isotopic ratio of oxygen-16 to oxygen-18, referred to as $\delta^{18}O$, in the precipitation that fell on Antarctica at the LGM, as part of a Bayesian analysis. This information can be collected from ice cores drilled to deep within the AIS. The ice in the core can be dated, enabling glaciologists to accurately reconstruct the history of the variation in the $\delta^{18}O$ record over time (Jouzel et al., 2013). This data is complex to interpret and model, and does not by itself tell us the ice-sheet shape. However, climate modellers have built global circulation models (GCMs) that can simulate how the $\delta^{18}O$ values in precipitation depends upon the height of the AIS as well as global climatic conditions (Werner et al., 2018). Global conditions are important as the distribution and transport of water molecules containing the heavier $\delta^{18}O$ molecules varies both spatially and temporally in the atmosphere depending on climatological conditions. Most notably, there is a clear linear relationship between $\delta^{18}O$ values and mean temperature at high latitudes (Jouzel et al., 2013). The relationship between temperature and site elevation has then been used to attempt to determine the height of the ice sheet, allowing estimation of the size and volume of the ice sheet (e.g. Sutter et al., 2020).

Previous approaches to estimating the AIS extent at the LGM (e.g. Briggs et al., 2014; Peltier, 2004; Whitehouse et al., 2012) have tended to rely primarily on either glaciological and/or geological data, or ice-sheet model simulations. There is no study that has attempted to combine both sources of information in a principled statistical manner for this time period, in large part because the climate simulators are necessarily complex and computationally expensive to run, as they need to resolve not just the climate, but also isotope transfer, sea ice area, air pressure at sea level and other variables. This computational expense makes estimation of the simulator parameters (such as the AIS shape) impossible using standard statistical methods, as we are limited to only a few simulator evaluations. The methods of both Werner et al. (2018) and Domingo et al. (2020) are similar to our own, using previous reconstructions, GCMs and proxy data to reduce uncertainty around the AIS and Greenland ice sheets at two different time periods, the LGM and Last Interglacial (LIG), respectively. These papers highlight the importance of better understanding past changes in global ice sheets, and demonstrate the role that statisticians can play in this research area; we believe we can build on these methods by incorporating expert judgement, and taking a more detailed approach to building a prior model. Although both this paper and Domingo et al. (2020) are based on the methodologies described in Wilkinson (2010) and Higdon et al. (2008), and are reconstructing ice sheets during a palaeo time period, our method includes a careful incorporation of data using a new and novel principal component analysis (PCA) technique, and a carefully constructed prior distribution.

In this paper, we take a Bayesian approach to the problem, and develop a methodology to combine three primary sources of information: data collected from ice-core samples; complex computer simulations of the climate that encode our scientific understanding of the interaction between climate and the ice-sheet shape; and the expert opinion of glaciologists at the British Antarctic Survey (BAS), which is the UK's national Antarctic research centre. We use the previous reconstructions of the AIS at the LGM as the basis of a prior model, with parameter priors determined in an elicitation exercise with glaciologists, before using the proxy ice-core data to update this model. The climate simulator provides the link between our model of the ice-sheet shape, and the ice-core data. The statistical challenge is solving the associated statistical inverse problem in order to estimate our posterior distribution for the ice-sheet shape.

There are two main difficulties in doing this. Firstly, each simulation takes 24 hr to run approximately 20 years on the supercomputer facility: ARCHER (2016). This cost (and our computational budget) means that we are limited in how many simulator evaluations we can perform. We develop a Gaussian process (GP) emulator of the simulator (Chang et al., 2014; Kennedy & O'Hagan, 2001; Wilkinson, 2010), a methodology that has become a widely used tool in the climate modelling community (Edwards et al., 2021; Holden et al. 2015, 2018; Lowry et al., 2021), which can be used as a cheap, fast approximation to the simulator.

The second challenge is the number of degrees of freedom in the complex computer simulations vs. the amount of data available. The latter is determined by the limited number of ice cores with $\delta^{18}O$ which extend back to the LGM. A flexible uninformative prior structure would result in an uninformative analysis, as we cannot hope to constrain so many parameters with such limited data. Instead, we seek a carefully elicited prior model that constrains the space of shapes to be

considered. This process is challenging as conceptualising uncertainty in high-dimensional spaces is hard. We thus develop the use of the principal components of the 40 previously published AIS reconstructions as a starting point for a reduced-dimensional basis. Some of the reconstructions are of greater importance than others due to being data-based reconstructions rather than ice-sheet model output, and some of them are highly correlated as they are produced by the same team, e.g. de Boer et al. (2017). To deal with this, we develop a *preferential PCA* approach that allows us to ensure the important reconstructions are given the greatest importance, and we show how this can be computed efficiently in high-dimensional problems (Section 2.1). The development of the prior distribution is described in Section 2.1, and the emulator is described in Section 3. In Section 3.2, we describe how to perform inference for this model to solve the inverse problem given in equation (1). In Section 4, we present results using synthetic data to validate the inferential approach, and then analyse the real data. The results from this work represent an improvement on the current state-of-the art reconstructions of past Antarctic elevation through the use of ice-core data by Buizert et al. (2021) and Werner et al. (2018), who use data from 7 and 11 ice cores, respectively. We have particular confidence in our results on mean elevation changes that are ascertained for Antarctic regions far from the coast such as for the central West Antarctic ice sheet (WAIS) and East Antarctic ice sheet (EAIS). Section 5 contains our discussion.

2 Methods

We let $\mathbf{x} \in \mathbb{R}^D$ denote the AIS shape at the LGM. We aim to estimate $\mathbf{x}$ using observations of the isotope ratio of $\delta^{18}O$, in samples denoted as $\mathbf{y} \in \mathbb{R}^n$, collected from n different ice-core sites at depths corresponding to the LGM (WAIS Divide Project Members et al., 2013). Only a few ice-core sites containing $\delta^{18}O$ data that extend to the LGM have been drilled in Antarctica, and so we have at most $n = 10$ data points available, as shown in online supplementary Figure 1.

We use the Hadley Centre Coupled Model version 3 (HadCM3) (Met-Office, 2016), with water isotope code (Tindall et al., 2009), as the climate simulator which implements various physical laws on a discrete computational grid with longitude/latitude grid cells of $3.75^\circ \times 2.5^\circ$ covering the entire globe, to which has been added a simulation of isotope transport. This allows us to include water isotope values as part of our simulations. The resolution of the computational grid used in the simulator is 96×73 due to the size of the longitude/latitude grid cells, meaning that the height of the AIS can be represented by a vector of length 7,008, i.e. $\mathbf{x} \in \mathbb{R}^{7,008}$, with each element recording the height of the AIS in a grid cell of the computational mesh used.

We let $f(\mathbf{x})$ denote the climate simulator's prediction of the observed oxygen isotope ratios at the n sites when the AIS shape at the LGM is $\mathbf{x}$. We use a simple additive error model to link $\mathbf{x}$ to $\mathbf{y}$ via f:

$$\mathbf{y} = f(\mathbf{x}) + \boldsymbol{\epsilon}. \tag{1}$$

Here, $\boldsymbol{\epsilon}$ is an error term representing all of the various sources of error (measurement error, model discrepancy, errors arising from discretisation, etc.).

Our aim is to use previous reconstructions of the AIS at the LGM, as well as expert judgement, to develop a prior distribution for the AIS shape, $\pi(\mathbf{x})$, and then to compute the posterior distribution of $\mathbf{x}$ given the data and the simulator: $\pi(\mathbf{x} \mid \mathbf{y})$.

2.1 Prior model

We use previous estimates of the shape of the AIS at the LGM, in combination with expert judgement, to build a prior model for $\mathbf{x}$. Our prior elicitation process began with an extensive literature review in which we sought and collected together previous estimates of the AIS at the LGM; this resulted in a collection of 40 different ice-sheet reconstructions created by a variety of authors (Argus et al., 2014; Bentley et al., 2014; Briggs et al., 2014; de Boer et al., 2017; Golledge et al. 2012, 2013; Peltier, 2004; Pollard & DeConto, 2009; Whitehouse et al., 2012; detail of these reconstructions can be found in online supplementary Tables 1 and 2). We vectorise these ice-sheet reconstructions and denote them as $\mathbf{x}_1, \ldots, \mathbf{x}_{40} \in \mathbb{R}^D$ with $D = 7{,}008$.

Given the limited data and computational budget (typically $n = 10$ observations and $m = 50$ GCM simulations), we cannot hope to directly estimate $\mathbf{x} \in \mathbb{R}^{7,008}$ if our model of $\mathbf{x}$ is too flexible. We thus seek a prior model that is still sufficiently flexible to allow us to make useful inferences

that shed additional light on the AIS at the LGM, but whose complexity is commensurate with the information available to fit the model. Thus, we seek a model for $\mathbf{x}$ that only has a small number, d say, of parameters to be estimated. We wish to use this model to represent physically plausible variation in $\mathbf{x}$, and it must be physically interpretable to enable expert elicitation. This led us to consider a linear basis expansion for $\mathbf{x}$ of the form

$$\mathbf{x} = \sum_{i=1}^{d} \boldsymbol{\theta}_i \lambda_i \mathbf{a}_i + \boldsymbol{\mu}, \tag{2}$$

where $\boldsymbol{\mu}$ is the prior mean of the AIS shape, which we estimate using the mean of the library of previous reconstructions. The basis vectors $\mathbf{a}_i \in \mathbb{R}^D$ define the directions in which the ice sheet shape may vary, and the scalars $\lambda_i \in \mathbb{R}$ specify the relative importance of each mode of variation and allow us to put the variation in $\boldsymbol{\theta}_i$ on a similar scale for each i. The pairs $\{\lambda_i, \mathbf{a}_i\}_{i=1}^{d}$ need to be known before any simulation with HadCM3 is carried out, and so we estimate these using the previous reconstructions $\{\mathbf{x}_i\}_{i=1}^{40}$. The scalars $\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_d$ are the parameters that we will estimate using the ice-core data and the simulator via equation (1). We describe the process for eliciting the prior distribution of $\boldsymbol{\theta}$ using expert judgement in light of the choice of λ_i, $\mathbf{a}_i$ in Section 3 below.

Ideally, we would seek an expansion that had parameters that are of active importance in the simulator response and are identifiable from the data. Whilst there are methods for determining such $\mathbf{a}_i$, for example active sub-space (Constantine, 2015), given the constraints on the computational budget, and lack of gradient information about the simulator, this makes these methods impracticable. Instead, we focus on finding an expansion of the form (2) that is interpretable and capable of describing the uncertainty about $\mathbf{x}$, as understood by glaciologists at BAS.

Using the library of 40 previous AIS shape reconstructions, we restrict the space of plausible shapes we consider to $S = \text{span}\{\mathbf{x}_1, \ldots, \mathbf{x}_{40}\}$, which is a 40-dimensional sub-space defined by the 40 previous reconstructions. However, this space still has too many degrees of freedom to estimate given the limited number of data points and computational budget. Thus, we seek a lower dimensional approximation to shapes in S, i.e. we seek $\tilde{S} = \text{span}\{\mathbf{a}_1, \ldots, \mathbf{a}_d\}$, where $d < 40$. If our aim was solely to ensure $\tilde{S}$ was close to S, then we could use the leading principal components of $\{\mathbf{x}_i\}_{i=1}^{40}$. However, the previous reconstructions are not all of equal importance to the glaciology community.

Four of the AIS reconstructions are data-based and are widely favoured by the palaeoclimatology community (e.g. when running climate models and in citations). These are ICE5G (Peltier, 2004), ICE6G (Argus et al., 2014), W12A (Whitehouse et al., 2012), and GLAC-1D (Briggs et al., 2014), which we denote as $\mathbf{x}_1$, $\mathbf{x}_2$, $\mathbf{x}_3$, $\mathbf{x}_4$, respectively. The remaining 36 ice-sheet reconstructions, $\mathbf{x}_5, \ldots, \mathbf{x}_{40}$, are from Pollard and DeConto (2009), Golledge et al. (2012, 2013), Bentley et al. (2014), and de Boer et al. (2017) and were constructed using ice-sheet simulation models. Principal component analysis of the collection of 40 reconstructions would give all previous reconstructions equal importance, and would find a low-dimensional representation that captures as much of the variance as possible. This leads to the information contained in the first four ice-sheet reconstructions, $\mathbf{x}_1, \ldots, \mathbf{x}_4$ (ICE5G, ICE6G, W12A, and GLAC-1D), being overwhelmed and $\mathbf{x}_9, \ldots, \mathbf{x}_{40}$, the 32 shapes from de Boer et al. (2017), dominating the modes of variation $\mathbf{a}_i$ in the prior model. These shapes came from the same ice-sheet model, with small adjustments made to the simulations. They are therefore all similar to each other, and if treated equally with the other eight shapes, may heavily influence our prior model. Instead, we develop an alternative form of PCA that allows us to stress the importance of the privileged shapes $\mathbf{x}_1, \ldots, \mathbf{x}_4$ and then to complement this with information from $\mathbf{x}_5, \ldots, \mathbf{x}_{40}$. We describe this approach generally in the next section, before specialising to this application case in Section 2.3. We have not found this problem or its solution discussed elsewhere, and believe it may prove to be a useful approach in other problems.

2.2 Preferential PCA

Suppose we are given two datasets consisting of n_1 and n_2 observations in $\mathbb{R}^D$, which we denote as $\boldsymbol{X}_1 \in \mathbb{R}^{D \times n_1}$ and $\boldsymbol{X}_2 \in \mathbb{R}^{D \times n_2}$, where the data matrices have been jointly column centred. Principal

component analysis is performed on $\boldsymbol{X}_1$ to find the leading principal components $\mathbf{a}_1, \ldots, \mathbf{a}_k$, which maximise $\mathbb{V}\mathrm{ar}(\boldsymbol{X}_1^\top \mathbf{a})$ subject to $\mathbf{a}_i^\top \mathbf{a}_j = \delta_{ij}$. We then want to find orthonormal vectors $\mathbf{a}_{k+1}, \mathbf{a}_{k+2}, \ldots$ that are orthogonal to $\mathbf{a}_1, \ldots, \mathbf{a}_k$, and which describe a maximal amount of variance in dataset 2. In other words, we wish to solve the following optimisation problem for $i = k+1, k+2, \ldots$:

$$\begin{aligned} \max \quad & \mathbb{V}\mathrm{ar}(\boldsymbol{X}_2^\top \mathbf{a}) \\ \text{subject to} \quad & \mathbf{a}^\top \mathbf{a} = 1 \\ & \mathbf{a}_j^\top \mathbf{a} = 0 \ \text{ for } j = 1, \ldots, i-1. \end{aligned} \tag{3}$$

This is a quadratically constrained quadratic program, an optimisation problem where both the objective function and the constraints are quadratic functions. This is a non-convex problem, so many of the standard numerical optimisation approaches fail. We present a solution to equation (3) that first reparameterises the problem to the null space of $\boldsymbol{A}^\top$, where $\boldsymbol{A} = (\mathbf{a}_1 \ \mathbf{a}_2 \ \ldots \ \mathbf{a}_k)$, the matrix with columns formed from the principal components of $\boldsymbol{X}_1$.

If we form the Lagrangian of problem (3), then the Karush–Kuhn–Tucker optimality conditions (see, e.g. Boyd et al., 2004) give that if $\mathbf{a}$ is a solution to equation (3), then we must have

$$\begin{aligned} 2\boldsymbol{X}_2\boldsymbol{X}_2^\top \mathbf{a} + 2\lambda \mathbf{a} + \sum_{j=1}^{i-1} \mu_j \mathbf{a}_j = 0 \\ \mathbf{a}^\top \mathbf{a} = 1 \\ \mathbf{a}_j^\top \mathbf{a} = 0 \ \text{ for } j = 1, \ldots, i-1, \end{aligned}$$

where $\lambda, \mu_j \in \mathbb{R}$. Because $\mathbf{a}_1, \ldots, \mathbf{a}_k$ were derived independently of $\boldsymbol{X}_2$, and specifically because they are not eigenvectors of the empirical covariance matrix $\frac{1}{n_2}\boldsymbol{X}_2\boldsymbol{X}_2^\top$, the usual optimisation approach used in PCA fails.

Instead, we can rewrite the problem. Note that the solution to equation (3) must lie in $\mathrm{null}(\boldsymbol{A}^\top) = \{\mathbf{x} : \boldsymbol{A}^\top \mathbf{x} = 0\}$, the null space of $\boldsymbol{A}^\top$. If $\{\mathbf{b}_1, \ldots, \mathbf{b}_{D-k}\}$ is an orthonormal basis for $\mathrm{null}(\boldsymbol{A}^\top)$, we can write

$$\mathbf{a} = \boldsymbol{B}\mathbf{w} \ \text{ for some } \mathbf{w} \in \mathbb{R}^{D-k}.$$

Problem (3) then becomes, for $i = k+1, k+2, \ldots$

$$\begin{aligned} \max \quad & \mathbb{V}\mathrm{ar}(\boldsymbol{X}_2^\top \boldsymbol{B}\mathbf{w}) \\ \text{subject to} \quad & \mathbf{w}^\top \mathbf{w} = 1 \\ & \mathbf{w}^\top \mathbf{w}_j = 0 \ \text{ for } j = k+1, \ldots, i-1. \end{aligned} \tag{4}$$

This is the standard principal component problem for $\boldsymbol{B}^\top \boldsymbol{X}_2$, and thus the solution is to take $\mathbf{w}_i$ to be the $(i-k)$th left singular vector of $\boldsymbol{B}^\top \boldsymbol{X}_2$ (or equivalently, eigenvector of $\boldsymbol{B}^\top \boldsymbol{X}_2 \boldsymbol{X}_2^\top \boldsymbol{B}$). We then set $\mathbf{a}_i = \boldsymbol{B}\mathbf{w}_i$. The problem of finding the null space of $\boldsymbol{A}^\top$ and solving the eigenvalue problem can be solved efficiently by the QR or singular value decomposition, depending on speed and accuracy requirements (Van Loan & Golub, 1996).

To summarise, we take $\mathbf{a}_1, \ldots, \mathbf{a}_k$ as the first k principal components of the reference dataset $\boldsymbol{X}_1$, and then take $\mathbf{a}_{k+1} = \boldsymbol{B}\mathbf{w}_1$, $\mathbf{a}_{k+2} = \boldsymbol{B}\mathbf{w}_2, \ldots$ where $\mathbf{w}_1, \mathbf{w}_2, \ldots$ are the leading principal components of $\boldsymbol{B}^\top \boldsymbol{X}_2$.

2.3 Prior distribution for the AIS at the LGM

The number of basis vectors in our prior model [equation (2)], and hence the effective dimension of the parameter we wish to learn, has to be decided in advance of running any climate simulations.

The decision needs to take into account the computational budget (which limits us to approximately 50 simulations of HadCM3), the amount of data available, and how we intend to estimate the posterior distribution. Incorporating additional basis vectors into the model increases its flexibility and expressive power, but results in more parameters to estimate, i.e. the usual bias-variance trade-off. To compute the posterior, we will use a Gaussian process emulator of the simulator that approximates the simulator mapping from ice-sheet shapes to observations. The general rule-of-thumb advocated in Loeppky et al. (2009) and elsewhere suggests that to train a GP on d-dimensional inputs requires at least $10d$ training points, which suggests we should allow at most 5 degrees of freedom in our prior model.

PCA of the important ice-sheet reconstructions $\boldsymbol{X}_1 = (\mathbf{x}_1 \ \ldots \ \mathbf{x}_4)$ (the reconstructions we wish to give preference to) shows that we can represent 94.2% of the variation in these shapes with the first three principal components, $\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3$, which we collect in matrix $\boldsymbol{A} = [\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3]$. We then use the second set of AIS reconstructions $\boldsymbol{X}_2 = (\mathbf{x}_5 \ \ldots \ \mathbf{x}_{40})$ to add additional information about the likely ice-sheet variation. Applying the preferential PCA approach to $\boldsymbol{X}_2$ gives candidate values $\mathbf{a}_4, \mathbf{a}_5, \ldots$. Table 1 gives the root mean square reconstruction errors when using either four, five, or six basis vectors to represent the 40 previous AIS reconstructions. As expected, using standard PCA results in a slightly lower root mean square error (RMSE) than the preferential PCA method but higher accuracy on the preferential shapes: we sacrifice accuracy in the overall reconstruction in order to prioritise the preservation of information contained in the preferential shapes $\mathbf{x}_1, \ldots, \mathbf{x}_4$. Given the budgetary constraints and the errors in Table 1, we decide to use five basis vectors.

Now we have selected which basis vectors to use in the prior model [equation (2)], we can complete the specification of the prior. The free parameters are $\mu \in \mathbb{R}^{7,008}$, the prior mean ice-sheet shape, which we take to be the mean of the 40 prior ice-sheet reconstructions $\mu = \frac{1}{40}\sum_{i=1}^{40} \mathbf{x}_i$, and $\boldsymbol{\theta} \in \mathbb{R}^5$. We can think of $\boldsymbol{\theta}$ as the coordinates with respect to the basis $\{\lambda_i \mathbf{a}_i\}_{i=1}^5$, controlling the contribution of each basis vector to the synthetic ice-sheet shapes. To project an ice-sheet shape $\mathbf{x}$ onto the basis vectors we set $\boldsymbol{\theta} = (\mathbf{x} - \boldsymbol{\mu})^\top \boldsymbol{A}\Lambda^{-1}$ where $\Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_5)$.

Finally, we need to specify a prior distribution for $\boldsymbol{\theta}$, which we did by conducting an informal two-stage expert elicitation process (O'Hagan et al., 2006) with five glaciologists from the British

Table 1. The average RMSE for the 40 collected ice-sheet shapes when using four, five, and six basis vectors

Number of basis vectors	Ice-sheet shape	PCA	Preferential PCA
4	ICE5-G	13.07	1.42
	ICE6-G	12.85	6.07
	W12A	55.98	35.63
	GLAC-1D	43.23	44.21
	All	13.52	14.05
5	ICE5-G	12.87	1.38
	ICE6-G	9.44	5.89
	W12A	2.07	34.53
	GLAC-1D	38.75	42.85
	All	11.34	11.93
6	ICE5-G	10.85	1.38
	ICE6-G	4.02	5.89
	W12A	2.04	34.53
	GLAC-1D	10.73	42.85
	All	9.89	10.22

Note. Although our method does not reduce the average RMSE, it does reduce the RMSE for the orographies in X_1 in most cases. PCA = principal component analysis; RMSE = root mean square error.

Antarctic Survey. To aid the elicitation exercise, we created an interactive version of our prior model, which displayed the ice-sheet orography as an interactive map (a snapshot of this interactive model can be seen in online supplementary Figure 2). The method of empirical orthogonal functions (EOFs) (Hannachi et al., 2007) is widely used in geophysics, and so the form of our prior model was understood well by the glaciologists we consulted. We assumed a prior distribution for $\boldsymbol{\theta}$ of the form $\boldsymbol{\theta} \sim \mathcal{N}_5(\mathbf{0}, \Sigma_\theta)$, with $\Sigma_\theta = \text{diag}(\sigma_{\theta,1}, \ldots, \sigma_{\theta,5})$ remaining to be chosen. In the first stage of the elicitation exercise, we adjusted the values of $\sigma_{\theta,i}$ (starting from $\sigma_{\theta,i} = 3$) until the glaciologists were satisfied the random samples from the prior model were plausible AIS shapes. We eventually settled on $\sigma^2_{\theta,1} = \sigma^2_{\theta,2} = \sigma^2_{\theta,4} = 0.5$, $\sigma^2_{\theta,3} = 0.6$, and $\sigma^2_{\theta,5} = 1$.

To train the GP emulator, we need a design of $\boldsymbol{\theta}$ samples to generate ice-sheet shapes to run through HadCM3. We used a space-filling maxi-min Latin Hypercube design (Joseph & Hung, 2008) generated in $[0, 1]^5$, which we then pushed through the inverse cumulative distribution function of the respective normal distributions to generate a space-filling design specific to our prior distribution. Given the tight constraints on computational budget, there was a strong desire not to simulate using any ice-sheet shapes that were felt to be implausible representations of the AIS at the LGM. So we conducted a second elicitation exercises, where we shared the shapes in our initial design with the glaciologists at BAS, who independently ruled out or accepted each ice-sheet shape as plausible or implausible on a sample-by-sample basis. This resulted in a design of 47 synthetic orographies to include in simulations to run through HadCM3, which were hopefully collectively representative of the space of plausible AIS shapes.

2.4 HadCM3 simulations

We created 47 orography files describing plausible AIS shapes at the LGM. We built 47 simulation files with control LGM atmospheric greenhouse gas concentrations and orbital forcing. These were then input into HadCM3 and run for 60 model years each on ARCHER, including 10 years of spin-up time. For more detail on HadCM3, see Valdes et al. (2017).

3 Emulation and inference

The prior model allows us to create synthetic ice-sheet shapes, and to simulate [via equation 1)] what isotope data we might have expected to observe given a particular AIS geometry. Our aim is to use the observed data $\mathbf{y}_{obs}$ to infer the posterior distribution of the basis parameters

$$\pi(\boldsymbol{\theta} \mid \mathbf{y}_{obs}), \tag{5}$$

which then induces a posterior distribution for the AIS shape $\pi(\mathbf{x} \mid \mathbf{y}_{obs})$.

We have data from 10 different sites in Antarctica: four located in the WAIS at Byrd (Blunier & Brook, 2001), Mount Moulton (Popp, 2008), Siple Dome (Brook et al., 2005; WAIS Divide Project Members et al., 2013), and WDC (Steig et al., 2013; WAIS Divide Project Members et al., 2013), and six in the EAIS at EDC (Jouzel et al., 2013), EDML (EPICA Community Members et al., 2006), Fuji Dome (Kawamura et al., 2007), Talos Dome (Stenni et al., 2011), Taylor Dome (Grootes et al., 2001; Steig et al., 2000), and Vostok (Petit et al., 1999). A map showing the location of these sites is given in online supplementary Figure 6. At each site ($i = 1, \ldots, 10$), we have a single record of the $\delta^{18}O$ value, y^i_{obs}, which we collate in the vector $\mathbf{y}_{obs} \in \mathbb{R}^{10}$.

One approach to estimate the posterior distribution [equation (5)] is to use Markov chain Monte Carlo (MCMC). However, this would require tens of thousands of simulations of the forward model. The computational cost of running HadCM3 limits the number of simulator evaluations we can perform. Instead, we build an approximation to $f(\mathbf{x})$, which we refer to as the emulator (or surrogate model) of f (O'Hagan, 2006). We used the agreed design of 47 ice-sheet shapes described in the previous section and ran HadCM3 using each of them as inputs. This provided us with a set of $\delta^{18}O$ values paired with the synthetic orographies. Let $\mathcal{D}$ denote this set of simulator evaluations $\mathcal{D} = \{(\mathbf{x}_i, f(\mathbf{x}_i)\}_{i=1}^{47}$. We use $\mathcal{D}$ to build an emulator of HadCM3 which approximates the HadCM3 relationship between ice-sheet shape and $\delta^{18}O$ observations with minimal computational cost. In Section 3.1, we describe this process, and in Section 3.2 the subsequent process for computing the posterior distribution using the emulator.

3.1 Emulation of HadCM3

Given an ice-sheet orography, $\mathbf{x}$, HadCM3 outputs a number of spatial fields predicting various aspects of the climate. Our interest lies solely in the predictions of the $\delta^{18}O$ value at each of the 10 ice-core sites, which we will denote as $f_1(\boldsymbol{\theta}), \ldots, f_{10}(\boldsymbol{\theta})$. We build independent GP emulators for each of the 10 outputs.

Gaussian processes can be thought of as infinite-dimensional prior distributions for the unknown function $f(\boldsymbol{\theta})$. The distribution is fully specified by a prior mean function $m(\boldsymbol{\theta}) = \mathbb{E}[f(\boldsymbol{\theta})]$, and prior covariance function $k(\cdot, \cdot)$ where $k(\boldsymbol{\theta}, \boldsymbol{\theta}') = \mathbb{C}\text{ov}[f(\boldsymbol{\theta}), f(\boldsymbol{\theta}')]$. Then, the finite-dimensional distributions of f have a multivariate normal distribution so that given any finite set of input locations, $\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_m$, the vector

$$\mathbf{f} = \begin{pmatrix} f(\boldsymbol{\theta}_1) \\ \vdots \\ f(\boldsymbol{\theta}_m) \end{pmatrix} \sim \mathcal{N}_m(\mathbf{m}, \mathbf{K}),$$

where $\mathcal{N}_m(\mathbf{m}, \mathbf{K})$ denotes the m-dimensional multivariate Gaussian distribution with mean vector $\mathbf{m} \in \mathbb{R}^m$, and covariance matrix $\mathbf{K} \in \mathbb{R}^{m \times m}$ with

$$\mathbf{m} = \begin{pmatrix} m(\boldsymbol{\theta}_1) \\ \vdots \\ m(\boldsymbol{\theta}_m) \end{pmatrix} \quad \text{and} \quad \mathbf{K}_{ij} = k(\boldsymbol{\theta}_i, \boldsymbol{\theta}_j).$$

The number of simulator evaluations available to train the GP is the primary determinant of predictive accuracy of the resulting emulator. The choice of prior mean and covariance function also affect predictive performance, but with so few simulations available, there is limited opportunity to optimise these choices. After experimentation with a variety of choices, we used a linear prior mean function of the form

$$m(\boldsymbol{\theta}) = \beta_0 + \boldsymbol{\theta}^{\mathrm{T}}\boldsymbol{\beta}, \tag{6}$$

where $(\beta_0, \boldsymbol{\beta}^{\mathrm{T}})^{\mathrm{T}} \in \mathbb{R}^6$. We tried a variety of kernels, including members of the Matérn family and the rational quadratic kernel (Rasmussen & Williams, 2006), as well as sums and products of standard kernels. We found the best performance was obtained using a radial basis function (RBF) kernel (also known as the exponentiated quadratic or squared exponential kernel) plus a white noise or nugget kernel, i.e. using

$$k(\boldsymbol{\theta}, \boldsymbol{\theta}') = \alpha^2 \exp\left\{-\frac{1}{2}(\boldsymbol{\theta} - \boldsymbol{\theta}')^T \mathrm{P}^{-2}(\boldsymbol{\theta} - \boldsymbol{\theta}')\right\} + \sigma^2\delta(\boldsymbol{\theta} - \boldsymbol{\theta}') \qquad \text{where} \quad \delta(y) = \begin{cases} 1 & \text{if } y = 0 \\ 0 & \text{otherwise.} \end{cases} \tag{7}$$

Here, α^2 and σ^2 are variance hyperparameters, and $\mathrm{P} = \text{diag}(\rho_1, \ldots, \rho_5)$ is a 5×5 diagonal matrix giving the length scale for each of the input coordinates. The length scale controls the decay in correlation when the distance between the input points increases. The nugget parameter, σ^2, describes the at a point variance. The RBF kernel results in a GP that has infinitely differentiable samples, whereas the nugget kernel produces a nowhere continuous white-noise process. We can think of the sum as producing a smoothly varying function of $\boldsymbol{\theta}$ plus a discontinuous white-noise term. To summarise the GP distribution for f, we write $f(\cdot) \mid \boldsymbol{\psi} \sim GP(m(\cdot), k(\cdot, \cdot))$, emphasising in the notation that the prior specification is dependent upon unknown hyper parameters $\boldsymbol{\psi} = \{\alpha, \mathrm{P}, \beta, \sigma^2\}$.

We train the emulator using a set of simulator evaluations, $\mathcal{D} = \{\boldsymbol{\theta}_j, f(\boldsymbol{\theta}_j)\}_{j=1}^m$. Given a value of $\boldsymbol{\psi}$, conditioning our prior GP specification on $\mathcal{D}$ results in another GP distribution for f, but with updated mean and covariance functions

$$f \mid \mathcal{D}, \boldsymbol{\psi} \sim GP(m^*(\cdot), k^*(\cdot, \cdot)).$$

See, for example, Rasmussen and Williams (2006) for details of m^* and k^*. The hyperparameters in GPs are often estimated by maximum likelihood (Myung, 2003), but we take a Bayesian

approach and give them prior distributions and marginalise them out of the analysis. We model each of the 10 outputs $f_1(\boldsymbol{\theta}), \ldots, f_{10}(\boldsymbol{\theta})$ as independent GPs, allowing different hyperparameter values for each output, i.e. using $\boldsymbol{\psi}^i = \{\alpha_i, \mathrm{P}_i, \boldsymbol{\beta}_i, \sigma_i^2\}$ for $i = 1, \ldots, 10$, all of which need to be marginalised. We considered using a multi-output GP to take account of the correlations between the 10 different outputs, but decided against, as doing so would significantly complicate the analysis, likely for little benefit given we only have 47 simulations runs with which to fit the GO. A standard multi-output GP (Alvarez et al., 2012) would require us to estimate a 10×10 covariance matrix, and would likely cause difficulties in developing a well-mixing MCMC scheme. We give the hyperparameters the following prior distributions:

$$\boldsymbol{\beta}_{ij} \sim \mathcal{N}(0, 100) \quad \boldsymbol{\sigma}_i^2 \sim \Gamma^{-1}(1, 1), \tag{8}$$

$$\alpha_i \sim \mathcal{N}(0, 1) \quad \rho_{ij} \sim \Gamma^{-1}(5, 5), \tag{9}$$

where $i = 1, \ldots, 10$ indexes across the different GP emulators for each site, and $j = 1, \ldots, 5$ indexes across the input dimensions. Here, $\Gamma^{-1}(a, b)$ denotes an inverse gamma prior distribution with shape parameter a and scale parameter b.

We fit the GP and infer the hyperparameters and posterior ice-sheet shape jointly using MCMC as described in the next section. First, we assess the predictive skill of the GP emulator using leave-one-out cross-validation (Vehtari et al., 2017). We fit the GP to 46 of the 47 simulations from HadCM3 and predict the mean and variance of the left out simulation. We repeat this leaving out each input in turn, and for all 10 ice-core sites at which we have measurements. The results are shown in Figure 1 along with the RMSE and the observed coverage of the 95% prediction intervals. We can see that the GP emulators perform well. The prediction intervals are relatively wide compared to the signal, but appear to be well calibrated. The nugget variance dominates the predictive uncertainty, and gives an effective minimum width to the prediction intervals. We tried constraining the nugget variance to be smaller, but this degraded the overall performance of the GP predictions. With a larger simulation budget it may be possible to reduce the predictive uncertainty.

3.2 Inference

We have demonstrated that we can successfully approximate the relationship between ice-sheet orography and the $\delta^{18}O$ anomalies that HadCM3 simulates. In this section we describe how to use the GP emulator and the prior model to infer the posterior distribution of the shape of the AIS at the LGM, which we denote by $\mathbf{x}$. Let $\boldsymbol{\theta}$ be the corresponding coordinates of $\mathbf{x}$ when projected onto the basis used by the prior model [equation (2)]. Our aim is to infer the posterior distribution for $\mathbf{x}$, $\pi(\mathbf{x} \mid \mathbf{y}^{obs}, \mathcal{D})$, which we do by computing $\pi(\boldsymbol{\theta} \mid \mathbf{y}^{obs}, \mathcal{D})$. We have

$$\begin{aligned}\pi(\boldsymbol{\theta} \mid \mathbf{y}^{obs}, \mathcal{D}) &\propto \pi(\theta)\pi(\mathbf{y}^{obs} \mid \boldsymbol{\theta}, \mathcal{D}) \\ &= \pi(\theta) \int \pi(\mathbf{y}^{obs} \mid \boldsymbol{\psi}, \boldsymbol{\theta}, \mathcal{D})\pi(\boldsymbol{\psi} \mid \boldsymbol{\theta}, \mathcal{D})\mathrm{d}\boldsymbol{\psi} \\ &= \pi(\theta) \int \pi(\mathbf{y}^{obs} \mid \boldsymbol{\psi}, \boldsymbol{\theta}, \mathcal{D})\pi(\mathcal{D} \mid \boldsymbol{\psi})\pi(\boldsymbol{\psi})\mathrm{d}\boldsymbol{\psi}.\end{aligned} \tag{10}$$

Our prior distribution for $\boldsymbol{\theta}$, $\pi(\boldsymbol{\theta})$, is

$$\boldsymbol{\theta}_{obs} \sim \mathcal{N}(\mathbf{0}, \mathrm{diag}(0.5, 0.5, 0.6, 0.5, 1))$$

as described in Section 2.1. The prior distribution for the GP hyperparameters, $\pi(\boldsymbol{\psi})$ is given by equation (9). The GP likelihood is given by equation (10) after using the kernel in equation (7) and the prior mean in equation (6). The remaining term is $\pi(\mathbf{y}^{obs} \mid \boldsymbol{\psi}, \boldsymbol{\theta}^*, \mathcal{D})$ which is the GP prediction for the observed data given $\boldsymbol{\theta}$. To complete our prior specification for the data, we need to choose a distribution for the error term, $\boldsymbol{\epsilon}$, in equation (1). This error represents the difference between the HadCM3 prediction when run with the true AIS shape, and the observations.

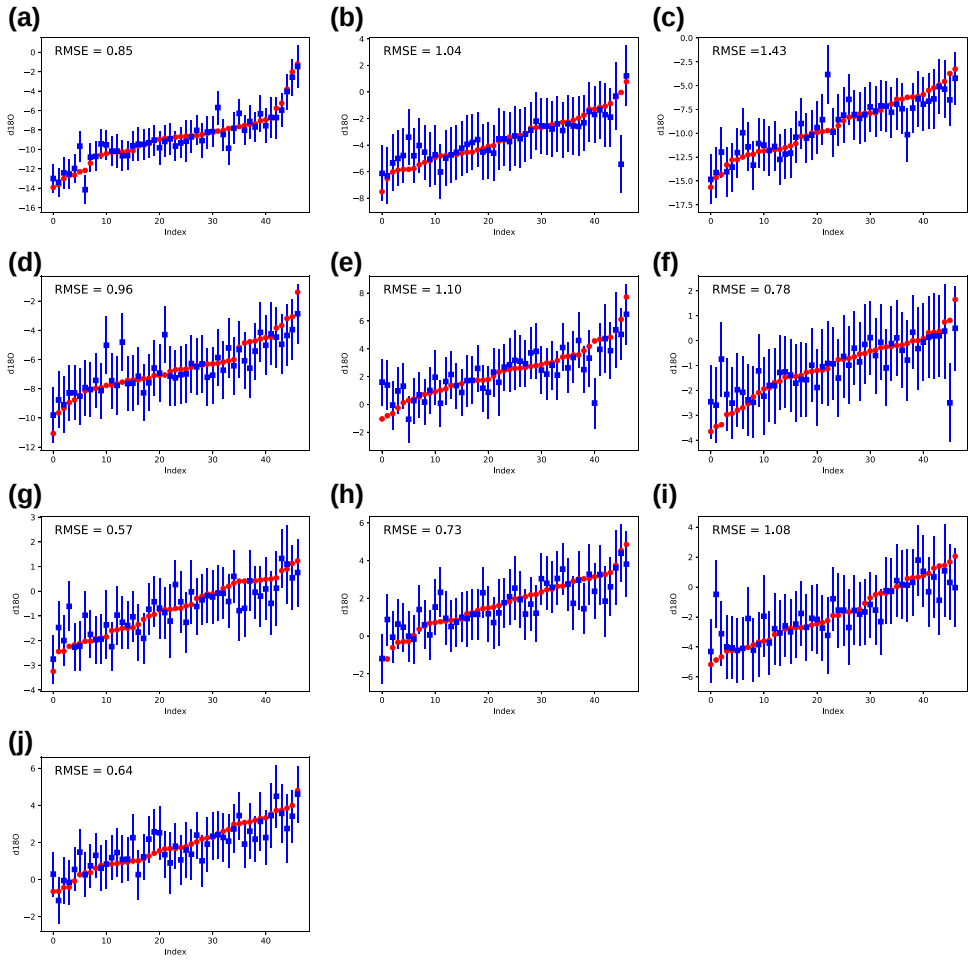

Figure 1. Leave-one-out cross-validation results. Each plots shows the HadCM3 simulated δ^{18}O anomalies (circular points) with the corresponding GP prediction (square points) as well as 95% prediction intervals (line segments). The GP used anistropic RBF kernel plus white noise, with a linear mean function, with hyperparameter prior distributions as given in equation (9). A 5D representation of the orographies was used for all 10 ice-core sites. The RMSE and 95% empirical coverage values are also reported. (a) Byrd. Coverage: 94%. (b) Mount Moulton. Coverage: 96%. (c) Siple. Coverage: 96%. (d) WDC. Coverage: 94%. (e) EDC. Coverage: 94%. (f) EDML. Coverage: 96%. (g) Fuji. Coverage: 98%. (h) Talos. Coverage: 94%. (i) Taylor. Coverage: 98%. (j) Vostok. Coverage: 100%.

It thus must account for observation error (caused by natural variation, as well as the spatial and temporal averaging in the ice deposition and model discretisation), as well as simulator discrepancy. Given the limited data available, we cannot hope to resolve these individual errors, and so we simply assume $\boldsymbol{\epsilon} \sim N(0, \sigma_1^2)$ with the variance σ_1^2 to be estimated with the other parameters.

We add in a fixed standard deviation term for $\boldsymbol{y}_{\text{pred}}$ that represents the measurement error of the ice cores, τ_1, and the spatial variance between modelling a latitude/longitude grid cell from HadCM3 and the location of the ice-core site, τ_2. We considered also including an error term on the dating process, as the measurements are at different dates in each ice core. However, as the LGM was a very broad period, with most global ice sheets in equilibrium for several thousand years (Clark et al., 2009), we decided that this was unnecessary. If modelling time periods with more rapid changes, such as the deglaciation period up to 10 Ka BP, then a dating error would be necessary. We give a value of $\tau_1 = 0.5‰$ for the measurement error from the results of Keller et al. (2018), and a value of $\tau_2 = 2.5‰$ for the spatial error from looking at the standard deviation between HadCM3 model grid cells surrounding the cell containing the ice-core site. Our model is

therefore

$$\boldsymbol{y}_i \mid \cdot \sim \mathcal{N}(\Theta'\boldsymbol{\beta}_i + f, \sigma^2_{1,i} + \tau_1^2 + \tau_2^2) \quad i = 1, \ldots, 10. \tag{11}$$

To compute the posterior distribution, we use Hamiltonian MCMC (Betancourt, 2017), which is a method of sampling from a probability distribution by drawing a series of correlated samples that converge to the target distribution (Geyer, 1992). We implement this in the probabilistic programming language Stan (Carpenter et al., 2017), which uses No-U-Turn sampling (NUTS) to efficiently explore the posterior distribution. We finally generate

$$y_{\text{pred},i} = \mathcal{N}(\boldsymbol{\theta}_{\text{obs}}\boldsymbol{\beta}_i + f, \sigma^2_{1,i} + \tau_1^2 + \tau_2^2), \tag{12}$$

using the posterior sample Θ_{obs} to create distributions of $\delta^{18}O$ anomalies. We can then compare the distribution of y_{pred} to the value y_{obs}, which would show if our model is simulating accurate values of $\delta^{18}O$ anomalies.

4 Results

We present results in two sections; Section 4.1 contains a description of the analysis of our model and a discussion of the model's ability to reconstruct ice-sheet shapes. Section 4.2 discusses the difference between the posterior ice-sheet shapes our model estimates, and the previously published reconstructions.

4.1 Sensitivity testing

We performed a sensitivity analysis of our model to explore the source of various uncertainties. We inflated and shrank the variance of the $\boldsymbol{\theta}_{\text{obs}}$, as well as the fixed variance term $\tau^2 = \tau_1^2 + \tau_2^2$, to check how they influence the model. Our posterior distributions were not affected much by changing the fixed variance, although the posterior variance of $\boldsymbol{\theta}_{\text{obs}}$ was altered by changes in the prior variances.

Figure 2 shows the generated predictions, y_{pred}, at each of the ice-core sites, with the observation as a red vertical line. The predictions all peak close to the true value, suggesting our model is successfully predicting the isotope values.

We carried out two sensitivity tests to check the robustness of our model. We removed one of the HadCM3 climate simulations, and set this simulation as our y_{obs} to generate values of $\boldsymbol{\theta}_{\text{obs}}$ and compare them to the known prior values. The posterior distributions of $\boldsymbol{\theta}_{\text{obs}}$ and generated distributions of $\boldsymbol{y}_{\text{pred}}$ can be seen in Figure 3a and b. We can see that our model is recovering the $\delta^{18}O$ simulations and prior variable values well; the histograms of the prior variables in Figure 3a are all different from the prior distributions and are peaking around the $\boldsymbol{\theta}$ values used to create the orography adopted for the missing HadCM3 simulation. Similarly, Figure 3b shows that the $\delta^{18}O$ are all peaking around the values from the missing HadCM3 simulation.

We further test the model by trying to recover a $\delta^{18}O$ anomaly from an ice core. We wish to see if, by modelling only nine of the cores, we can create $\delta^{18}O$ anomalies close to the missing ice-core observation. We do this by removing the observation from equation 11, so that we are not learning anything about it from the input and output data. We then use equation 12 to generate $\delta^{18}O$ anomalies for the missing ice core and compare it to the observed value. The distribution of y_{pred} for the missing ice-core observation peaks close to the true value, as can be seen in online supplementary Figure 3. However, the range of values is extremely large; as the $\boldsymbol{\beta}_i$'s were given vague priors with standard deviation of 100, the model is not able to update much from these. Further tests using other ice cores showed that this applied to cores in both the East and West Antarctic ice sheets; the model struggled to recover an unknown core based on the sampling of the other nine. We were limited by the number of ice cores available dating back to the LGM; more ice-core data for the LGM, sampled evenly across both ice sheets, would lead to a more robust model. We test this by creating a psuedo-ice-core value to see how this affected distributions of y_{pred}; results can be seen in online supplementary Figure 4. We can see the distributions

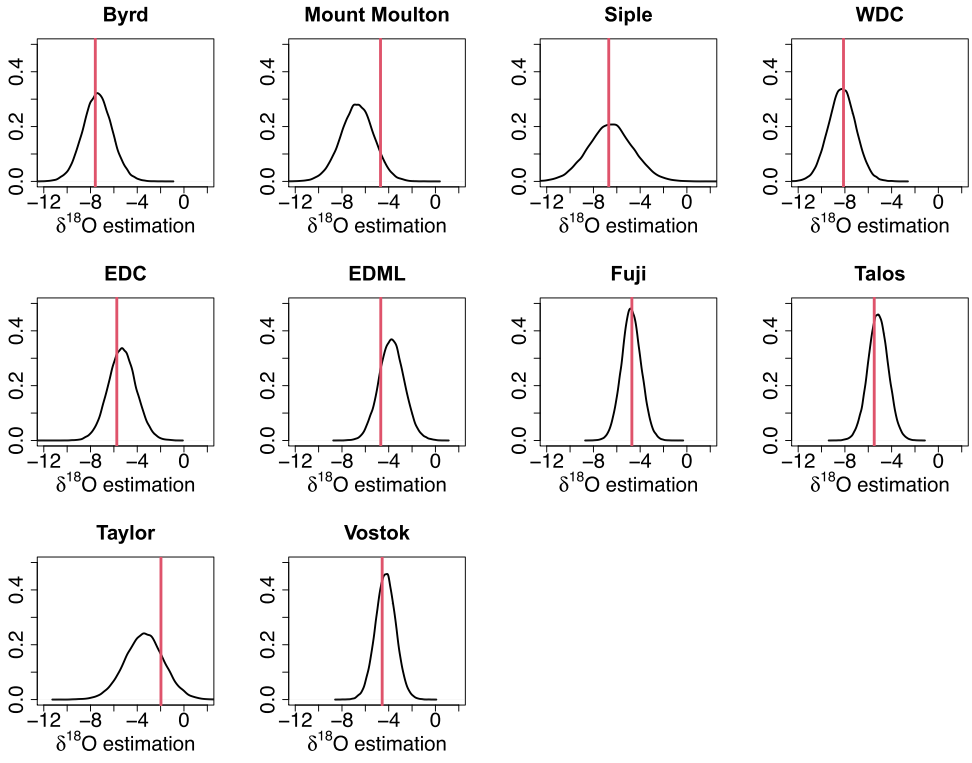

Figure 2. Comparing the generated densities of y_{pred} against the ice-core values. The black lines are the densities of y_{pred}. The vertical red lines are the observed $\delta^{18}O$ anomalies. The densities all peak close to the true observation from the ice cores, suggesting we have successfully predicted the isotope values.

are peaking closer to observations in many cases, demonstrating the improvement additional data would have on our analysis.

4.2 Posterior shapes

We use the samples of the posterior distributions of $\boldsymbol{\theta}_{\text{obs}}$ to create posterior ice sheet shapes and consider how plausible the shapes are. Figure 4c shows the standard deviation of the posterior reconstructions. One of our aims with this method was to reduce uncertainty around the AIS at the LGM; comparing the standard deviation of the posterior ice-sheet shapes from our prior model in Figure 4b, we can see there is a lot less variance in our posterior shapes compared to the prior model, with Figure 4d showing the same standard deviation on a smaller scale to show the values in more detail.

We ran the mean posterior ice-sheet shape through HadCM3 to test the validity of our analysis. The resulting simulation produced δ^{18}O anomalies far closer to the observations than any of our training simulations. The smallest RMSE of our original simulations was 4.46; in contrast, the posterior orography resulted in a RMSE for the δ^{18}O output of 1.96. Even in the situation where the model is perfect, we would not expect the posterior mean RMSE to be zero, due to measurement errors and the remaining uncertainty in the ice-sheet shape. In addition, there are almost certainly biases in HadCM3 (as well as errors in the hundreds of other climate parameters we had to specify in order to do LGM simulations), which will cause the posterior mean RMSE to differ significantly from zero. But the fact we have a much reduced RMSE shows the value of analysis in extracting information from our three sources.

4.2.1 Comparing our orographies to previous reconstructions

We compare our posterior shapes to four of the ice-sheet reconstructions used to build our prior model, ICE-5G (Peltier, 2004), ICE-6G (Argus et al., 2014), W12A (Whitehouse et al., 2012), and

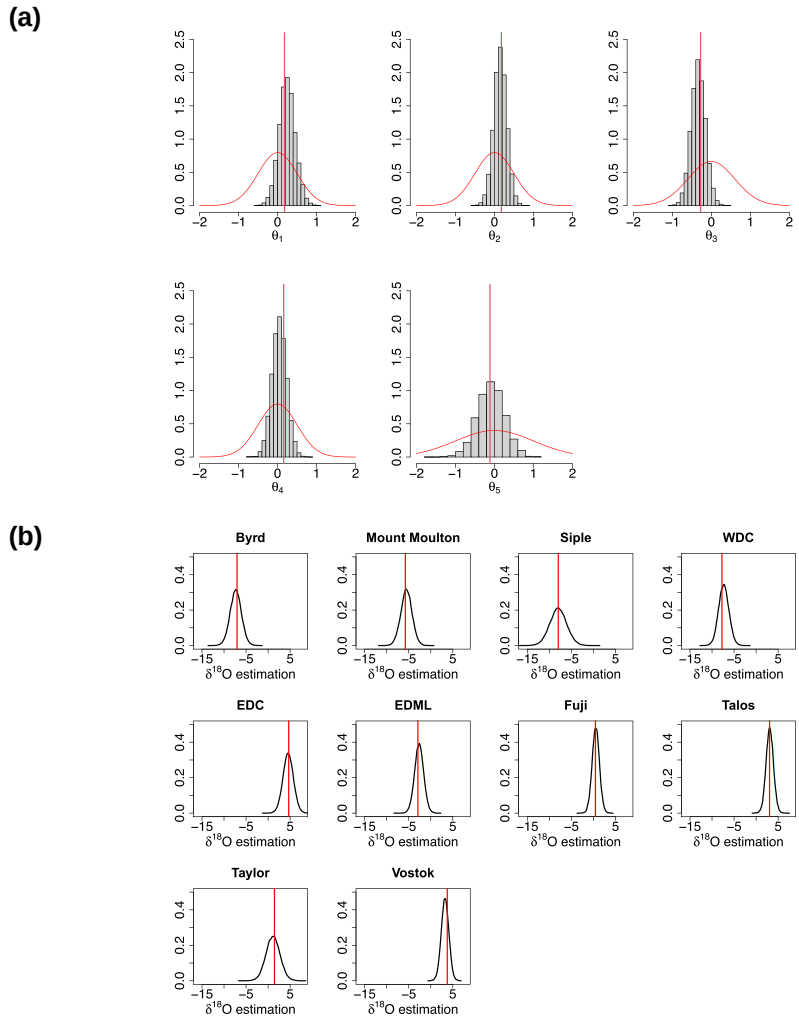

Figure 3. (a) Histograms of the posterior distributions of $\boldsymbol{\theta}_{\text{obs}}$, with density curves of the prior distributions superimposed, when trying to recover a HadCM3 simulation. The true values of the synthetic $\boldsymbol{\theta}_i$'s are added on as red vertical lines. We can see the posterior distributions are very different to the priors, and are peaking close to the observations. (b) The distributions of y_{pred} when recovering a HadCM3 simulation, with the synthetic observations added on as red vertical lines. The densities all peak close to the synthetic observations, demonstrating that our model has successfully recovered the HadCM3 simulation removed from our input data matrix.

GLAC-1D (Briggs et al., 2014). These are the four reconstructions in the dataset $\boldsymbol{X}_1$, used in the first round of PCA. This allows us to see where the posterior reconstructions differ from the original ice-sheet reconstructions. We compare our results to the four reconstructions we originally started with to give some examples of how our method differs to previous work.

Figure 5 shows the difference between the posterior mean of our reconstruction and the previous reconstructions. We can see that aside from the coastal regions, our reconstructions suggest a higher elevation for the AIS at the LGM, particularly in the Eastern ice-sheet (EAIS). ICE-5G has a higher elevation than our posterior mean in the Western ice-sheet (WAIS), but ICE-5G is known to have a thicker AIS than the more recent reconstructions (ICE-6G, GLAC-1D, W12A) due to methodological updates as described in Argus et al. (2014). Our model also has a lower elevation around much of the Antarctic coastline than ICE-6G. Despite this, samples from our posterior distribution are ice-sheet orographies which have consistently higher elevation over the EAIS than the reconstructions used in our prior model, suggesting this is where some of the 'missing ice' (Gowan et al., 2021; Simms et al., 2019), previously unaccounted

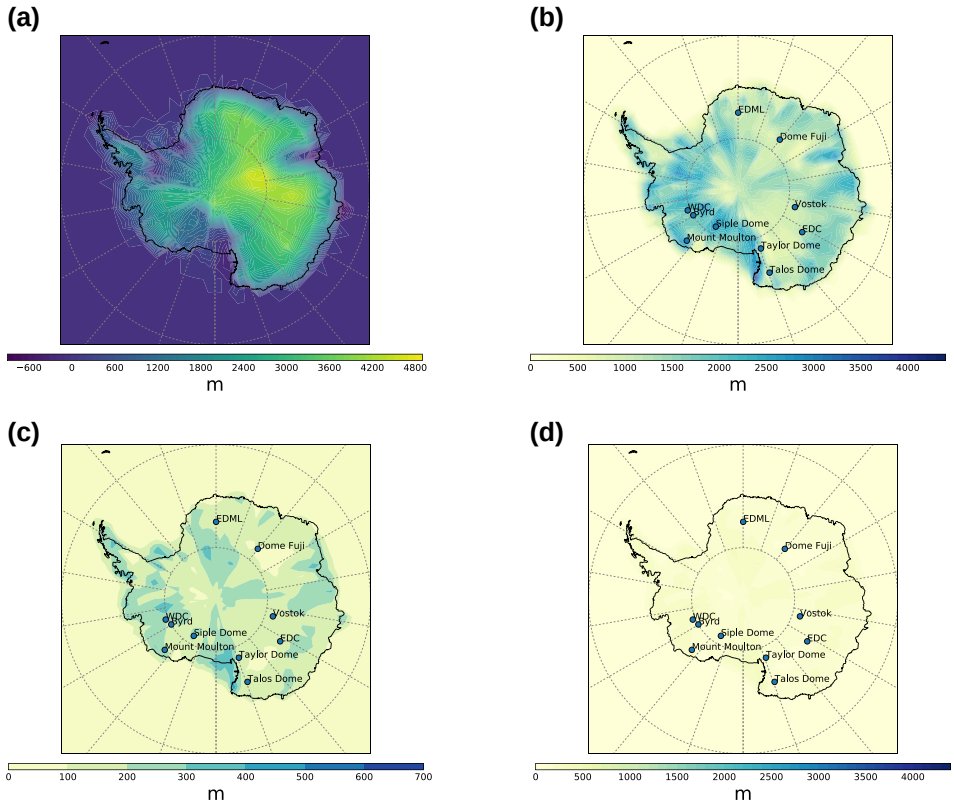

Figure 4. (a) Mean posterior shape from our model. (b) Standard deviation of our prior model. (c) Standard deviation of our posterior ice-sheet shapes, with (d) showing the same values on a finer scale. Comparing (b) to (c) shows our method has greatly reduced the uncertainty around the AIS at the LGM.

for in other ice-sheet reconstructions, may be found. Although smaller than ICE-5G, the posterior mean shape from our model is considerably larger than in the other three reconstructions. Simms et al. (2019) states that the post-LGM sea-level rise is not balanced by the amount of ice melted, with the 'missing ice' corresponding to approximately 15 m of global sea-level rise. They conclude that 'either a large reservoir of water (e.g. a missing LGM ice sheet) has yet to be discovered or current estimates of one or more of the known LGM ice sheets are too small. Our work corroborates the hypothesis that some of this missing ice may have been present in the AIS.

Table 2 shows the difference between the elevation of our posterior mean reconstruction and four of the original ice-sheet reconstructions over the AIS, below 60°. We estimate the AIS height to be significantly higher than three of the four (the most recent three) most highly cited reconstructions. Calculating the additional ice volume present in our reconstructions, and converting this to an equivalent sea-level rise, requires knowledge of bed elevation through time, and due to the viscoelastic nature of the Earth's mantle this depends on past ice-sheet change. The ice volume associated with each reconstructed ice-sheet surface must therefore be determined using models that calculates solid Earth deformation in response to time-integrated surface load change (e.g. Whitehouse, 2018). This detailed calculation requires significant additional work, and is beyond the scope of this paper. But a crude *back-of-the-envelope* calculation suggests that our best estimate of the AIS contains approximately 1.1 million km^3 of additional ice than is represented in ICE-6G, and that this difference is approximately equivalent to a 3 m difference in global sea levels at the LGM. However, we stress that this is likely to be an underestimate of the sea level contribution of our revised AIS. To be more precise will require a more careful analysis using a glacial isostatic adjustment model.

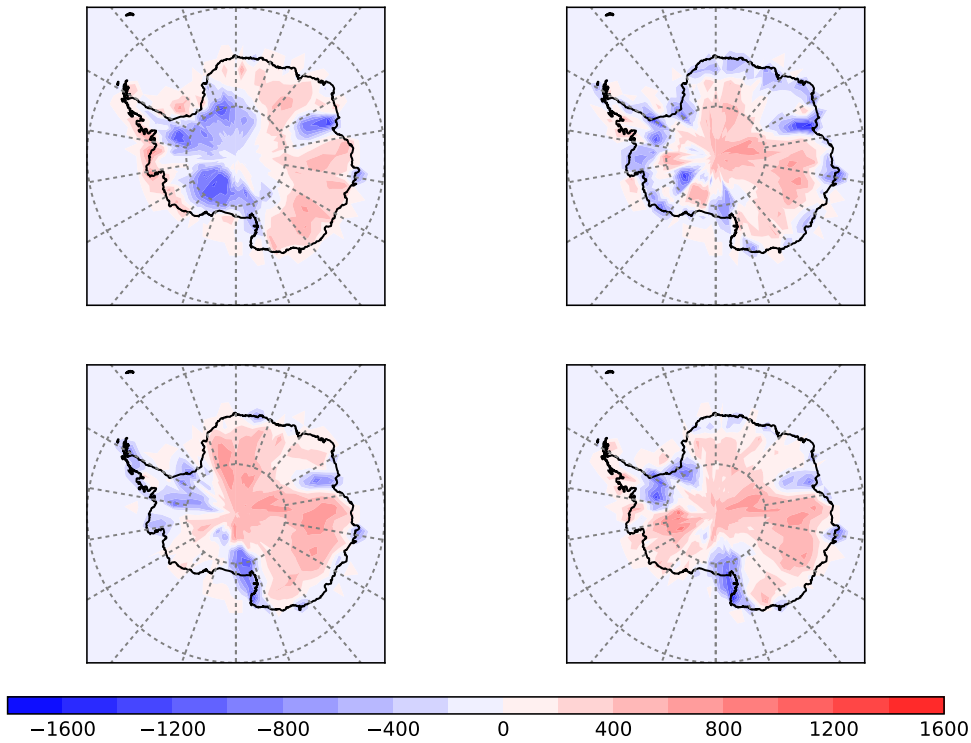

Figure 5. The difference between the posterior mean height our model creates and the four ice-sheet reconstructions in X_1. Clockwise from the top left they are ICE-5G (Peltier, 2004), ICE-6G (Argus et al., 2014), GLAC-1D (Briggs et al., 2014), and W12A (Whitehouse et al., 2012). We can see that our average ice-sheet shape has a higher elevation in the EAIS and lower elevation in the coastal regions than three of the reconstructions; the exception is ICE-5G, that has much higher elevation in the WAIS than our model output.

Table 2. The mean height difference in metres between the posterior mean posterior reconstruction from our analysis and prior AIS reconstructions

AIS reconstruction	Mean height difference from mean posterior shape (in metres)
ICE-5G	−80.22
ICE-6G	77.13
W12A	126.72
GLAC-1D	122.04

Note. A positive value denotes our posterior reconstruction being larger than the prior reconstructions. This calculation is performed without taking into account any glacial isostatic adjustment response. AIS = Antarctic ice sheet.

5 Discussion

We have demonstrated the success our method has at modelling the relationship between the AIS orography at the LGM and $\delta^{18}O$ anomalies from ice cores, and considered the geographical interpretations of these results including the posterior ice-sheet shapes that our methods have produced. Here, we discuss our findings, and consider some issues and further work.

We have created a novel approach to weighted PCA. This allows the use of data when there is an obvious split in the relevance or importance of sources or variables. It also allows for the incorporation of more data at a later date, and can assist in ranking data by prior beliefs about its influence on the subject of interest, and the use of related but unequal data sets. Situations such as ours

are not exclusive to the climate science community, and this weighted PCA method is applicable to many other areas, including research areas using data sets of varying levels of favourable quality such as other environmental sciences, medicine or epidemiology.

We have demonstrated a way of combining prior knowledge, expert elicitation, GP emulation and Bayesian inference to model the uncertainty around aspects of the palaeo-climate. Building on the previous work of Domingo et al. (2020), who used a similar emulation technique to reconstruct the Greenland ice sheet at the LIG, we have refined how the prior model was built and the process of selecting a set of orographies to input into HadCM3. We have shown the success of building an emulator of a climate model, which gives the possibility of future uses in other aspects of statistical analysis of climatology. We chose a relatively stable and recent time period—the LGM was chronologically a very broad event; conditions changed very slowly, allowing us to disregard any dating errors in the isotopes. The stable nature of the time period meant we could model it statistically. If applying this method to a less stable period, when ice sheets were more in flux and the climate changing more quickly, then further uncertainties would have to be incorporated in to the calibration model.

Our estimate of the posterior distribution uses a probabilistic approximation to the simulator and a reduced-dimensional representation of the prior reconstructions. As a result, although the analysis is based on physically constrained quantities, these constraints are only embedded approximately. Consequently it is theoretically possible for our model to output implausible ice sheet shapes, but this was not observed during our analysis.

In particular, the edges of the ice sheets that our model creates require smoothing. The coarseness of the HadCM3 latitude/longitude grid and the lack of physics in our model means it is hard to control where the ice sheet cuts off in our output, resulting in shapes that are jagged and unrealistic, with occasional holes where the ice sheet thins. Our posterior shapes are on the same coarse latitude/longitude grid; this means our posterior ice-sheet shapes can therefore only be viewed as guides for the location of ice, rather than a precise map of the AIS at the LGM.

Our posterior shapes show a clear difference from the reconstructions in our prior model; they are larger than most of the past reconstructions, with lower elevation around the coast of the continent and higher elevation in the central part of the ice sheet, particular in the EAIS. The plausibility of these shapes, and their potential contribution to sea level rise up to the present day, is a topic for further collaboration. Simms et al. (2019) emphasise the missing ice problem around the LGM, and the possiblity of that ice being present in the AIS. Our work corroborates this hypothesis, providing some guidance on where that ice may have been. Giving a more precise and careful estimate would take some additional work and will be done elsewhere.

In summary, the methods proposed in this paper could be applied to any time period or climate variable given that there is adequate proxy data and prior knowledge. Indeed they could be used for any application that has expert beliefs, a complex simulator that is too computationally expensive to run numerous times, and a set of observations with which to compare emulator output. Furthermore, in developing our suite of models and methods we have tackled some problems, and provided tools, that are likely to be relevant to others working within and beyond climate science.

Acknowledgments

The authors thank Dr Robert Arthern, Dr Richard Hindmarsh (deceased), Dr Dominic Hodgson, Dr Robert Mulvaney, and Dr James Smith at the British Antarctic Survey for taking part in the expert elicitation process.

Funding

This work was completed as part of a PhD funded by the Grantham Centre for Sustainable Futures.

Data availability

Data and code are available at https://doi.org/10.5281/zenodo.8246227

Conflict of interests: None declared.

Supplementary material

Supplementary material is available online at *Journal of the Royal Statistical Society: Series C*.

References

Alvarez M. A., Rosasco L., & Lawrence N. D. (2012). Kernels for vector-valued functions: A review. *Foundations and Trends in Machine Learning*, *4*(3), 195–266. https://doi.org/10.1561/2200000036

ARCHER (2016). *ARCHER supercomputer*. http://www.archer.ac.uk

Argus D. F., Peltier W. R., Drummond R., & Moore A. W. (2014). The Antarctica component of postglacial rebound model ICE-6G_C (VM5a) based on GPS positioning, exposure age dating of ice thicknesses, and relative sea level histories. *Geophysical Journal International*, *198*(1), 537–563. https://doi.org/10.1093/gji/ggu140

EPICA Community Members, Barbante C., Barnola J. M., Becagli S., Beer J., Bigler M., Boutron C., Blunier T., Castellano E., Cattani O., Chappellaz J., Dahl-Jensen D., Debret M., Delmonte B., Dick D., Falourd S., Faria S., Federer U., Fischer H., Freitag J., … Wolff E. (2006). One-to-one coupling of glacial climate variability in Greenland and Antarctica. *Nature*, *444*(7116), 195–198. https://doi.org/10.1038/nature05301

Bentley M. J., Cofaigh C. Ã., Anderson J. B., Conway H., Davies B., Graham A. G., Hillenbrand C. -D., Hodgson D. A., Jamieson S. S., Larter R. D., Mackintosh A., Smith J. A., Verleyen E., Ackert R. P., Bart P. J., Berg S., Brunstein D., Canals M., Colhoun E. A., … Zwartz D. (2014). A community-based geological reconstruction of Antarctic ice sheet deglaciation since the Last Glacial Maximum. *Quaternary Science Reviews*, *100*(Supplement C), 1–9. https://doi.org/10.1016/j.quascirev.2014.06.025

Betancourt M. (2017). 'A conceptual introduction to Hamiltonian Monte Carlo', arXiv, arXiv:1701.02434, preprint: not peer reviewed.

Blasco J., Alvarez-Solas J., Robinson A., & Montoya M. (2021). Exploring the impact of atmospheric forcing and basal boundary conditions on the simulation of the Antarctic ice sheet at the Last Glacial Maximum. *The Cryosphere*, *15*, 215–231.

Blunier T., & Brook E. J. (2001). Timing of millennial-scale climate change in Antarctica and greenland during the last glacial period. *Science*, *291*(5501), 109–112. https://doi.org/10.1126/science.291.5501.109

Boyd S., Boyd S. P., & Vandenberghe L. (2004). *Convex optimization*. Cambridge University Press.

Briggs R. D., Pollard D., & Tarasov L. (2014). A data-constrained large ensemble analysis of Antarctic evolution since the Eemian. *Quaternary Science Reviews*, *103*, 91–115. https://doi.org/10.1016/j.quascirev.2014.09.003

Brook E. J., White J. W., Schilla A. S., Bender M. L., Barnett B., Severinghaus J. P., Taylor K. C., Alley R. B., & Steig E. J. (2005). Timing of millennial-scale climate change at Siple Dome, West Antarctica, during the last glacial period. *Quaternary Science Reviews*, *24*(12–13), 1333–1343. https://doi.org/10.1016/j.quascirev.2005.02.002

Buizert C., Fudge T., Roberts W. H., Steig E. J., Sherriff-Tadano S., Ritz C., Lefebvre E., Edwards J., Kawamura K., Oyabu I., Motoyama H., Kahle E. C., Jones T. R., Abe-Ouchi A., Obase T., Martin C., Corr H., Severinghaus J. P., Beaudette R., … Schwander J. (2021). Antarctic surface temperature and elevation during the Last Glacial Maximum. *Science*, *372*(6546), 1097–1101. https://doi.org/10.1126/science.abd2897

Carpenter B., Gelman A., Hoffman M. D., Lee D., Goodrich B., Betancourt M., Brubaker M., Guo J., Li P., & Riddell A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, *76*(1), 1–32. https://doi.org/10.18637/jss.v076.i01

Chang W., Haran M., Olson R., & Keller K. (2014). Fast dimension-reduced climate model calibration and the effect of data aggregation. *The Annals of Applied Statistics*, *8*(2), 649–673. https://doi.org/10.1214/14-AOAS733

Clark P. U., Dyke A. S., Shakun J. D., Carlson A. E., Clark J., Wohlfarth B., Mitrovica J. X., Hostetler S. W., & McCabe A. M. (2009). The Last Glacial Maximum. *Science*, *325*(5941), 710–714. https://doi.org/10.1126/science.1172873

Constantine P. G. (2015). *Active subspaces: Emerging ideas for dimension reduction in parameter studies*. SIAM.

de Boer B., Haywood A. M., Dolan A. M., Hunter S. J., & Prescott C. L. (2017). The transient response of ice volume to orbital forcing during the warm late Pliocene. *Geophysical Research Letters*, *44*(20), 10,486–10,494. 2017GL073535. https://doi.org/10.1002/2017GL073535

Domingo D., Malmierca-Vallet I., Sime L., Voss J., & Capron E. (2020). Using ice cores and Gaussian process emulation to recover changes in the Greenland ice sheet during the last interglacial. *Journal of Geophysical Research: Earth Surface*, *125*(5), e2019JF005237. http://doi.org/10.1029/2019JF005237

Edwards T. L., Brandon M. A., Durand G., Edwards N. R., Golledge N. R., Holden P. B., Nias I. J., Payne A. J., Ritz C., & Wernecke A. (2019). Revisiting Antarctic ice loss due to marine ice-cliff instability. *Nature*, *566*(7742), 58–64. https://doi.org/10.1038/s41586-019-0901-4

Edwards T. L., Nowicki S., Marzeion B., Hock R., Goelzer H., Seroussi H., Jourdain N. C., Slater D. A., Turner F. E., Smith C. J., McKenna C. M., Simon E., Abe-Ouchi A., Gregory J. M., Larour E., Lipscomb W. H., Payne

A. J., Shepherd A., Agosta C., … Zwinger T. (2021). Projected land ice contributions to twenty-first-century sea level rise. *Nature*, *593*(7857), 74–82. https://doi.org/10.1038/s41586-021-03302-y

WAIS Divide Project Members, Fudge T. J., Steig E. J., Markle B. R., Schoenemann S. W., Ding Q., Taylor K. C., McConnell J. R., Brook E. J., Sowers T., White J. W. C., Alley R. B., Cheng H., Clow G. D., Cole-Dai J., Conway H., Cuffey K. M., Edwards J. S., Lawrence Edwards R., Edwards R., … Wong G. J. (2013). Onset of deglacial warming in West Antarctica driven by local orbital forcing. *Nature*, *500*(7463), 440–444. https://doi.org/10.1038/nature12376

Geyer C. J. (1992). Practical Markov chain Monte Carlo. *Statistical Science*, *7*(4), 473–483. http://doi.org/10.1214/ss/1177011137

Golledge N. R., Fogwill C. J., Mackintosh A. N., & Buckley K. M. (2012). Dynamics of the Last Glacial Maximum Antarctic ice-sheet and its response to ocean forcing. *Proceedings of the National Academy of Sciences*, *109*(40), 16052–16056. https://doi.org/10.1073/pnas.1205385109

Golledge N. R., Levy R. H., McKay R. M., Fogwill C. J., White D. A., Graham A. G., Smith J. A., Hillenbrand C. -D., Licht K. J., Denton G. H., Ackert R. P., Maas S. M., & Hall B. L. (2013). Glaciology and geological signature of the Last Glacial Maximum Antarctic ice sheet. *Quaternary Science Reviews*, *78*(Supplement C), 225–247. https://doi.org/10.1016/j.quascirev.2013.08.011

Gowan E. J., Zhang X., Khosravi S., Rovere A., Stocchi P., Hughes A. L., Gyllencreutz R., Mangerud J., Svendsen J. -I., & Lohmann G. (2021). A new global ice sheet reconstruction for the past 80 000 years. *Nature Communications*, *12*(1), 1199. https://doi.org/10.1038/s41467-021-21469-w

Grootes P., Steig E., Stuiver M., Waddington E., Morse D., & Nadeau M. J. (2001). The Taylor Dome Antarctic 18O record and globally synchronous changes in climate. *Quaternary Research*, *53*(3), 289–298. https://doi.org/10.1006/qres.2001.2276

Hannachi A., Jolliffe I. T., & Stephenson D. B. (2007). Empirical orthogonal functions and related techniques in atmospheric science: A review. *International Journal of Climatology: A Journal of the Royal Meteorological Society*, *27*(9), 1119–1152. https://doi.org/10.1002/(ISSN)1097-0088

Higdon D. M., Gattiker J. R., Williams B. J., & Rightley M. L. J. (2008). Computer model calibration using high-dimensional output. *Journal of the American Statistical Association*, *103*(482), 570–583. https://doi.org/10.1198/016214507000000888

Holden P. B., Edwards N., Ridgwell A., Wilkinson R., Fraedrich K., Lunkeit F., Pollitt H., Mercure J.-F., Salas P., Lam A., Knobloch F., Chewpreecha U., & Vinuales J. (2018). Climate-carbon cycle uncertainties and the Paris Agreement. *Nature Climate Change*, *8*(7), 609–613. http://doi.org/10.1038/s41558-018-0197-7

Holden P. B., Edwards N. R., Garthwaite P. H., & Wilkinson R. D. (2015). Emulation and interpretation of high-dimensional climate model outputs. *Journal of Applied Statistics*, *42*(9), 2038–2055. https://doi.org/10.1080/02664763.2015.1016412

Joseph V. R., & Hung Y. (2008). Orthonormal-MAXIMIN Latin hypercube designs. *Statistica Sinica*, *18*(1), 171–186. http://www.jstor.org/stable/24308251

Jouzel J., Delaygue G., Landais A., Masson-Delmotte V., Risi C., & Vimeux F. (2013). Water isotopes as tools to document oceanic sources of precipitation. *Water Resources Research*, *49*(11), 7469–7486. https://doi.org/10.1002/2013WR013508

Kawamura K., Parrenin F., Lisiecki L., Uemura R., Vimeux F., Severinghaus J. P., Hutterli M. A., Nakazawa T., Aoki S., Jouzel J., Raymo M. E., Matsumoto K., Nakata H., Motoyama H., Fujita S., Goto-Azuma K., Fujii Y., & Watanabe O. (2007). Northern Hemisphere forcing of climatic cycles in Antarctica over the past 360,000 years. *Nature*, *448*(7156), 912–916. https://doi.org/10.1038/nature06015

Keller E. D., Baisden W. T., Bertler N. A., Emanuelsson B. D., Canessa S., & Phillips A. (2018). Calculating uncertainty for the RICE ice core continuous flow analysis water isotope record. *Atmospheric Measurement Techniques*, *11*(8), 4725–4736. https://doi.org/10.5194/amt-11-4725-2018

Kennedy M. C., & O'Hagan A. (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *63*(3), 425–464. https://doi.org/10.1111/1467-9868.00294

Khan N. S., Horton B. P., Engelhart S., Rovere A., Vacchi M., Ashe E. L., Törnqvist T. E., Dutton A., Hijma M. P., & Shennan I. (2019). Inception of a global atlas of sea levels since the Last Glacial Maximum. *Quaternary Science Reviews*, *220*, 359–371. https://doi.org/10.1016/j.quascirev.2019.07.016

Loeppky J. L., Sacks J., & Welch W. J. (2009). Choosing the sample size of a computer experiment: A practical guide. *Technometrics*, *51*(4), 366–376. https://doi.org/10.1198/TECH.2009.08040

Lowry D. P., Krapp M., Golledge N. R., & Alevropoulos-Borrill A. (2021). The influence of emissions scenarios on future Antarctic ice loss is unlikely to emerge this century. *Communications Earth & Environment*, *2*(1), 1–14. https://doi.org/10.1038/s43247-021-00289-2

Met-Office (2016). *Met Office climate prediction model: HadCM3*. http://www.metoffice.gov.uk/research/modelling-systems/unified-model/climate-models/hadcm3

Myung I. J. (2003). Tutorial on maximum likelihood estimation. *Journal of Mathematical Psychology*, *47*(1), 90–100. https://doi.org/10.1016/S0022-2496(02)00028-7

O'Hagan A. (2006). Bayesian analysis of computer code outputs: A tutorial. *Reliability Engineering & System Safety*, *91*(10–11), 1290–1300. https://doi.org/10.1016/j.ress.2005.11.025

O'Hagan A., Buck C. E., Daneshkhah A., Eiser J. R., Garthwaite P. H., Jenkinson D. J., Oakley J. E., & Rakow T. (2006). *Uncertain judgements: Eliciting experts' probabilities*. Wiley.

Peltier W. R. (2004). Global glacial isostasy and the surface of the ice-age earth: The ICE-5G (VM2) model and GRACE. *Annual Review of Earth and Planetary Sciences*, *32*(1), 111–149. https://doi.org/10.1146/earth.2004.32.issue-1

Petit J. R., Jouzel J., Raynaud D., Barkov N. I., Barnola J. M., Basile I., Bender M., Chappellaz J., Davis M., Delaygue G., Delmotte M., Kotlyakov V. M., Legrand M., Lipenkov V. Y., Lorius C., Pépin L., Ritz C., Saltzman E., & Stievenard M. (1999). Climate and atmospheric history of the past 420,000 years from the Vostok ice core, Antarctica. *Nature*, *399*(6735), 429–436. https://doi.org/10.1038/20859

Pollard D., & DeConto R. M. (2009). Modelling West Antarctic ice sheet growth and collapse through the past five million years. *Nature*, *458*(7236), 329–332. https://doi.org/10.1038/nature07809

Popp T. J. (2008). *The speed and timing of climate change: Detailed ice core stable isotope records from NorthGRIP, Greenland and Mt. Moulton, West Antarctica* [PhD thesis]. University of Colorado.

Rasmussen C. E., & Williams C. K. I. (2006). *Gaussian processes for machine learning*. MIT Press.

Simms A. R., Lisiecki L., Gebbie G., Whitehouse P. L., & Clark J. F. (2019). Balancing the Last Glacial Maximum (LGM) sea-level budget. *Quaternary Science Reviews*, *205*, 143–153. https://doi.org/10.1016/j.quascirev.2018.12.018

Steig E. J., Ding Q., White J. W. C., Küttel M., Rupper S. B., Neumann T. A., Neff P. D., Gallant A. J. E., Mayewski P. A., Taylor K. C., Hoffmann G., Dixon D. A., Schoenemann S. W., Markle B. R., Fudge T. J., Schneider D. P., Schauer A. J., Teel R. P., Vaughn B. H., … Korotkikh E. (2013). Recent climate and ice-sheet changes in West Antarctica compared with the past 2,000 years. *Nature Geoscience*, *6*(5), 372–375. https://doi.org/10.1038/ngeo1778

Steig E. J., Morse D. L., Waddington E. D., Stuiver M., Grootes P. M., Mayewski P. A., Twickler M. S., & Whitlow S. I. (2000). Wisconsinan and Holocene climate history from an ice core at Taylor Dome, western Ross Embayment, Antarctica. *Geografiska Annaler: Series A, Physical Geography*, *82*(2–3), 213–235. https://doi.org/10.1111/j.0435-3676.2000.00122.x

Stenni B., Buiron D., Frezzotti M., Albani S., Barbante C., Bard E., Barnola J. M., Baroni M., Baumgartner M., Bonazza M., Capron E., Castellano E., Chappellaz J., Delmonte B., Falourd S., Genoni L., Iacumin P., Jouzel J., Kipfstuhl S., … Udisti R. (2011). Expression of the bipolar see-saw in Antarctic climate records during the last deglaciation. *Nature Geoscience*, *4*(1), 46–49. https://doi.org/10.1038/ngeo1026

Sutter J., Eisen O., Werner M., Grosfeld K., Kleiner T., & Fischer H. (2020). Limited retreat of the Wilkes Basin ice sheet during the last interglacial. *Geophysical Research Letters*, *47*(13), e2020GL088131. https://doi.org/10.1029/2020GL088131

Tindall J. C., Valdes P. J., & Sime L. C. (2009). Stable water isotopes in HadCM3: Isotopic signature of El Niño Southern Oscillation and the tropical amount effect. *Journal of Geophysical Research: Atmospheres*, *114*(D4). https://doi.org/10.1029/2008JD010825

Valdes P. J., Armstrong E., Badger M. P. S., Bradshaw C. D., Bragg F., Crucifix M., Davies-Barnard T., Day J. J., Farnsworth A., Gordon C., Hopcroft P. O., Kennedy A. T., Lord N. S., Lunt D. J., Marzocchi A., Parry L. M., Pope V., Roberts W. H. G., Stone E. J., … Williams J. H. T. (2017). The BRIDGE HadCM3 family of climate models: HadCM3@Bristol v1.0. *Geoscientific Model Development*, *10*(10), 3715–3743. https://doi.org/10.5194/gmd-10-3715-2017

Van Loan C. F., & Golub G. (1996). *Matrix computations*. Johns Hopkins Studies in Mathematical Sciences. The Johns Hopkins University Press.

Vehtari A., Gelman A., & Gabry J. (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, *27*(5), 1413–1432. https://doi.org/10.1007/s11222-016-9696-4

Werner M., Jouzel J., Masson-Delmotte V., & Lohmann G. (2018). Reconciling glacial Antarctic water stable isotopes with ice sheet topography and the isotopic paleothermometer. *Nature Communications*, *9*(1), 1–10. https://doi.org/10.1038/s41467-018-05430-y

Whitehouse P. L. (2018). Glacial isostatic adjustment modelling: Historical perspectives, recent advances, and future directions. *Earth Surface Dynamics*, *6*(2), 401–429. https://doi.org/10.5194/esurf-6-401-2018

Whitehouse P. L., Bentley M. J., Milne G. A., King M. A., & Thomas I. D. (2012). A new glacial isostatic adjustment model for Antarctica: Calibrated and tested using observations of relative sea-level change and present-day uplift rates. *Geophysical Journal International*, *190*(3), 1464–1482. https://doi.org/10.1111/gji.2012.190.issue-3

Wilkinson R. D. (2010). Bayesian calibration of expensive multivariate computer experiments. In L. T. Biegler, G. Biros, O. Ghattas, M. Heinkenschloss, D. Keyes, B. K. Mallick, L. Tenorio, B. V. B. Waanders, K. Wilcox, & Y. Marzouk (Eds.), *Large-scale inverse problems and quantification of uncertainty* (pp. 195–215). John Wiley and Sons.

Yokoyama Y., & Purcell A. (2021). On the geophysical processes impacting palaeo-sea-level observations. *Geoscience Letters*, *8*(1), 1–19. http://doi.org/10.1186/s40562-021-00184-w