



Ice Cores and Emulation: Learning More About Past Ice Sheet Shapes

Fiona Elizabeth Turner

Submitted for the degree of Doctor of Philosophy
School of Mathematics and Statistics

August 2020

Supervisors:
Richard Wilkinson
Caitlin Buck
Louise Sime
Julie M. Jones

University of Sheffield

Abstract

Understanding the effect warming has on ice sheets is vital for accurate projections of climate change. A better understanding of how the Antarctic ice sheets have changed size and shape in the past would allow us to improve our predictions of how they may adapt in the future; this is of particular relevance in predicting future global sea level changes. This research makes use of previous reconstructions of the ice sheets, ice core data and Bayesian methods to create a model of the Antarctic ice sheet at the Last Glacial Maximum (LGM). We do this by finding the relationship between the ice sheet shape and water isotope values.

We developed a prior model which describes the variation between a set of ice sheet reconstructions at the LGM. A set of ice sheet shapes formed using this model was determined by a consultation with experts and run through the general circulation model HadCM3, providing us with paired data sets of ice sheet shapes and water isotope estimates. The relationship between ice sheet shape and water isotopes is explored using a Gaussian process emulator of HadCM3, building a statistical distribution describing the shape of the ice sheets given the isotope values outputted by the climate model. We then use MCMC to sample from the posterior distribution of the ice sheet shape and attempt to find a shape that creates isotopic values matching as closely as possible to the observations collected from ice cores. This allows us to quantify the uncertainty in the shape and incorporate expert beliefs about the Antarctic ice sheet during this time period. Our results suggests that the AIS may have been thicker at the LGM than previous reconstructions predict.

Acknowledgements

Firstly I'd like to thank my supervisors: Richard Wilkinson, Caitlin Buck, Louise Sime and Julie Jones. Somehow having four people supervising me was never anything less than a pleasure. Having such a group of people to turn to for their expertise and help made the last four years much easier than they could otherwise have been. Thank you for somehow managing to never give contradictory advice.

Thank you to Max, Irene, Kira and everyone else at the British Antarctic Survey who helped me as I came to grips with climatology, your patience and guidance were truly appreciated. My thanks to Louise Sime, Nick Golledge and Boer de Bas who provided me with the ice sheet reconstructions used in Chapter 3, and to the National Snow & Ice Data Center (<https://nsidc.org/>), where I accessed the ice core data used in Chapter 5.

This project would never have happened without the funding and assistance from the Grantham Centre for Sustainable Futures. Such interdisciplinary work as this often falls down the cracks of traditional funding institutes; thank you to everyone at the centre for your support throughout the years.

To Sonette and Robbie, for the copious amounts of wine we have shared these last four years as you listened to my ramblings. To all of my other friends, both in Sheffield and other places, for welcome distractions and vital encouragement. To my mum, dad, Mike, and Dan, for always believing in me and assuring me that I could be anything I want in life - sorry that that has turned out to be a perpetual student.

Abbreviations

AIS:	Antarctic Ice Sheets
AP:	Antarctic Peninsula
EAIS:	East Antarctic Ice Sheet
GCM:	General Circulation Model
GP:	Gaussian Process
HadCM3:	Hadley Centre Coupled Model - version 3
IPCC:	Intergovernmental Panel on Climate Change
Ka BP:	Thousands of Years Before Present
LGM:	Last Glacial Maximum
LOOCV:	Leave-One-Out Cross-Validation
PC(A):	Principal Component (Analysis)
PI:	Pre Industrial
SVD:	Singular Value Decomposition

Symbol List

$\mathbf{a}_i \in \mathbb{R}^{7008}$	principal components, often called loadings, found using null space method and used in our prior model, $i = 1, \dots, 5$
$\mathbf{y}_i \in \mathbb{R}^{47}$	vector of $\delta^{18}O$ anomalies at ice core sites from HadCM3 output, $i = 1, \dots, 10$
$\mathbf{z} \in \mathbb{R}^5$	the prior model variables
\mathbf{z}_{obs}	vector of \mathbf{z} values that would create the $\delta^{18}O$ observations from the ice cores
$\mathbf{z}'_{obs} = [\mathbf{1} \quad \mathbf{z}_{obs}]$	matrix used as part of the linear model component of the calibration model
$\beta \in \mathbb{R}^6$	linear model parameters for calibration model
θ	ice sheet shape output by the prior model
μ	row mean of the collection of ice sheet reconstructions
σ_1	HadCM3 model standard deviation
τ_1	isotope measurement standard deviation
τ_2	spatial standard deviation between HadCM3 and ice core sites
D	space that the $\delta^{18}O$ HadCM3 output describes
$E \in \mathbb{R}^{47}$	vector of elevation values at an ice core site in HadCM3 simulations
N_1	number of HadCM3 simulations
N_2	number of ice core sites being modelled
S	vector space in which we define ice sheet shape
$X \in \mathbb{R}^{7008 \times 40}$	Collection of ice sheet reconstructions
$Z \in \mathbb{R}^{47 \times 5}$	set of \mathbf{z}_i values used to make orographies included in HadCM3 simulations
$Z' = [\mathbf{1} \quad Z]$	matrix used as part of the linear model component of the calibration model

Contents

1	Introduction	1
1.1	Bayesian framework	5
2	Motivation and Literature Review	7
2.1	Data Issues in Antarctica	8
2.1.1	Past Ice Sheet Shapes	11
2.2	Reconstructing Climates with Proxy Data	13
2.2.1	The Use of Water Isotopes in Palaeo-climate Reconstruction	16
2.3	Bayesian methods in palaeo-climatology	18
2.4	Summary of chapter	20
3	Building The Prior Model	23
3.1	Principal Component Analysis	25
3.1.1	The first principal component	27
3.1.2	The second principal component	29
3.2	Reconstruction Errors	30
3.3	Performing PCA on Ice Sheet Reconstructions	31
3.3.1	The Newton-Raphson Method	47

3.3.2	The LU Decomposition	51
3.3.3	The nullspace method	53
3.4	Comparing two variable reduction approaches	58
3.5	Building the Prior Model	65
3.6	Summary of Chapter	66
4	Elicitation of Ice Sheet Shapes	69
4.1	Preparation	72
4.1.1	The Interactive Model	74
4.1.2	Planning the Elicitation Meeting	75
4.2	Conducting the Meeting	77
4.3	The Second Elicitation Process	80
4.4	Summary of Chapter	82
5	Building The Emulator	87
5.1	Running HadCM3	88
5.1.1	Issues with HadCM3	89
5.1.2	Processing Scripts	90
5.2	Exploratory Data Analysis	91
5.2.1	Comparing simulations of WAIS and EAIS	93
5.3	Gaussian Process Emulation	94
5.3.1	Choosing a kernel function	99
5.3.2	Building our emulator	99
5.4	Incorporating the EAIS sites	106
5.4.1	Using an Anisotropic Kernel	111
5.5	Summary of chapter	111

6	Calibration of the Model	115
6.1	Markov Chain Monte Carlo methods	116
6.1.1	Hamiltonian Monte Carlo (HMC)	116
6.2	1D model	117
6.3	Using all five prior variables	120
6.4	The multivariate model for the WAIS ice core sites	122
6.4.1	Comparing model output to linear regression	126
6.5	Incorporating the Gaussian process emulator	127
6.5.1	1-D example	128
6.5.2	1-D inverse example	132
6.5.3	Automatic Relevance Determination	135
6.6	Applying to the ice core data	135
6.6.1	Modelling combined and separate ice sheets	137
6.6.2	Checking hyperparameters	139
6.7	Analysing model output	143
6.7.1	Sensitivity analysis	143
6.8	Testing the model	150
6.8.1	Pseudo data	150
6.8.2	Recovering a HadCM3 simulation	152
6.8.3	Recovering an ice core site	155
6.9	Summary of Chapter	155
7	Geographical Interpretations	165
7.1	Posterior shapes	165
7.2	Comparing our orographies to previous reconstructions	167
7.3	Summary of Chapter	181

8	Conclusions	183
8.1	Discussion	184
8.1.1	Scientific Findings	184
8.1.2	Statistical Findings	185
8.2	Future Work	186
	Appendices	189
A	Design of eighty shapes from elicitation.	192
B	Post-processing scripts for HadCM3 output.	206
B.1	data_process_PIron.py	206
B.2	data_process.py	211
B.3	my_cmaps.py	217
B.4	my_defs.py	224
B.5	my_plots.py	230
B.6	process_annual.py	234
B.7	process_monthly_means.py	238
B.8	startup.py	240

Chapter 1

Introduction

Acceleration of ice flow and retreat in Antarctica, which has the potential to lead to sea level rise of several metres within a few centuries, ... may be the onset of an irreversible ice sheet instability. Uncertainty related to the onset of ice sheet instability arises from limited observations, inadequate model representation of ice sheet processes, and limited understanding of the complex interactions between the atmosphere, ocean and the ice sheet.

*IPCC Special Report on the Ocean and Cryosphere
in a Changing Climate, 2019*

A better understanding of the historic changes that have occurred in ice sheet shapes has become increasingly relevant as the consequences of man made climate change becomes more apparent. Ice sheets are defined as a mass of glacial ice larger than 50,000 km². Currently ice sheets are only present in Antarctica and Greenland; however, during the last glacial period (115,000 - 11,700 Ka BP), ice sheets existed in Europe and both Americas. A rise in global sea levels of 1m, primarily due to melting ice sheets, is considered likely by 2100. Uncertainty in the amount of sea rise is mostly determined by lack of knowledge about the possible contributions from ice sheets, particularly those in Antarctica (Pörtner et al., 2019). As ice mass loss across the continent has increased in recent years (IMBIE, 2018), the need to understand the contributions the Antarctic ice sheets (AIS) have made in the past to global sea levels is more apparent. A map of Antarctica can be seen in Figure 1.1, showing the location of the East Antarctic ice sheet (EAIS) and the West Antarctic ice sheet (WAIS), as well as other key geographical features.

Reconstructions of past ice sheet shapes have been produced by many authors in order to understand the state of the planet's climate at different time periods: Peltier (1994), Briggs et al. (2014) and Whitehouse et al. (2012) among others. We are interested in the shape of the Antarctic ice sheet at the Last Glacial Maximum (~ 21 Ka BP; LGM), the point of maximum ice sheet extent in the last glacial period (Clark et al., 2009). The LGM and the shape of ice sheets are both points of interest in palaeo-climatology; by understanding the deglaciation process in Antarctica, we can more fully understand the effect of external forcing on Antarctic and other ice sheets (Briggs et al., 2014). External forcings are a climate forcing agent outside of the climate system itself, such as solar or orbital variations.



Figure 1.1: A map of Antarctica showing the major geographic landmarks, from Abrahamsen (2012).

Further motivation for investigating the LGM is given in 2.1.1, including the issue of “missing ice” and its contribution to changes in sea level. However, it should be stated that quantifying our own ice sheet reconstructions in this thesis into ice sheet volume is beyond the scope of this work; due to the statistical nature of our models, we will be analysing the shape and size of the ice sheet, but are unable to give an estimate of how much ice our reconstructions hold. As our method models only surface elevation, we are unable to take into account the underlying bedrock, and thus cannot find an estimate for ice volume.

Building ice sheet reconstructions often relies on the use of climate models as simulators for the planet during the period of interest. However, simulators can be very expensive to run, and the number of simulations you can achieve is severely limited. The model we wish to use is HadCM3, an isotope-enabled coupled General Circulation Model (GCM), created by the Met Office in 1999 (Met-Office, 2016) and used frequently by the climatology community, including in the latest IPCC Assessment Report (Stocker et al., 2013). It should be stated that this work is not intended as a review or critique of HadCM3, or a test of how well it simulates the climate. Indeed, such work would be a thesis in itself. Here, we are not trying to fix HadCM3, or understand it, but rather use it as a key tool in our methods. We intend to improve understanding about the AIS by applying a Bayesian framework to ice sheet modelling. The computational cost of running HadCM3 means we must turn to a statistical method that has become an increasingly useful tool in palaeoclimatology, an emulator.

Emulators are cheaper, faster models that approximate the simulator. This allows us to learn about the relationship between the input and output of interest, using

a set of observations from our simulator and treating the emulator as a ‘black box’ (Kennedy and O’Hagan, 2001). We can then learn more about the process in question without being reliant on a slow and expensive simulator. Here we will be working with Gaussian process emulators, which models the outputs as a Gaussian process on the space determined by the inputs.

1.1 Bayesian framework

A Bayesian approach to statistical inference describes all uncertainties with probability distributions, and tells us how to update our initial uncertainty (our prior distribution) in light of data and a model (giving us our posterior distribution). Mathematically given a data point, x , some parameter θ and a statistical model $p(x|\theta)$ s.t.

$$x \sim p(x|\theta),$$

if we have prior distribution for θ

$$\theta \sim p(\theta)$$

then given a sample

$$X = [x_1, \dots, x_n]$$

the posterior distribution $p(\theta|X)$ can be calculated

$$p(\theta|X) \propto p(X|\theta)p(\theta).$$

$p(\theta)$ is known as the prior distribution, representing our beliefs about the parameter of interest, θ . $p(X|\theta)$ is the likelihood function, representing how likely X is under the model if θ is the parameter, and $p(\theta|X)$ is the posterior distribution, an updated distribution for θ given the information found from X and the model.

Applying this to a palaeo-climatological context, we are interested in the shape of the AIS at the LGM, which we will denote as θ . Previous ice sheet reconstructions are used to form a prior distribution of the ice sheet shape, $p(\theta)$. The observations we have are in the form of water isotope ratios from ice cores. A description of how these are calculated can be found in Section 2.2.1. Isotope data are an example of proxy data: preserved physical characteristics of past climate that can stand in for direct measurement. The arguments for the use of such proxy data are strong: by analysing preserved analogues, past climates can be used to understand future climate change and assess climate models' validity (Salzmann et al., 2009). An emulator of HadCM3 will be used to find the relationship between ice sheet shape and our proxy data, allowing us to calculate the distribution $p(X|\theta)$. We can then combine these to find the posterior distribution, and from that predict the distribution of the "true" ice sheet shape, $\hat{\theta}$.

In the next chapter, we give a review of the literature as well as motivation for our work. By considering the issues with modelling Antarctica, past work on reconstructing climates with proxy data, and current Bayesian methods used in palaeoclimatology, we show that there is a real need for further incorporation of statistical thinking in ice sheet modelling and climatology in general.

Chapter 2

Motivation and Literature Review

Modelling the Earth's climate is reliant on knowledge of both the internal and external forcings, accurate initial and boundary conditions, and can be aided by a long record of observations. Understanding past changes in climate can reduce uncertainties around predictions of future changes. In climate modelling, initial conditions are starting values for variables such as wind, temperature and moisture, whilst boundary conditions are values decided on by a modeller for variables such as solar radiation intensity. Weather is largely dependent on initial conditions, whilst climate, a longer time scale, is dependent on these boundary conditions.

Focus is often placed on the higher latitudes as changes in polar ice caps have an effect on global sea levels and ocean circulation (Bracegirdle and Stephenson, 2012). A knowledge of what is happening at the poles is therefore necessary in order to

understand what may occur both regionally and globally. An important aspect of this is the amount of ice present in the ice sheets. In order to analyse how ice sheets have changed over millennia, accurate reconstructions of past shapes are required. However, knowledge of the change of ice sheet margins is limited, which inhibits our ability to reconstruct the extent of the ice sheets (Haywood et al., 2016).

When analysing the climate in Antarctica, specific issues arise, from the short time-scale of the data to the scarcity of measurements over vast regions of the continent, and the particularly variable nature of the climate itself (Fogt et al. (2016), Haywood et al. (2016), Schneider et al. (2006)). The use of more advanced statistical techniques, using the data we do have in new ways, could offer a way to work around these issues and result in a more detailed and accurate climate model.

In this chapter, we look at the existing literature to motivate this project. The importance of the Antarctic Ice Sheets (AIS) and the relevance of the Last Glacial Maximum (LGM) is discussed, and the use of Bayesian methodology in the palaeoclimatology community is reviewed.

2.1 Data Issues in Antarctica

Here we consider the issues in relation to studying Antarctica's climate in various forms, as well as ice sheet mass change.

A major issue when it comes to modelling the Antarctic climate is the short time-scale of data. Measurements were not collected until the 1950s when stations were set up on the continent; data that spanned the entire landmass were not available until satellites were launched in 1979. At least 30 years of data are required to understand a climate, although a longer record is required if the climate is going through a period of rapid change (Notz, 2015). Having such a short record in Antarctica makes it incredibly difficult to detect changes in the region and understand what effect climate change is having on the continent. With barely forty years of climate data, it is difficult to characterise the natural variability of Antarctica, or to significantly recognise an anthropogenic contribution (Jones et al., 2016). With such little information available, climatologists cannot draw accurate conclusions on any potential changes in the region; this lack of observations means the Antarctic climate is the least understood on Earth (Fogt et al., 2016).

The lack of observations in the region is especially problematic as even more data is needed for Antarctica compared to other regions due its large natural climatic variability. Trends often cannot be definitively attributed to any external forcing, as changes in the climate can still be shown to be within the expected range of internal variability. Natural variability is also difficult to estimate with such a short record. This means that until recently it was not possible to detect any effect that climate change may be having on the region. The analysis conducted by Abram et al. (2013) of ice melt on the Antarctic Peninsula concluded that although there was warming across the continent throughout the last century, it was “*not unprecedented in the context of natural climate variability*” [pg. 409]. Similarly, Turner et al. (2016) found that changes in the continents climate could not be associated with global climate change, but rather reflect the extreme natural internal variability of the region. More

recently, Clem et al. (2020) concluded that the affects of anthropogenic warming over Antarctica has been obscured by decadal variability throughout this century.

Antarctica, at 14 million km², is twice as large as Australia. Studies at research stations are conducted at single locations and therefore cannot represent the conditions of the whole continent. This is apparent when analysing a specific time period. Regions will display leads and lags in a trend; the period, for instance, of maximum ice coverage or minimum temperature will not be the same for the whole of the land mass (Haywood et al., 2016). When observed as a whole therefore, any warming effects often average out as some regions experience changes that others do not. There is clear evidence of warming in the AP; Schneider et al. (2006) showed that some of the strongest warming on the planet during the last half century occurred there. However, they concluded that it was difficult to give any context to this change due to the short instrumental records. Turner et al. (2016) also found an absence of significant warming in the last two decades due to the extreme natural variability, although the more recent work of Screen et al. (2018) did state that human influence was detectable in the region. It is clear that more data must be incorporated somehow into models in order to learn more about the changes in climate and to assess whether they are outside of natural variability.

Our interest lies in the changes in ice sheet mass. A motivation for looking at past time periods is to better understand changes in the present day ice sheets. IMBIE (2018) have found a cumulative mass change in the ice sheets between 1992 and 2017, although they recognise these estimates become more uncertain as the area of each ice sheet region increases. Although WAIS and AP definitely decreased in

size during this period by $-159 \pm 26 \text{ Gt yr}^{-1}$ and $-20 \pm 15 \text{ Gt yr}^{-1}$ respectively, their study concluded that the mass change of the EAIS was $5 \pm 46 \text{ Gt yr}^{-1}$. This large uncertainty shows how difficult it is to accurately model the AIS. Studies of recent mass change such as this one are conducted through the use of satellites and short term data; when past ice sheet shapes are measured, we are reliant on other methods.

There are numerous issues with the data collected so far in Antarctica. The size of the continent as well as the large variability in its climate means we require more data than other parts of the globe in order to accurately model it. The short record of data available means this is difficult to do.

2.1.1 Past Ice Sheet Shapes

Extensive work has already been done reconstructing past ice sheet shapes or other climate variables with the use of climate models (Holloway et al., 2016). 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 ice sheets at two different time periods. 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.

As discussed in Malmierca-Vallet et al. (2018), understanding how ice sheets respond to changes in climate is fundamental to predicting more accurately how high latitudes will react to increases in greenhouse gases. These regions often act as amplifiers, with the poles experiencing some of the clearest evidence of global warming. Extensive modelling of past ice sheet shapes is critical to improve predictions of global sea level rise, as the change in ice sheet volume is directly linked to sea level changes.

The LGM is of particular interest due to the uncertainty around the ice sheet shapes during this period, causing an issue of “missing ice”. Simms et al. (2019) studied the LGM sea-level budget, as current estimates of sea-level rise after the LGM are not balanced by the amount of ice believed to have melted since then. They found that this amounts to a shortfall of $15.6 \pm 9.6\text{m}$ of global sea-level rise. It would seem that the contributions of one or more ice sheets has been underestimated; as the extent of the AIS is relatively unconstrained during this time period, this is historically where the missing ice is believed to have been. Reducing uncertainty around the WAIS and EAIS at this time period is therefore of great relevance; if the AIS has lost more ice mass than scientists have previously predicted up to the present day, this may mean it is more responsive to changes in the climate.

As described in Whitehouse et al. (2012), several LGM reconstructions have been published, all with a large range in amplitude and distribution of ice mass. This variation between reconstructions is due to the range of methods and data used to constrain the models, and suggests that the errors in the models are large. An overview of the methods used to build the collected published reconstructions we use in our research can be found in Chapter 3. This thesis develops a Bayesian

approach that models this uncertainty, highlighting where the reconstructions agree and disagree with one another, give us an opportunity to reduce this uncertainty, and use the prior knowledge offered by existing reconstructions to better constrain estimates of the AIS.

Previous reconstructions of the AIS can vary greatly. Given the importance in understanding the ice sheets at the LGM in the context of present day and future sea level rise, it is important that the uncertainties within and between existing ice sheet reconstructions are better understood, and a better understanding of the AIS at the LGM is obtained.

2.2 Reconstructing Climates with Proxy Data

As discussed in Section 2.1.1, past ice sheet reconstructions are built with a combination of proxy data and climate models, and they play an important role in our own methodology. We are reliant on the use of proxy data in order to model different parts of Antarctica's climate before observational data is available, both for the previous centuries and much longer time periods.

There have been many calls for the incorporation of proxy data into climate models over Antarctica; this would help solve the issue of a lack of data as well as providing information about past climates when the environment was similar to current or near-future conditions. Jones et al. (2016) believed that the use of such data would help to

progress understanding of the Antarctic climate, as well as improve proxy modelling such as with water isotopes. The use of water isotopes in palaeo-climatology research is discussed in more detail in Section 2.2.1. The use of samples such as ice cores, moss banks and tree rings in modelling is becoming more common, having been used by Abram et al. (2013), Thomas et al. (2009) and Schneider et al. (2006) amongst others in order to model past Antarctic climate systems.

The use of proxy data sources allows us to learn more about past conditions, such as surface air temperature. Examples of this for the recent past includes Abram et al. (2013), who used deuterium isotopes (δD , one of the two stable isotopes of hydrogen; isotopes are defined more fully in Section 2.2.1) as a proxy for mean annual temperature to reconstruct snow melt during the twentieth century, and concluded that warming is now present over the West Antarctic Ice Sheet. Schneider et al. (2006) also used δD , as well as oxygen-18 ($\delta^{18}O$) ice core records, to reconstruct the mean surface temperature over Antarctica. They support the use of proxy data, stating that because the physical processes underpinning their variability is well understood, there is a strong advantage for using isotopic records. By normalising the records to cover 1800-1983, combining the ice cores by a weighted average and calibrating the reconstruction they found proof of long-term warming over the continent.

Whether warming is present solely in small regions has also been studied with the use of proxy data. Thomas et al. (2009) used a 150 year isotope record as a proxy for temperature and found that the warming reported by Schneider et al. (2006) was not a local phenomenon, but was part of a warming trend over the whole region going back to 1900, and that these trends are stronger than the range of natural

variability. Thomas et al. (2013) went further and created a 308-year record 1702-2010 of climate variability in West Antarctica using δD from an ice core drilled on the West Antarctic coast. Although they found statistically significant warming in the area since 1957, they also discovered similar warming and cooling trends throughout the record. They found that warming had not yet left the natural range of climate variability for the region.

The papers above have focused on reconstructing the timelines of past centuries using proxy data. Cuffey et al. (2016), Holloway et al. (2016) and D'Andrilli et al. (2017) all used proxy data to create records of Antarctic climate over much longer time periods, going back to the LGM. Cuffey et al. (2016) reconstructed the temperature history over West Antarctica from the LGM to present day using water isotopes and temperature readings in boreholes. They commented that understanding the transition from glacial to interglacial is important, as it allows us to learn how forcings such as a rise in atmospheric CO_2 effects the climate. D'Andrilli et al. (2017) created a record of organic matter from a West Antarctic ice core dating back to the LGM. This allowed them to improve their understanding of the influence that organisms have on climate change. Although focusing on different aspects of climate, they all show the importance of such data to provide information on the nature of Antarctica through time. Without the use of these data sources, we would know little about how the AIS responds to glacial and interglacial periods.

These papers all showcase the benefits of incorporating proxy records in to the evaluation of the Antarctic climate. We can extend climate records and compare recent warming with past trends, yet these papers show that having more data is not all

that is needed; the way the data is used is also important. Although proxy data and models have been brought together through data assimilation, that is not the point of this thesis. We instead wish to use a novel Bayesian approach that will combine proxy data and climate models to better model the ice sheets.

2.2.1 The Use of Water Isotopes in Palaeo-climate Reconstruction

As mentioned in 1.1, we will be making use of proxy data in the form of stable water isotopes found in ice cores. Water isotopes have long been viewed as an essential tool for palaeo-climate research. As stated in Haywood et al. (2016), they have been used to estimate palaeotemperatures and consequently ice volume, and have become the “*multitool of palaeoceanographic inquest*” [pg. 4].

Jouzel et al. (2013) explains in detail the way water isotopes are used in climatology and what we can learn from them. Water is mostly formed of H_2^{16}O , with tiny amounts being isotopic molecules H_2^{18}O , H_2^{17}O and HD^{16}O . The distribution of these molecules varies both spatially and temporally in the atmosphere depending on climatological parameters. Most notably there is a “*linear relationship between annual values of δD and $\delta^{18}\text{O}$ and the annual mean temperature ... at middle and high latitudes*” [pg. 7469], where

$$\delta = (R_{\text{sample}} - R_{\text{SMOW}})/R_{\text{SMOW}}$$

with R_{sample} and R_{SMOW} the isotopic ratios of the sample and the Vienna Standard Mean Ocean Water. δD and $\delta^{18}O$ are also linearly related to each other along the Meteoric Water Line (MWL, the relationship between hydrogen and oxygen isotope ratio in water). We use the paper Sime et al. (2009) to define the “*deuterium excess parameter*”, or *d*-excess, as $d = \delta D - 7.85 \times \delta^{18}O$, which is sensitive to differences between data and the MWL.

This formula has been used to reconstruct past climates from water archives. In Antarctica and Greenland, δD and $\delta^{18}O$ have been measured to calculate temperature change at the drilling site (Jouzel et al., 2013). Comparison with *d*-excess has then allowed a reconstruction of conditions at the ocean surface. More recently ^{17}O -excess has been developed, providing new information about oceanic sources due to differences between $\delta^{17}O$ and $\delta^{18}O$ (Landais et al., 2008).

As shown here, water isotopes can play a key part in ice sheet reconstruction. Through their use, we can look back further than observational data would allow, and estimate how the ice sheets have transitioned from glacial to interglacial periods. However, there must be careful consideration to ensure the data are used as effectively as possible. Further details of the ice cores we use in this work are given in Section 5.2.

2.3 Bayesian methods in palaeo-climatology

Work has already been done on learning about past ice sheet shapes through statistical analysis. The work of Chang et al. (2016b), Chang et al. (2016a), Chang et al. (2014a) and Chang et al. (2014b) are good examples of the techniques used. In Chang et al. (2016b), the most recent paper, they generate projections of the change in ice volume in the West Antarctic Ice Sheet (WAIS) using both modern- and palaeo-observations. Here we use the word observations in the statistical sense, not geographical. Throughout this thesis, ice core data will be referred to as observations despite them not adhering to the geographic definition in order to differentiate it from any synthetic data either from HadCM3 or simulated ourselves.

Chang et al. (2016b) encourage the use of a combination of proxy and observational data; the use of solely modern observations may not lead to well-constrained projections, as they lack information on long-term parameters. Data that only covers the last 60 years or so will not accurately estimate the parameters in climate models that represent any long-term changes or trends; palaeo-data is clearly needed for this purpose. As they go on to say, “*studies using heuristic approaches suggest that utilising information from palaeo-data can reduce uncertainties in these long-term behaviour related parameters*” [pg. 3]. Through a use of emulation and calibration, their results are less uncertain than those using solely modern data.

Their method consists of two steps: emulation and calibration. They use the PSU3D-ICE model, run with 625 different settings for the four parameters of interest, with

each model simulation starting at 40Ka BP and spun up to 5,000 years in to the future. Due to computational and inferential challenges, the dimensions of model output are reduced through PCA. A Gaussian process emulator is constructed for each principal component to be used, giving an approximation to this model output. Information from these emulators is then combined with observations to infer input parameters for the model. A simulated example shows that this method can recover true parameters with little uncertainty, and provides information on parameter interactions. Projections are also constrained better when using real data, with the probability of unrealistic events occurring reduced.

There are some caveats to this method which the authors recognise. The atmospheric conditions used are very simple, assuming a linear increase in temperature over 150 years before remaining constant. More detailed scenarios need to be considered for warming, and the grid size used is also very coarse. They conclude that further work should focus on reducing the size of the latitude and longitude cells, although this will make the emulation runs more expensive.

The work in Domingo et al. (2020) gives even more motivation for our work. Using a similar approach to us, they used PCA to reduce a set of reconstructions to components describing the Greenland Ice Sheet during the Last Interglacial. A set of orographies were submitted as part of a set of simulations in HadCM3, and emulation and history matching were used to identify orographies that match proxy data observations. Their results show that this method of data-model comparison has great potential, although it also shows the importance of well-constrained ice core records.

As computational power has increased, statistical methods have become more applicable for palaeo-climatology. Starting with Parnell (2005) and Kopp et al. (2009), Gaussian processes have become a useful tool when analysing climates in past time periods. A detailed description of GP emulators is given in Chapter 5. By focusing on the LGM, a more recent time period, we build on the work of Domingo et al. (2020) by demonstrating this data-model comparison method when more ice core records are available. The breadth of time that the LGM took place in means the dating errors in these ice cores are relatively low and hopefully will result in a reduced uncertainty around the AIS shape.

2.4 Summary of chapter

Understanding the past size and shape of the Antarctic ice sheets is imperative to understand how they will respond to future warm climates. The LGM is an important time period to study due to the uncertainty around global ice sheets and the “missing ice” problem. By reducing uncertainty around the AIS during this period and giving more certain estimations of where this missing ice may have been, we can better constrain the contribution Antarctica has made to rising sea levels up to present day.

The literature shows a clear need to incorporate more proxy data into climate modelling. It also shows that the application of statistical techniques already widely used in other areas could allow us to learn more from the data than current methods are capable of. With the increasingly recognised need for the incorporation of

more advanced statistical techniques in palaeo-climatology, interdisciplinary research is slowly becoming the norm. We aim to show that a thorough and informed way of combining Bayesian methodology, climate modelling and proxy data is possible.

In the next chapter, we use previous reconstructions of the AIS at the LGM to build a statistical model describing the variations between these reconstructions. This will model our prior beliefs of the AIS, showing which parts of the ice sheet reconstruction are most uncertain, and will be the first stage of our Bayesian framework.

Chapter 3

Building The Prior Model

We wish to create a prior model of the shape of the Antarctic ice sheets at the LGM. This will be a statistical model that outputs synthetic ice sheet shapes, with no physical laws underlying it. A sample of these shapes will be chosen through consultation with experts and run through HadCM3 to provide us with estimated water isotope ratio values.

To ensure our prior model is as informative as possible, we incorporate expert knowledge about the ice sheets at the LGM. We conducted a literature review to collect reconstructions of the ice sheets. This led to a set of forty ice sheet shapes from Peltier (2004), Argus et al. (2014), Whitehouse et al. (2012), Briggs et al. (2014), Pollard and DeConto (2009), Golledge et al. (2012), Golledge et al. (2013), Bentley

et al. (2014) and de Boer et al. (2017), which will be used to create a prior distribution over the space of possible shapes. We can use these shapes to describe a vector space of possible shapes. The sum of two vectors in the space of shapes, and the multiplication of a vector by a scalar, remains in this vector space.

We want to keep the number of variables in our prior distribution to as few as possible; using the collected reconstructions in their original form is unwise as a model with forty variables would be too computationally expensive. We also want to have variables that are orthogonal to each other, so that they describe different directions of the vector space. As many of these ice sheet reconstructions are made using the same data and models, they are not uncorrelated and must be transformed. We wish to do this whilst retaining as much information, or variance, as possible.

From a statistical viewpoint, our method is as follows: we wish to find a shape $y \in \mathbb{R}^{96 \times 73}$, where

$$y_{ij} = \text{height at latitude}_i, \text{ longitude}_j.$$

$\mathbb{R}^{96 \times 73}$, which is the dimensions of the HadCM3 latitude/longitude grid, is a vector space. The ice sheet shapes live within a subspace S . We use the reconstructions from the literature to define

$$S = \text{span}\{X_1, \dots, X_{40}\}.$$

where X_i is a matrix describing the ice sheet reconstruction. $\dim(S) = 40$, meaning we would need forty coordinates to describe a shape in S . This is too many

dimensions, so we aim to find a compact way to describe shapes in S , i.e. find \hat{S} s.t.

$$\hat{S} = \text{span}\{\mathbf{a}_1, \dots, \mathbf{a}_k\} \quad \text{for some } k < 40, \quad \text{such that } \hat{S} \simeq S.$$

To find these basis vectors \mathbf{a}_i , we use principal component analysis (Jolliffe, 1986).

There are some drawbacks to the method. By defining the space our prior model will lie in by the collection of ice sheet reconstructions, we are discarding any ice sheet shapes that lie outside of this space. We are therefore reliant on these reconstructions being plausible estimates of the ice sheets at the LGM, and anything they do not describe being implausible.

In this chapter, we use a variation on principal component analysis to build a prior model of the Antarctic Ice Sheets at the LGM. We discuss the various methods considered, and the checks used to decide on the number of variables to be included in our model.

3.1 Principal Component Analysis

Principal component analysis (PCA) provides us with a statistical transformation of multivariate data. Given a data set with n observations of p variables it finds a new set of n uncorrelated observations in the new variables, the principal components. These describe the variation in the data ordered by how much of this variation they

describe; the first principal component describes the direction of the largest amount of variation, the second principal component describes the direction of the second largest amount and so on. The first $k < p$ principal components are then chosen to represent the data, allowing us to reduce the number of dimensions whilst discarding a minimal amount of variation.

Here, our data is in the form of forty 96×73 matrices, which is the dimensions of the HadCM3 latitude/longitude grid; each grid cell represents an area of $3.75^\circ \times 2.5^\circ$. We reshape the ice sheet shapes into a single matrix of dimension 7008×40 and treat the 7008 latitude/longitude grid cells as “observations” in each of the forty ice sheet reconstructions, which are our “variables”. We therefore have $n = 7008$ and $p = 40$.

In algebraic terms the process is: given our data set

$$X = [\mathbf{X}_1, \dots, \mathbf{X}_{40}] \in \mathbb{R}^{7008 \times 40}$$

with row mean $\boldsymbol{\mu} \in \mathbb{R}^{7008}$, which we remove before performing the analysis to centre the data, empirical covariance matrix for all forty shapes

$$\Sigma_{40} = \hat{X} \hat{X}^T \in \mathbb{R}^{7008 \times 7008},$$

where

$$\hat{X} = X - \left(\frac{1}{n} X \mathbf{1}\right) \mathbf{1}^T$$

is the centred matrix, and $\mathbf{X}_i = (x_{i1}, \dots, x_{i7008})$, we want to find a set of vectors $\mathbf{a}_k \in \mathbb{R}^{7008}$ such that

$$Y = A^T X, \quad A = (\mathbf{a}_1, \dots, \mathbf{a}_k) \tag{1}$$

where $Y \in \mathbb{R}^{k \times 40}$ is a matrix of independent, uncorrelated vectors describing X .

Geometrically, $Y = \mathbf{a}_k^T X$ is the projection of X on to a new basis described by \mathbf{a}_k . Values of Y give the coordinates of each observation along the set of vectors \mathbf{a}_k , where \mathbf{a}_k are the principal components.

3.1.1 The first principal component

The first principal component is defined to be the vector $\mathbf{a}_1 \in \mathbb{R}^{7008}$ where the projection $\mathbf{a}_1^T X$, for some single random vector X , has maximal variance subject to the normalising constraint $\mathbf{a}_1^T \mathbf{a}_1 = 1$. We therefore wish to maximise

$$\text{Var}(\mathbf{a}_1^T X) = \mathbf{a}_1^T \Sigma_{40} \mathbf{a}_1. \quad (2)$$

We use a Lagrange multiplier to find the maximum value of this, taking into account the normalising constraint. A Lagrange multiplier is a method of finding the local minima and maxima of a function, $f(x)$, for some variable x and given a set of constraints, $g(x) = 0$. Assuming the function and constraints all have continuous first derivatives, a Lagrange multiplier λ is introduced giving a Lagrange function:

$$L(x) = f(x) - \lambda g(x).$$

A stationary point of the Lagrange function is then found through differentiation. The function we wish to maximise when performing PCA is

$$L(\mathbf{a}_1) = \mathbf{a}_1^T \Sigma_{40} \mathbf{a}_1 - \lambda(\mathbf{a}_1^T \mathbf{a}_1 - 1).$$

We need to maximise L with respect to \mathbf{a}_1 and λ . Differentiating with respect to λ and setting the result to zero gives

$$\mathbf{a}_1^T \mathbf{a}_1 - 1 = 0$$

which is the normalising constraint. Differentiating with respect to \mathbf{a}_1 gives

$$\frac{\partial L}{\partial \mathbf{a}_1} = 2\Sigma_{40}\mathbf{a}_1 - 2\lambda\mathbf{a}_1.$$

Setting this to 0, we find

$$(\Sigma_{40} - \lambda I)\mathbf{a}_1 = 0 \Rightarrow \Sigma_{40}\mathbf{a}_1 = \lambda\mathbf{a}_1.$$

λ is therefore an eigenvalue of Σ_{40} . Σ_{40} is a 7008×7008 covariance matrix and has 39 eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_{39}$, having lost a degree of freedom due to centring the data, all of which are non-negative as Σ_{40} is positive semidefinite. Assuming the eigenvalues are distinct, we have $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_{39} \geq 0$.

Returning to Equation 2, we have

$$\begin{aligned}\text{Var}(\mathbf{a}_1^T X) &= \mathbf{a}_1^T \Sigma_{40} \mathbf{a}_1 \\ &= \mathbf{a}_1^T \lambda I \mathbf{a}_1 \\ &= \lambda.\end{aligned}$$

As we want to maximise $\text{Var}(\mathbf{a}_1^T X)$, we choose λ to be λ_1 , the largest eigenvalue of Σ_{40} . Therefore, \mathbf{a}_1 is the eigenvector corresponding to the largest eigenvalue.

3.1.2 The second principal component

The second principal component, \mathbf{a}_2 , is also found using a Lagrange multiplier with the additional constraint that $X^T \mathbf{a}_2$ be uncorrelated with $X^T \mathbf{a}_1$. This means setting

$$\mathbf{a}_2^T \Sigma_{40} \mathbf{a}_1 = 0$$

However, since $\Sigma_{40} \mathbf{a}_1 = \lambda_1 \mathbf{a}_1$ we can instead use the constraint $\mathbf{a}_2^T \mathbf{a}_1 = 0$. This means that \mathbf{a}_2 and \mathbf{a}_1 should be orthogonal.

This requires two Lagrange multipliers, giving us

$$L(\mathbf{a}_2) = \mathbf{a}_2^T \Sigma_{40} \mathbf{a}_2 - \lambda(\mathbf{a}_2^T \mathbf{a}_2 - 1) - \alpha \mathbf{a}_2^T \mathbf{a}_1.$$

Differentiating with respect to the Lagrange multipliers and \mathbf{a}_2 then gives

$$\frac{\partial L}{\partial \lambda} = 0 \Rightarrow \mathbf{a}_2^T \mathbf{a}_2 = 1,$$

$$\frac{\partial L}{\partial \alpha} = 0 \Rightarrow \mathbf{a}_2^T \mathbf{a}_1 = 0$$

and

$$\frac{\partial L}{\partial \mathbf{a}_2} = 0 \Rightarrow 2\Sigma_{40}\mathbf{a}_2 - 2\lambda\mathbf{a}_2 - \alpha\mathbf{a}_1 = 0. \quad (3)$$

Premultiplying by \mathbf{a}_1^T , this becomes

$$2\mathbf{a}_1^T \Sigma_{40} \mathbf{a}_2 - 2\lambda \mathbf{a}_1^T \mathbf{a}_2 - \alpha \mathbf{a}_1^T \mathbf{a}_1 = 0. \quad (4)$$

$\mathbf{a}_1^T \Sigma_{40} \mathbf{a}_2 = \mathbf{a}_2^T \Sigma_{40} \mathbf{a}_1$, since Σ_{40} is a real symmetric matrix. This means

$$\mathbf{a}_1^T \Sigma_{40} \mathbf{a}_2 = \mathbf{a}_2^T \Sigma_{40} \mathbf{a}_1 = \mathbf{a}_2^T \lambda \mathbf{a}_1 = \lambda \mathbf{a}_2^T \mathbf{a}_1 = 0,$$

as \mathbf{a}_1 is an eigenvector of Σ_{40} . Thus the first component of Equation 3 is zero. The second component also equals zero due to the constraint that \mathbf{a}_1 and \mathbf{a}_2 be orthogonal to each other. As $\mathbf{a}_1^T \mathbf{a}_1 = 1$ we find that $\alpha = 0$. Equation 3 therefore becomes $(\Sigma_{40} - \lambda I)\mathbf{a}_2 = 0$, so \mathbf{a}_2 is the eigenvector with the second largest eigenvalue, λ_2 .

Continuing this method, we find that the j^{th} principal component is the eigenvector corresponding to the j^{th} largest eigenvalue.

3.2 Reconstruction Errors

Once we have decided on a set of principal components, it is important that they accurately represent the original data. We check this by reconstructing the data

using the new projection Y found with Equation 1. This is done by calculating

$$AY + \mu = \hat{X},$$

where A is the matrix of k eigenvectors \mathbf{a}_k and \hat{X} is the reconstruction of X . To check the accuracy of $\hat{X} = [\hat{X}_1, \dots, \hat{X}_n]$, the root mean square error is found with

$$R = \sqrt{\frac{1}{n} \sum_{i=1}^n (\hat{X}_i - X_i)^2}, \quad n = 40.$$

Using all of the principal components would give a value of 0, as no information has been discarded. If some of the components have been discarded, we try to keep R as small as possible subject to k not being too large.

3.3 Performing PCA on Ice Sheet Reconstructions

We try two methods to create a prior model using the collected ice sheet reconstructions. For Method 1, we perform PCA on all forty shapes together and form the prior model out of subsets of these principal components. We build a linear model using the first four, five and six of the principal components, discarding 13.98%, 9.70% and 3.98% of the variation respectively. A screeplot of the forty principal components is in Figure 3.1; we can see that after the first six, the amount of variance each component represents becomes negligible. The other thirty-four principal components are not used; we treat them as if they represent noise in the data. We are limited in how


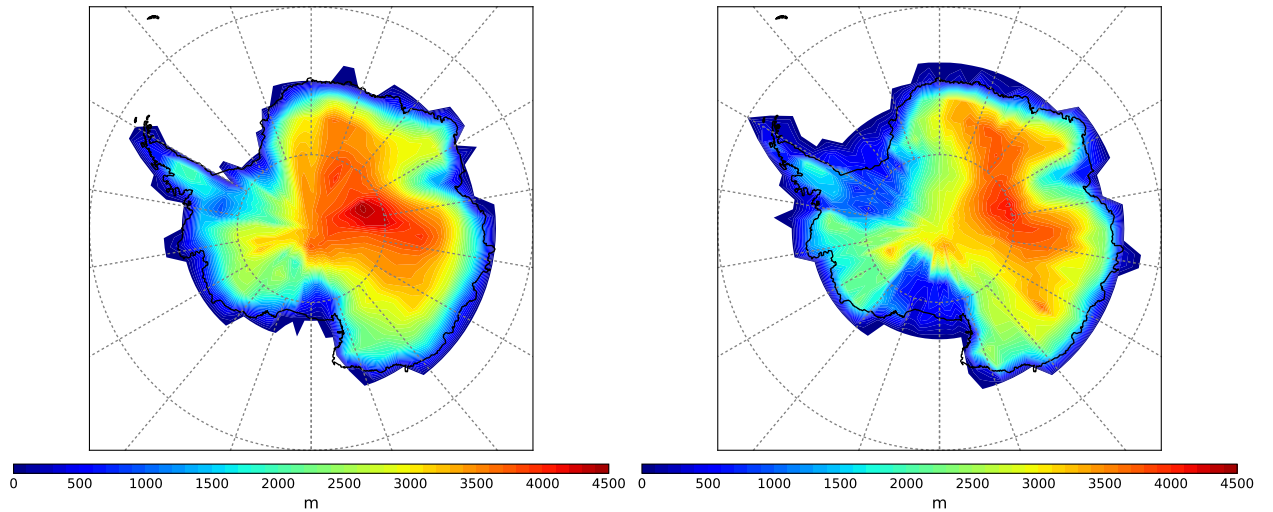


Figure 3.1: A screeplot of the eigenvalues of the principal components from method 1, showing the decrease in the amount of variance each component represents. After the sixth, the amount of variance each subsequent component represents is negligible.

many principal components we can use by the number of simulations we will be able to run in HadCM3; using the rule of needing to run ten simulations for each variable in the model if using a GP emulator (Loeppky et al., 2009), we cannot have more than six in our prior model.

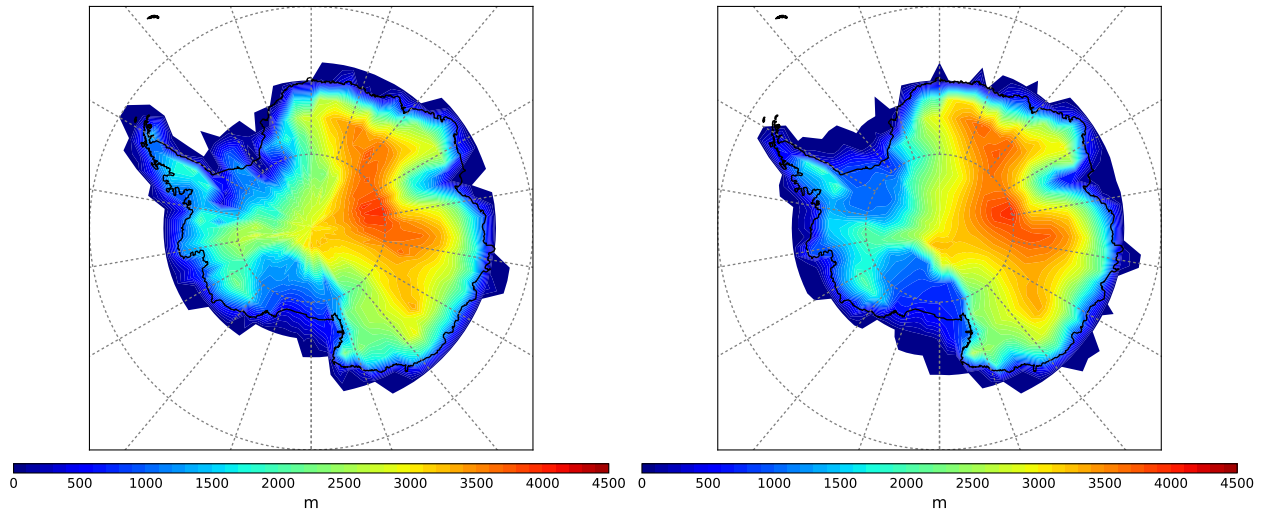
However, not all reconstructions are equally widely used by the climatology community. The first four ice sheet shapes we collected are data-based reconstructions and widely favoured by the palaeo-climatology community when running climate models.

These are ICE5G (Peltier, 2004), ICE6G (Argus et al., 2014), W12A (Whitehouse et al., 2012) and GLAC-1D (Briggs et al., 2014); plots of them can be seen in Figure 3.2 and an overview of the methods used to build them can be seen in Table 3.1. The other thirty six shapes come from ice sheet model output. These are Pollard and DeConto (2009), Golledge et al. (2012), Golledge et al. (2013), Bentley et al. (2014) and de Boer et al. (2017); they are plotted in Figure 3.3 and summarised in Table 3.2. When these come from the same model, particularly in de Boer et al. (2017), the resulting shapes will be strongly related to one another. Given this we want to explore approaches that ensure they do not carry the same weight as the other shapes. PCA describes variation in the data, but we want to prioritise the variation described by ICE5G, ICE6G, W12A and GLAC-1D. If we use all forty together, then they are treated equally. We calculated the RMSE with reconstructions of the forty collected ice sheet shapes; the values can be seen in Tables 3.3, 3.4 and 3.5. Whilst this method gives low reconstruction errors, some of the largest errors are for the four shapes that we wish to have most influence on the prior model. We therefore want to find a method of constructing basis vectors that would emphasise these shapes compared to the other thirty-six ice sheet reconstructions.



(a) Ice sheet reconstruction ICE5G from Peltier (2004).

(b) Ice sheet reconstruction ICE6G from Argus et al. (2014).




(c) Ice sheet reconstruction W12A from Whitehouse et al. (2012).

(d) Ice sheet reconstruction GLAC-1D from Briggs et al. (2014).


Figure 3.2: The first four ice sheet reconstructions from Peltier (2004), Argus et al. (2014), Whitehouse et al. (2012) and Briggs et al. (2014).

	Methodology	Resolution	Domain	References
ICE-5G	Refinement of ICE-4G. Coral histories in Barbados used to construct relative sea level history, then tuned with other data from the west Pacific. Ice sheet thickness inferred from glacial isostasy with mathematical models.	$3.75^\circ \times 2.5^\circ$	Global	Peltier (2004)
ICE-6G	Adjusted ICE-5G with new GPS observations and ice core records. Ice loss removed from East Antarctica due to lack of constraints on plateau.	$1^\circ \times 1^\circ$	Global	Argus et al. (2014)
W12A	Combined proxy data with a mathematical ice sheet model. Boundary conditions found using proxy data or models. Regions with little or no data estimated by relating to areas with many observations.	$0.5^\circ \times 0.5^\circ$	Antarctica (below 60°)	Whitehouse et al. (2012)
GLAC-1D	Glacial systems model (GSM) used to build a reconstruction. The GSM was run, providing a data ensemble, which was then constrained with observational data and filtered to rule out unlikely results.	$1^\circ \times 0.5^\circ$	Global	Briggs et al. (2014)


Table 3.1: A summary of four of the ice sheet reconstructions to be used in this project.




(a) Ice sheet reconstruction from Pollard and DeConto (2009).



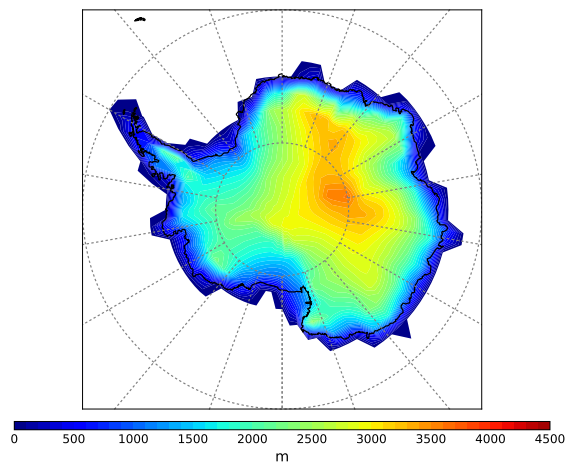
(b) Ice sheet reconstruction from Golledge et al. (2012).




(c) Ice sheet reconstruction from Golledge et al. (2013).




(d) Ice sheet reconstruction from Bentley et al. (2014).




(e) Ice sheet reconstruction from de Boer et al. (2017).




(f) Ice sheet reconstruction from de Boer et al. (2017).




(g) Ice sheet reconstruction from de Boer et al. (2017).




(h) Ice sheet reconstruction from de Boer et al. (2017).




(i) Ice sheet reconstruction from de Boer et al. (2017).




(j) Ice sheet reconstruction from de Boer et al. (2017).




(k) Ice sheet reconstruction from de Boer et al. (2017).




(l) Ice sheet reconstruction from de Boer et al. (2017).




(m) Ice sheet reconstruction from de Boer et al. (2017).




(n) Ice sheet reconstruction from de Boer et al. (2017).




(o) Ice sheet reconstruction from de Boer et al. (2017).




(p) Ice sheet reconstruction from de Boer et al. (2017).




(q) Ice sheet reconstruction from de Boer et al. (2017).




(r) Ice sheet reconstruction from de Boer et al. (2017).




(s) Ice sheet reconstruction from de Boer et al. (2017).




(t) Ice sheet reconstruction from de Boer et al. (2017).




(u) Ice sheet reconstruction from de Boer et al. (2017).




(v) Ice sheet reconstruction from de Boer et al. (2017).




(w) Ice sheet reconstruction from de Boer et al. (2017).




(x) Ice sheet reconstruction from de Boer et al. (2017).




(y) Ice sheet reconstruction from de Boer et al. (2017).




(z) Ice sheet reconstruction from de Boer et al. (2017).




(aa) Ice sheet reconstruction from de Boer et al. (2017).




(ab) Ice sheet reconstruction from de Boer et al. (2017).




(ac) Ice sheet reconstruction from de Boer et al. (2017).




(ad) Ice sheet reconstruction from de Boer et al. (2017).




(ae) Ice sheet reconstruction from de Boer et al. (2017).




(af) Ice sheet reconstruction from de Boer et al. (2017).




(ag) Ice sheet reconstruction from de Boer et al. (2017).



(ah) Ice sheet reconstruction from de Boer et al. (2017).



(ai) Ice sheet reconstruction from de Boer et al. (2017).



(aj) Ice sheet reconstruction from de Boer et al. (2017).

Figure 3.3: The thirty-six ice sheet reconstructions used to create the second set of principal components. These are from Pollard and DeConto (2009), Golledge et al. (2012), Golledge et al. (2013), Bentley et al. (2014) and de Boer et al. (2017).

Methodology	Resolution	Domain	References
Combined ice sheet/ice shelf model with new treatment of grounding-line dynamics and ice-shelf buttressing	40km \times 40km	Antarctica (below 55°)	Pollard and DeConto (2009)
Numerical ice-sheet model, using terrestrial and marine geological data for constraints	50km \times 50km	Antarctica (below 56°)	Golledge et al. (2012)
Combine interpretations of ice sheet model results with geological data	100km \times 100km	Antarctica (below 55°)	Golledge et al. (2013)
Synthesis of geological and glaciological datasets to determine position of grounding-line and ice sheet upper surface	150km \times 150km	Antarctica (below 55°)	Bentley et al. (2014)
Climate model snapshots forced with different orbital forcing scenarios	400km \times 400km	Antarctica (below 54°)	de Boer et al. (2017)

Table 3.2: A summary of the thirty-six “other” ice sheet reconstructions to be used in this project.

Therefore, for Method 2 we split our data into two subsets. The first subset contains the four ice sheet shapes that we wish to be most influential on our prior model. The second set contains the other thirty-six shapes, whose information we will incorporate once we have found variables from the first set. PCA is applied to the first subset of ice sheet shapes to keep the reconstruction errors as small as possible. This gives four principal components that represent 50.45%, 34.48%, 9.23% and 5.84% of the variation in those data alone. The last of these is discarded, as it only represents a small proportion of the data. The first three principal components, which we call $A = [\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3]$, collectively contain 94.16% of the variation in these four shapes and will be used to build the prior model.

The three principal components from this first set of ice sheet shapes are shown in Figure 3.4. The first principal component represents the contrasts between patches in the WAIS and the rest of the continent. The second component compares the centre of the continent with the coast and the third contrasts patches throughout Antarctica.

We now attempt to make variables from the second set of 36 ice sheet reconstructions. The variation of these shapes needs to be added into the prior model in the form of at least one vector that is orthogonal to the principal components in A . We need to take the original three principal components into account when creating further basis vectors as we do not want to represent the same variation in the data twice in our prior model. In Method 2, we try to find a way of performing PCA whilst taking into account the set A .

In mathematical terms, we have three orthogonal vectors \mathbf{a}_1 , \mathbf{a}_2 and $\mathbf{a}_3 \in \mathbb{R}^{7008}$, the

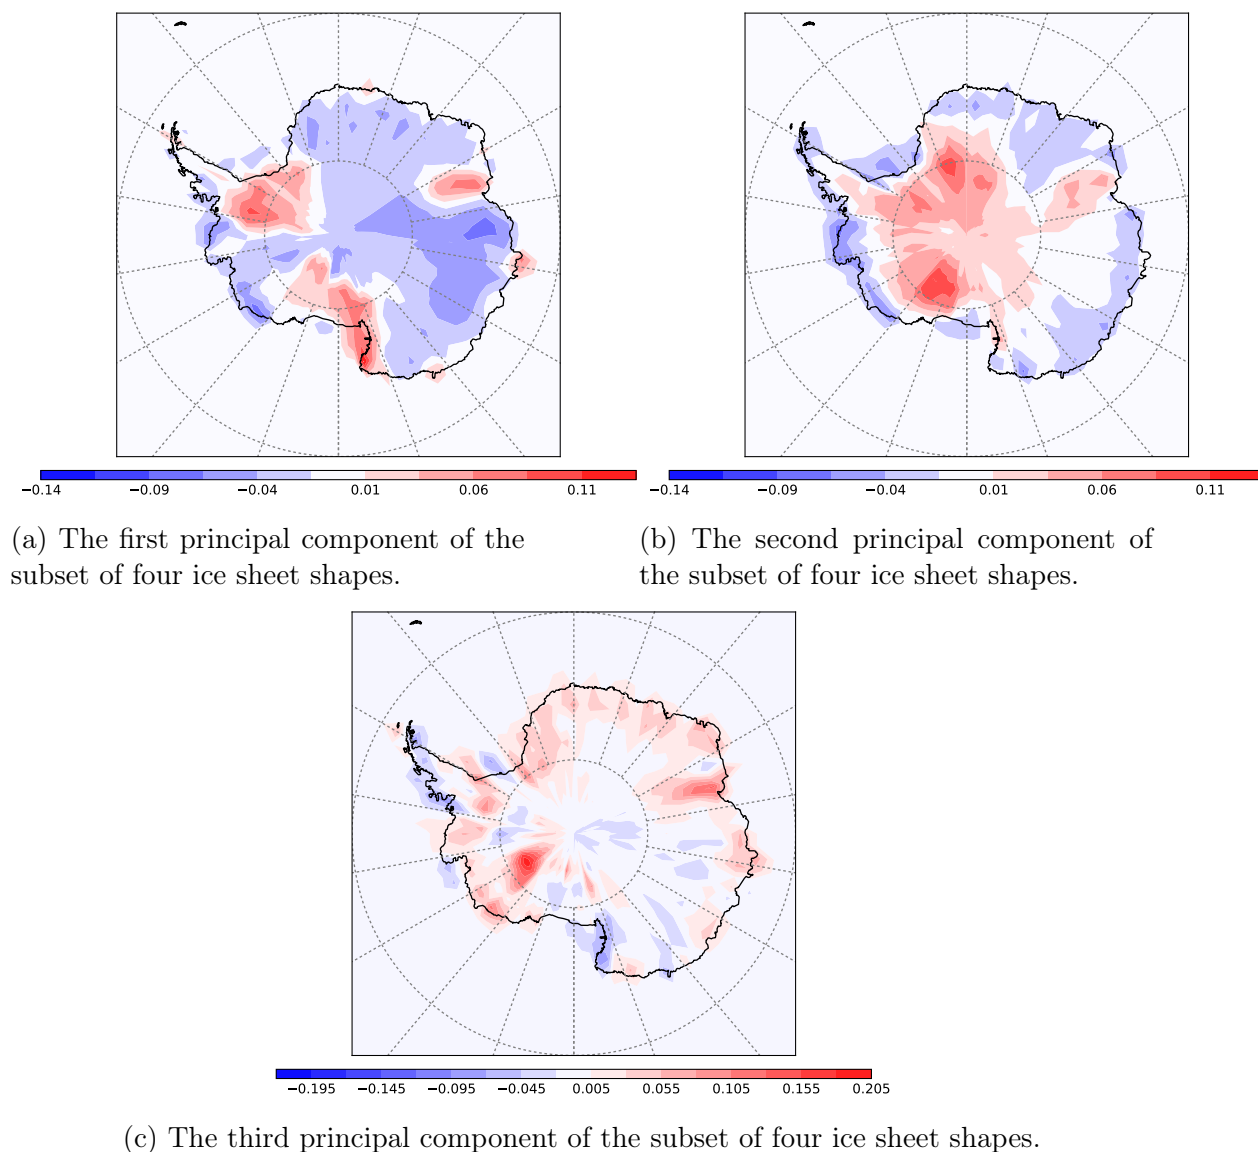


Figure 3.4: The first three principal components for the subset of four ice sheet shapes as defined by Peltier (2004), Argus et al. (2014), Whitehouse et al. (2012) and Briggs et al. (2014). These components are the first three basis vectors for our prior model.

principal components from the first four ice sheet reconstructions, and a data set $X_{36} \in \mathbb{R}^{7008 \times 36}$, containing the other thirty-six ice sheet reconstructions. We wish to find a vector, \mathbf{a}_4 , that describes the variation in X_{36} such that

$$\mathbf{a}_4 \perp \mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3$$

and maximises the variance of the projection $\mathbf{a}_4^T X_{36}$. We therefore want to choose \mathbf{a}_4 that maximises $\text{Var}(\mathbf{a}_4^T X_{36}) = \mathbf{a}_4^T \Sigma_{36} \mathbf{a}_4$, where $\Sigma_{36} \in \mathbb{R}^{7008 \times 7008}$ is the empirical covariance matrix of X_{36} , with the constraint that $\mathbf{a}_4^T \mathbf{a}_1 = \mathbf{a}_4^T \mathbf{a}_2 = \mathbf{a}_4^T \mathbf{a}_3 = 0$. We add in a normalising constraint, $\mathbf{a}_4^T \mathbf{a}_4 = 1$ and try the Lagrange function again to find \mathbf{a}_4 . We now wish to maximise the function

$$L(\mathbf{a}_4) = \mathbf{a}_4^T \Sigma_{36} \mathbf{a}_4 - \lambda(\mathbf{a}_4^T \mathbf{a}_4 - 1) - \alpha \mathbf{a}_4^T \mathbf{a}_1 - \beta \mathbf{a}_4^T \mathbf{a}_2 - \gamma \mathbf{a}_4^T \mathbf{a}_3.$$

Differentiating with respect to the Lagrange multipliers λ , α , β and γ gives the constraints. Differentiating with respect to \mathbf{a}_4 gives

$$\frac{\partial L}{\partial \mathbf{a}_4} = 2\Sigma_{36} \mathbf{a}_4 - 2\lambda \mathbf{a}_4 - \alpha \mathbf{a}_1 - \beta \mathbf{a}_2 - \gamma \mathbf{a}_3.$$

We wish to find the value of \mathbf{a}_4 for which this is equal to zero. We also want to find the values of the Lagrange multipliers. By premultiplying by \mathbf{a}_1^T , we find

$$2\mathbf{a}_1^T \Sigma_{36} \mathbf{a}_4 - 2\lambda \mathbf{a}_1^T \mathbf{a}_4 - \alpha \mathbf{a}_1^T \mathbf{a}_1 - \beta \mathbf{a}_1^T \mathbf{a}_2 - \gamma \mathbf{a}_1^T \mathbf{a}_3 = 0.$$

Unlike when performing PCA, the first component is not equal to zero as \mathbf{a}_1 is not an eigenvector of Σ_{36} . We therefore have $\alpha = 2\mathbf{a}_1^T \Sigma_{36} \mathbf{a}_4$. By similarly premultiplying by \mathbf{a}_2^T , \mathbf{a}_3^T and \mathbf{a}_4^T we find that $\beta = 2\mathbf{a}_2^T \Sigma_{36} \mathbf{a}_4$, $\gamma = 2\mathbf{a}_3^T \Sigma_{36} \mathbf{a}_4$ and $\lambda = 2\mathbf{a}_4^T \Sigma_{36} \mathbf{a}_4$.

This gives us the equation

$$\Sigma_{36}\mathbf{a}_4 - \mathbf{a}_4^T \Sigma_{36} \mathbf{a}_4 \mathbf{a}_4 - \mathbf{a}_1^T \Sigma_{36} \mathbf{a}_4 \mathbf{a}_1 - \mathbf{a}_2^T \Sigma_{36} \mathbf{a}_4 \mathbf{a}_2 - \mathbf{a}_3^T \Sigma_{36} \mathbf{a}_4 \mathbf{a}_3 = 0, \quad (5)$$

which we must solve for \mathbf{a}_4 .

3.3.1 The Newton-Raphson Method

We now need to find the value of \mathbf{a}_4 for which Equation 5 holds. We first try this using the Newton-Raphson method, which finds successively better approximations to the roots of a vector-valued function, $\mathbf{x} : \mathbf{f}(\mathbf{x}) = 0$ for some variable \mathbf{x} . The method starts with the function \mathbf{f} , its derivative $\nabla \mathbf{f} \in \mathbb{R}^{p \times p}$ if $\dim \mathbf{f}(\mathbf{x}) = p$, and an initial estimate \mathbf{x}_0 . The better approximation, \mathbf{x}_1 , is found with

$$\mathbf{x}_1 = \mathbf{x}_0 - (\nabla \mathbf{f}(\mathbf{x}_0))^{-1} \mathbf{f}(\mathbf{x}_0).$$

This process is repeated as

$$\mathbf{x}_{n+1} = \mathbf{x}_n - (\nabla \mathbf{f}(\mathbf{x}_n))^{-1} \mathbf{f}(\mathbf{x}_n)$$

until a root has been found.

A 3-D toy example

To show how the Newton-Raphson method works, we use a 3-D toy example. We have a variance matrix

$$\Sigma_1 = \begin{pmatrix} 3 & 0 & 0 \\ 0 & 0.1 & 0 \\ 0 & 0 & 0.1 \end{pmatrix}$$

whose first eigenvector is

$$\mathbf{a}_1^T = (1 \ 0 \ 0).$$

We wish to find a new vector, \mathbf{a}_2 from our second variance matrix

$$\Sigma_2 = \begin{pmatrix} 2 & 1 & 0 \\ 1 & 2 & 0 \\ 0 & 0 & 0.1 \end{pmatrix}$$

that is orthogonal to \mathbf{a}_1 and maximises $\mathbf{a}_2^T \Sigma_2 \mathbf{a}_2$. The first eigenvalue of Σ_2 is $(1, 1, 0)^T$, which is not orthogonal to \mathbf{a}_1 . We therefore use Newton-Raphson to find the vector that solves

$$\Sigma_2 \mathbf{a}_2 - \mathbf{a}_2^T \Sigma_2 \mathbf{a}_2 \mathbf{a}_2 - \mathbf{a}_1^T \Sigma_2 \mathbf{a}_2 \mathbf{a}_1 = 0.$$

We find the derivative of this:

$$2[\Sigma_2 - 2(\Sigma_2 \mathbf{a}_2) \mathbf{a}_2^T - (\mathbf{a}_2^T \Sigma_2 \mathbf{a}_2) I - (\Sigma_2 \mathbf{a}_1) \mathbf{a}_1^T]$$

and run the function

$$\mathbf{a}_n = \mathbf{a}_{n-1} - (\nabla f(\mathbf{a}_{n-1}))^{-1} f(\mathbf{a}_{n-1})$$

a hundred times. This gives a new vector

$$\mathbf{a}_2^T = \begin{pmatrix} 0 & 1 & 0 \end{pmatrix},$$

which is the y -axis. This is the direction we would expect \mathbf{a}_2 to take, as \mathbf{a}_1 is the x -axis and Σ_2 has nearly all variation in the x and y directions.

Applying to our own data

We now try this method on our own data. In order to find the derivative of Equation 5, we use the summation convention (Einstein, 1916). When an index variable appears twice in a single term and is a free variable, it implies summation of that term over all the values of the index. This means that

$$y = \sum_{i=1}^n c_i x_i = c_1 x_1 + c_2 x_2 + \dots + c_n x_n$$

can be simplified to

$$y = c_i x_i.$$

We differentiate each component of $\frac{\partial L}{\partial \mathbf{a}_4} = f(\mathbf{a}_4)$ individually with respect to \mathbf{a}_4 . The derivative of the first component of Equation 5 is

$$\frac{\partial(\Sigma_{36,ij}\mathbf{a}_{4j})}{\partial \mathbf{a}_{4m}} = \Sigma_{36,ij}\delta_{jm} = \Sigma_{36,im},$$

where

$$\delta_{jm} = \begin{cases} 1 & \text{if } j = m, \\ 0 & \text{if } j \neq m. \end{cases}$$

The derivative of the second component is

$$\begin{aligned} \frac{\partial(\mathbf{a}_{4k}\Sigma_{36,kj}\mathbf{a}_{4j}\mathbf{a}_{4i})}{\partial \mathbf{a}_{4m}} &= \delta_{km}\Sigma_{36,kj}\mathbf{a}_{4j}\mathbf{a}_{4i} + \mathbf{a}_{4k}\Sigma_{36,kj}\delta_{mj}\mathbf{a}_{4i} + \mathbf{a}_{4k}\Sigma_{36,kj}\mathbf{a}_{4j}\delta_{im} \\ &= \Sigma_{36,mj}\mathbf{a}_{4j}\mathbf{a}_{4i} + \mathbf{a}_{4k}\Sigma_{36,km}\mathbf{a}_{4i} + \mathbf{a}_{4k}\Sigma_{36,kj}\mathbf{a}_{4j}\delta_{im}. \end{aligned}$$

The derivative of the third component is

$$\begin{aligned} \frac{\partial(\mathbf{a}_{1k}\Sigma_{36,kj}\mathbf{a}_{4j}\mathbf{a}_{1i})}{\partial \mathbf{a}_{4m}} &= \mathbf{a}_{1k}\Sigma_{36,kj}\delta_{jm}\mathbf{a}_{1i} \\ &= \mathbf{a}_{1k}\Sigma_{36,km}\mathbf{a}_{1i}. \end{aligned}$$

We get similar results for the last two components, $\beta\mathbf{a}_2$ and $\gamma\mathbf{a}_3$. Putting these components together we get:

$$\begin{aligned} \nabla f(\mathbf{a}_4) &= 2\Sigma_{36} - 2[(\mathbf{a}_4^T\Sigma_{36}\mathbf{a}_4)I - (\Sigma_{36}\mathbf{a}_4)\mathbf{a}_4^T + (\mathbf{a}_4^T\Sigma_{36})^T\mathbf{a}_4^T] - 2(\mathbf{a}_1\Sigma_{36})^T\mathbf{a}_1^T \\ &\quad - 2(\mathbf{a}_2\Sigma_{36})^T\mathbf{a}_2^T - 2(\mathbf{a}_3\Sigma_{36})^T\mathbf{a}_3^T \\ &= 2[\Sigma_{36} - 2(\Sigma_{36}\mathbf{a}_4)\mathbf{a}_4^T - (\mathbf{a}_4^T\Sigma_{36}\mathbf{a}_4)I] - (\Sigma_{36}\mathbf{a}_1)\mathbf{a}_1^T - (\Sigma_{36}\mathbf{a}_2)\mathbf{a}_2^T - (\Sigma_{36}\mathbf{a}_3)\mathbf{a}_3^T. \end{aligned} \tag{6}$$

However, when we try to use the Newton-Raphson method on Equations 5 and 6, it proves too computationally expensive to find the inverse of $\nabla f(\mathbf{a}_4) \in \mathbb{R}^{7008 \times 7008}$ due to the size of our data set. We therefore try another method to find $\nabla f(\mathbf{a}_4)^{-1}$, so that we can run the Newton-Raphson method.

3.3.2 The LU Decomposition

The LU decomposition factors a matrix into a lower triangular matrix, L , and an upper triangular matrix, U . It is then easy to find the inverse of the matrix. Set $Z = \nabla f(\mathbf{a}_4)$, the matrix to be decomposed, and $z = (\nabla f(\mathbf{a}_4))^{-1}$, the inverse we wish to calculate. Performing the LU decomposition gives

$$LU = Z,$$

so

$$Zz = LUz = L(Uz) = I,$$

where I is the $n \times n$ identity matrix. We have n^2 equations; to solve these, we use forward and back substitution. Starting with

$$Ly = I,$$

where $y = Ux$, we can use forward substitution to find

$$y_1 = \frac{1}{l_{11}},$$

$$y_i = \frac{1}{l_{ii}} \left[1 - \sum_{j=1}^{i-1} l_{ij} y_j \right]$$

for $i = 2, \dots, n$. We can then solve

$$Uz = y$$

for x using back substitution. This gives

$$z_n = \frac{y_n}{u_{nn}},$$

$$z_i = \frac{1}{u_{ii}} \left[y_i - \sum_{j=i+1}^n u_{ij} z_j \right]$$

for $i = n-1, \dots, 1$. We can now try the Newton-Raphson method, replacing $(\nabla f(\mathbf{a}_4))^{-1}$ with \mathbf{z} . However, trying this proves again to be too computationally expensive due to the size of our dataset.

A 3-D toy example

Applying the LU decomposition to the 3-D example, we find the upper and lower triangular matrices of $\nabla f(\mathbf{a}_0)$, where $\mathbf{a}_0^T = (1, 0, 0)$, are

$$L = \begin{pmatrix} 1 & \frac{1}{3} & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad U = \begin{pmatrix} -12 & 0 & 0 \\ 2 & \frac{-2}{3} & 0 \\ 0 & 0 & -3.8 \end{pmatrix}.$$

We can then use the R commands “forwardsolve” and “backsolve” to find the inverse.

This gives

$$\begin{pmatrix} 0 & 0.5 & 0 \\ -0.25 & -1.5 & 0 \\ 0 & 0 & -0.2632 \end{pmatrix}.$$

We can check this is accurate as finding the inverse in this example is computationally very cheap. Using the “solve” command, we find

$$(\nabla f(\mathbf{a}_0))^{-1} = \begin{pmatrix} 0 & 0.5 & 0 \\ -0.25 & -1.5 & 0 \\ 0 & 0 & -0.2632 \end{pmatrix}.$$

Using the newly found x instead of $(\nabla f(\mathbf{a}_0))^{-1}$, we run the Newton-Raphson function. Again, we find the new vector is $\mathbf{a}_2^T = (0, 1, 0)$.

3.3.3 The nullspace method

The Newton-Raphson method has proved too computationally expensive because of the size of our dataset. We therefore need to find a different way of finding basis vectors other than solving Equation 5. The optimisation problem that we wish to solve is a quadratically constrained quadratic program (QCQP), an optimisation problem where both the objective function and the constraints are quadratic functions. These have the form

$$\text{minimize } f_0(x)$$

for some variable x , subject to $f_i(x) \leq 0$ for all $i = 1, \dots, m$

where $f_i : \mathbb{R}^n \rightarrow \mathbb{R}$ has the form

$$f_i(x) = x^T P_i x + q_i^T x + r_i.$$

This is simple to solve if P_0 is positive semi-definite (Basu et al., 2017). However, our function is defined as

$$\text{minimize} \quad -\mathbf{a}_4^T \Sigma_{36} \mathbf{a}_4$$

subject to

$$\mathbf{a}_4^T \mathbf{a}_4 = 1, \quad \mathbf{a}_i^T \mathbf{a}_4 = 0, \quad i = 1, 2, 3.$$

Our problem is non-convex as P_0 in this case is negative semi-definite, and it is therefore not trivial to solve. An added difficulty is the size of our data set, which means we want to avoid doing too much numerical work. However, we can solve the problem by considering the vector space that A describes.

The condition $A^T \mathbf{a}_4 = 0$ is equivalent to saying $\mathbf{a}_4 \in \text{null}(A^T) \equiv B$, where B is the nullspace of A . The nullspace is the set of all vectors \mathbf{v} which satisfies $A^T \mathbf{v} = 0$, forming a subset of \mathbb{R}^n . We can write any inner product space as the direct sum of the nullspace of a linear operator and its compliment. This means that we can write $\mathbb{R}^n = A \oplus A^\perp = A \oplus B$. Thus $\mathbf{a}_4 \in B$ as $\mathbf{a}_4 = \sum \mathbf{w}_i \mathbf{b}_i = B \mathbf{w}$ where $B = [\mathbf{b}_1, \dots, \mathbf{b}_k]$ for some vector $\mathbf{w} \in \mathbb{R}^k$. The problem then becomes:

$$\text{maximise} \quad \mathbf{w}^T B^T \Sigma_{36} B \mathbf{w}$$

$$\text{subject to} \quad \mathbf{w}^T B^T B \mathbf{w} = \mathbf{w}^T \mathbf{w} = 1,$$

as

$$B^T B = \mathbb{I},$$

since the \mathbf{b}_i are orthonormal. As $\Sigma_{36} = X_{36} X_{36}^T$ and $B^T \Sigma_{36} B = B^T X_{36} X_{36}^T B = (X_{36}^T B) X_{36}^T B$, this is now our original eigenvalue problem from equation 2. We can solve this by setting \mathbf{w} to the eigenvectors of $X_{36}^T B$, which we can find with the singular value decomposition (SVD). This is the factorisation of a matrix. For a $n \times p$ matrix M , the SVD factorises it into the form $U \Sigma V^T$ where U is an orthonormal $n \times n$ matrix, Σ is a diagonal $n \times n$ matrix containing the eigenvalues of M , and V is a $p \times n$ matrix. The right singular vectors, the columns of V , are eigenvectors of M .

We perform the singular value decomposition on $X_{36}^T B$ and take the first three vectors in V as \mathbf{w} . We then multiply them by B to find the new basis vectors. These are orthogonal to A , the matrix of the first three basis vectors, and are orthonormal. We now have a set of six vectors: three principal components that describe the four main ice sheet reconstructions and three basis vectors that describe the other thirty-six shapes. However, incorporating more basis vectors into our model adds in more dimensions for us to work in, and we only have a limited number of simulations we can run through HadCM3 to explore this model space. We therefore must consider how many of the new basis vectors to incorporate in to our prior model.

The three new basis vectors are plotted in Figure 3.5. The first vector describes variation between the Eastern plateau and the rest of the continent. The second vector represents variation between patches in the East and West and the rest of

Antarctica. The third vector shows contrasts between the coast and a patch in the West, with emphasis over an area next to the peninsula.

A 3-D toy example

We apply the nullspace method to our 3D example. We simulate two data sets of multivariate normally distributed variables: X_1 , with distribution $\mathcal{N}(0, \Sigma_1)$, and X_2 with distribution $\mathcal{N}(0, \Sigma_2)$. In R code, our exact steps are

```
Z1 = matrix(c(3, 0, 0, 0, 0.1, 0, 0, 0, 0.1), ncol=3)
```

```
data1 = mvtnorm::rmvnorm(10000, rep(0, 3), Z1)
```

```
Z2 = matrix(c(2, 1, 0, 1, 2, 0, 0, 0, 0.1), ncol=3)
```

```
data2 = mvtnorm::rmvnorm(10000, rep(0, 3), Z2)
```

We know our first eigenvector of Σ_1 is

$$\mathbf{a}_1 = \begin{pmatrix} 1 & 0 & 0 \end{pmatrix}.$$

We find its nullspace using the R command `Null`

$$\mathbf{b} = \begin{pmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{pmatrix},$$

and multiply this with X_2 . We use the SVD to find eigenvectors of $X_2^T \Sigma_2$, and multiply this with the set of nullspace vectors \mathbf{b} .

```
a1 = c(1,0,0)
```

```
b1 = MASS::Null(a1)
```

```
Y = svd(data2%*%b1)$v
```

```
a2 = b1%*%Y
```

Our new eigenvector is

$$\mathbf{a}_2 = \begin{pmatrix} 0 & 1 & 0 \end{pmatrix}.$$

We then reconstruct our data and find the root mean square errors; these are all very close to zero. For reference, the code to do the reconstruction and find the errors is as follows:

```
A = cbind(a1,a2)
```



```
data = rbind(data1,data2)
```

```
datarecon = data%*%A%*%t(A)
```

```
rmse <- sqrt(rowMeans((data-datarecon)^2)).
```

This is a problem that we have not found elsewhere, and believe we have created a novel solution. Our approach is new, and may be of use in other areas.

3.4 Comparing two variable reduction approaches

We wish to compare the two methods we have used to find basis vectors and see which gives the most accurate reconstruction of the forty shapes, looking in particular at the errors for the first four shapes. We also wish to see how much the accuracy of the methods increase when we use four, five or six basis vectors in our model, as we want our prior model to have as few variables as possible without discarding too much information from the data. We are limited in how many shapes from our model we can run through the climate model HadCM3, so we require our basis to have as few dimensions as possible or risk not representing the entire basis with our set of simulations.

We compare the methods by calculating the Root Mean Square Error (RMSE) of the reconstructed shapes. We calculate the difference between ice sheet shapes and the

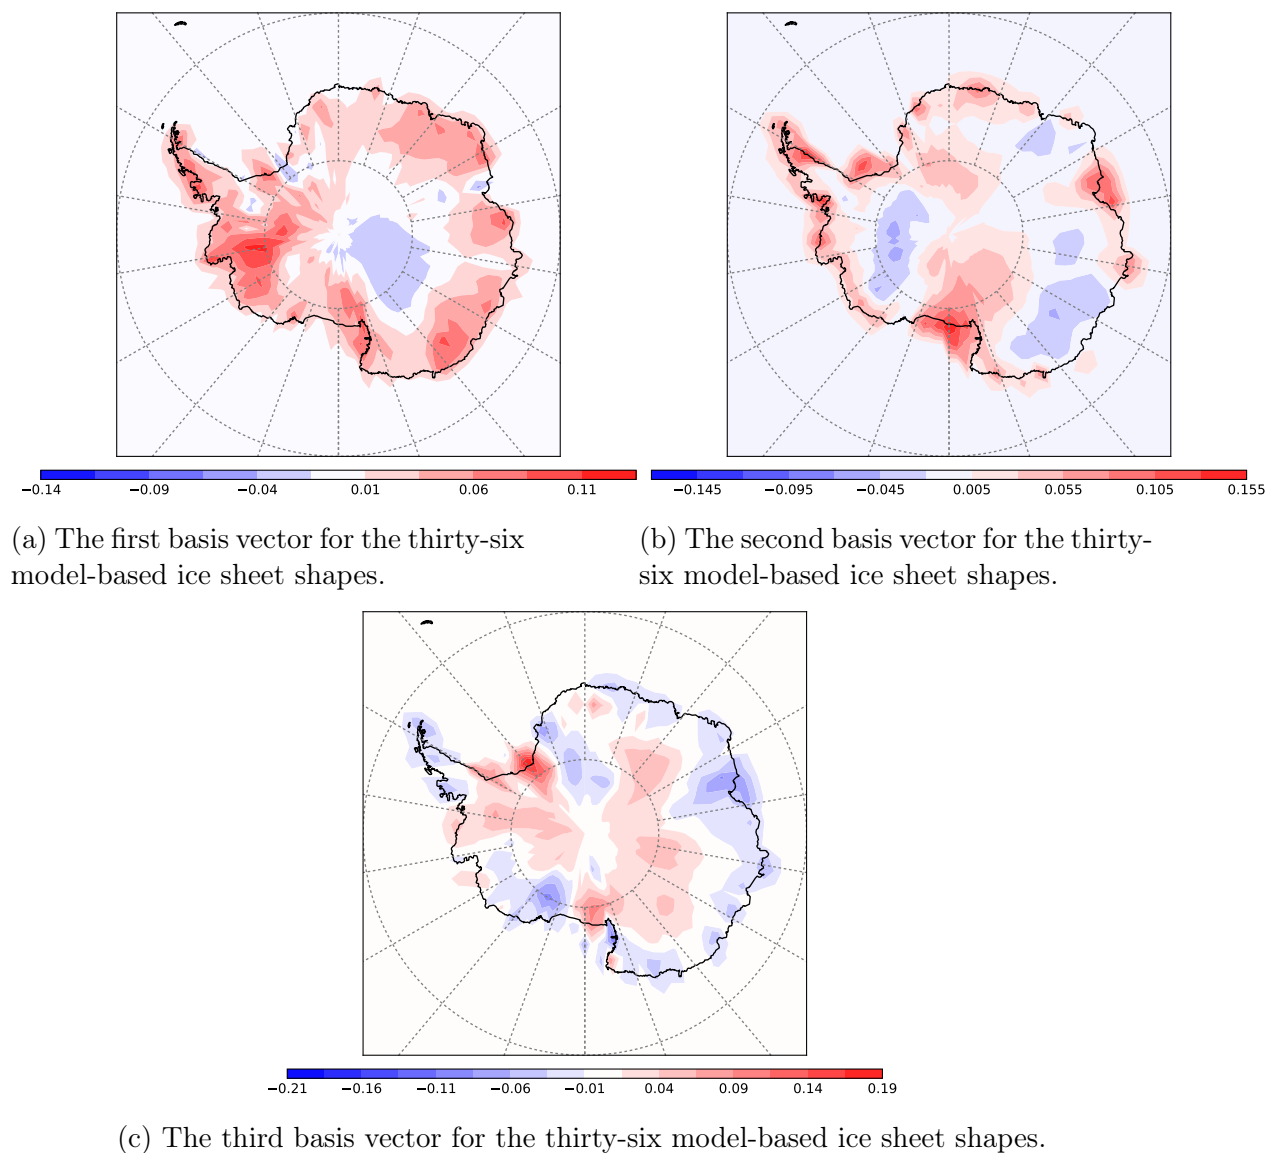


Figure 3.5: The three basis vectors found for the thirty-six model-based ice sheet shapes (Pollard and DeConto (2009), Golledge et al. (2012), Golledge et al. (2013), Bentley et al. (2014) and de Boer et al. (2017)), using the null space method (see Section 3.3). We must consider carefully how many of these to incorporate into our prior model.

reconstructions using the basis vectors created using the two different methods. We also compare how much the errors decrease when using four, five or six basis vectors for each of the two methods.

Tables 3.3, 3.4 and 3.5 show the RMSE values for the forty collected ice sheet shapes. We wish to compare the errors for each shape, with emphasis on having smaller errors for the subset of four ice sheet reconstructions that we consider most important in influencing our prior model.

When four basis vectors are used, method 2 has much smaller errors for the shapes from Peltier (2004), Argus et al. (2014) and Whitehouse et al. (2012) than method 1. The error for Briggs et al. (2014) is about the same for both methods. For the other thirty-six shapes, the errors are similar between both methods apart from the shape from Pollard and DeConto (2009), where method 1 has a much smaller error than method 2.

When five basis vectors are used, the error in method 1 for Whitehouse et al. (2012) is now much smaller than method 2. The errors for the first two shapes remain smaller for method 2 than method 1, and the error for the fourth shape is still similar between both of the methods. Method 3 now gives a much smaller error for the shape from Pollard and DeConto (2009), and the other thirty-five shapes have similar error sizes.

For six basis vectors, the errors for the first four shapes are almost identical to using five basis vectors for method 2, whilst method 1 has large decreases for three of the

first four. Whitehouse et al. (2012) and Briggs et al. (2014) both now have smaller errors for method 1 than the other method. Errors for the other thirty-six shapes are quite varied between the two methods; Pollard and DeConto (2009) is almost three times smaller in method 2 than in method 1. As there is little difference between five and six variables for the most important shapes in method 2, it is decided that the prior model will not include the sixth basis vector. This is due to the limited number of simulations we can run through HadCM3; with five prior variables instead of six, we can run fifty simulations to thoroughly test our model rather than sixty.

The average RMSE was also considered for both methods; the values can be seen in the final rows of Tables 3.3, 3.4 and 3.5. Method 1 has the smallest mean errors; this is to be expected, as PCA on all forty shapes would find the optimal scenario overall. We performed the other method wanting our basis vectors to be particularly accurate reconstructing the first four shapes, and have sacrificed accuracy for the other shapes in order to obtain this. The average loss of accuracy is minimal, so we decide to use the basis vectors from our second method. The errors for all methods and all scenarios have been small; given the average height of the Antarctic ice sheet reconstructions is around 2300m, the errors we have calculated suggest that our variables are very accurate at reconstructing the collection of ice sheet shapes.

Shape	Method 1	Method 2
Peltier (2004)	13.07	1.42
Argus et al. (2014)	12.85	6.07
Whitehouse et al. (2012)	55.98	35.63
Briggs et al. (2014)	43.23	44.21
Pollard and DeConto (2009)	40.23	62.86
Golledge et al. (2012)	24.95	23.69
Golledge et al. (2013)	21.81	23.36
Bentley et al. (2014)	35.70	38.19
de Boer et al. (2017)	7.43	8.80
de Boer et al. (2017)	6.27	6.90
de Boer et al. (2017)	8.62	8.92
de Boer et al. (2017)	9.01	9.68
de Boer et al. (2017)	6.35	6.62
de Boer et al. (2017)	9.12	9.26
de Boer et al. (2017)	8.25	8.55
de Boer et al. (2017)	12.29	13.43
de Boer et al. (2017)	13.10	16.51
de Boer et al. (2017)	12.13	15.07
de Boer et al. (2017)	8.24	9.38
de Boer et al. (2017)	6.11	6.77
de Boer et al. (2017)	10.66	13.14
de Boer et al. (2017)	13.30	14.67
de Boer et al. (2017)	6.19	6.26
de Boer et al. (2017)	9.31	9.51
de Boer et al. (2017)	7.65	8.93
de Boer et al. (2017)	6.34	6.98
de Boer et al. (2017)	9.24	9.47
de Boer et al. (2017)	9.64	10.43
de Boer et al. (2017)	6.73	6.87
de Boer et al. (2017)	9.23	9.38
de Boer et al. (2017)	8.46	8.81
de Boer et al. (2017)	12.25	13.42
de Boer et al. (2017)	10.07	13.01
de Boer et al. (2017)	11.89	14.86
de Boer et al. (2017)	7.99	8.85
de Boer et al. (2017)	6.07	6.62
de Boer et al. (2017)	10.37	12.79
de Boer et al. (2017)	13.87	15.19
de Boer et al. (2017)	7.16	7.25
de Boer et al. (2017)	9.83	10.08
Average RMSE	13.52	14.05

Table 3.3: The RMSE for each of the forty collected ice sheet shapes when using four basis vectors. We wish to prioritise reducing the errors for the first four shapes.

Shape	Method 1	Method 2
Peltier (2004)	12.87	1.38
Argus et al. (2014)	9.44	5.89
Whitehouse et al. (2012)	2.07	34.53
Briggs et al. (2014)	38.75	42.85
Pollard and DeConto (2009)	24.01	18.46
Golledge et al. (2012)	24.51	22.19
Golledge et al. (2013)	21.67	23.20
Bentley et al. (2014)	34.37	33.63
de Boer et al. (2017)	7.42	8.34
de Boer et al. (2017)	6.27	6.74
de Boer et al. (2017)	8.31	8.84
de Boer et al. (2017)	8.68	8.81
de Boer et al. (2017)	6.35	6.27
de Boer et al. (2017)	9.09	9.21
de Boer et al. (2017)	7.86	8.20
de Boer et al. (2017)	12.09	12.65
de Boer et al. (2017)	12.54	13.35
de Boer et al. (2017)	11.44	11.83
de Boer et al. (2017)	8.23	9.15
de Boer et al. (2017)	6.10	6.54
de Boer et al. (2017)	10.28	9.71
de Boer et al. (2017)	12.63	12.09
de Boer et al. (2017)	6.15	6.29
de Boer et al. (2017)	9.29	9.43
de Boer et al. (2017)	7.64	8.56
de Boer et al. (2017)	6.34	6.83
de Boer et al. (2017)	8.87	9.25
de Boer et al. (2017)	9.27	9.34
de Boer et al. (2017)	6.72	6.74
de Boer et al. (2017)	9.21	9.35
de Boer et al. (2017)	8.08	8.44
de Boer et al. (2017)	12.06	12.61
de Boer et al. (2017)	9.84	10.84
de Boer et al. (2017)	11.17	11.62
de Boer et al. (2017)	7.96	8.78
de Boer et al. (2017)	6.07	6.50
de Boer et al. (2017)	10.02	9.41
de Boer et al. (2017)	13.09	12.58
de Boer et al. (2017)	7.15	7.24
de Boer et al. (2017)	9.79	9.73
Average RMSE	11.34	11.93

Table 3.4: The RMSE of the forty collected ice sheet shapes when using five basis vectors. We wish to prioritise reducing the errors for the first four shapes.

Shape	Method 1	Method 2
Peltier (2004)	10.85	1.38
Argus et al. (2014)	4.02	5.89
Whitehouse et al. (2012)	2.04	34.53
Briggs et al. (2014)	10.73	42.85
Pollard and DeConto (2009)	18.75	6.26
Golledge et al. (2012)	19.26	22.14
Golledge et al. (2013)	21.52	23.20
Bentley et al. (2014)	32.86	22.10
de Boer et al. (2017)	7.37	7.44
de Boer et al. (2017)	5.92	6.72
de Boer et al. (2017)	8.02	8.83
de Boer et al. (2017)	7.97	8.41
de Boer et al. (2017)	5.98	6.16
de Boer et al. (2017)	8.93	8.02
de Boer et al. (2017)	7.82	6.42
de Boer et al. (2017)	12.09	7.88
de Boer et al. (2017)	12.45	7.14
de Boer et al. (2017)	11.28	8.10
de Boer et al. (2017)	7.92	7.67
de Boer et al. (2017)	6.04	6.37
de Boer et al. (2017)	8.98	8.66
de Boer et al. (2017)	11.51	11.82
de Boer et al. (2017)	6.02	5.87
de Boer et al. (2017)	9.01	7.99
de Boer et al. (2017)	7.53	7.69
de Boer et al. (2017)	5.97	6.79
de Boer et al. (2017)	8.30	9.25
de Boer et al. (2017)	8.43	8.81
de Boer et al. (2017)	6.56	6.55
de Boer et al. (2017)	9.07	8.09
de Boer et al. (2017)	8.05	6.51
de Boer et al. (2017)	12.06	7.92
de Boer et al. (2017)	9.84	6.59
de Boer et al. (2017)	11.05	7.87
de Boer et al. (2017)	7.56	8.11
de Boer et al. (2017)	5.98	6.39
de Boer et al. (2017)	8.63	8.68
de Boer et al. (2017)	2.05	12.18
de Boer et al. (2017)	7.15	6.99
de Boer et al. (2017)	9.18	8.72
Average RMSE	9.89	10.22

Table 3.5: The RMSE of the forty collected ice sheet shapes when using six basis vectors. We wish to prioritise reducing the errors for the first four shapes.

3.5 Building the Prior Model

The basis vectors can now be used to build the prior model. We use a simple linear model of the form

$$\boldsymbol{\theta} = z_1 \mathbf{a}_1 + z_2 \mathbf{a}_2 + z_3 \mathbf{a}_3 + z_4 \mathbf{a}_4 + z_5 \mathbf{a}_5 + \boldsymbol{\mu},$$

where $\boldsymbol{\theta} \in \mathbb{R}^{7008}$ is ice sheet shape, the $\mathbf{a}_i \in \mathbb{R}^{7008}$ are the basis vectors and $\boldsymbol{\mu} \in \mathbb{R}^{7008}$ is the mean ice sheet shape from the forty collected reconstructions. We need a set of scalars, $\mathbf{z} \in \mathbb{R}^5$, that will control how much of each basis vector contributes to the synthetic outputted ice sheet shapes. Using these \mathbf{z} , we want our model to be able to reconstruct the forty ice sheet shapes used to create the variables, as we want our model to output sensible estimates of the size and shape of the ice sheets. As the \mathbf{a}_i are the first i columns of V , we set \mathbf{z} to be a combination of U and Σ . However, as we performed the SVD on the two sets of ice sheet shapes separately we need to construct a new U ourselves. We have

$$X = U\Sigma V^T,$$

where X is our set of ice sheet shapes, V is the matrix of our five basis vectors and Σ the diagonal matrix of their corresponding eigenvalues. From this we find

$$U\Sigma = XV = X_{proj}$$

where X_{proj} is the data projected through V . We then have

$$U = X_{proj}\Sigma^{-1}.$$

We thus set $\mathbf{z} = (X_{proj}\Sigma^{-1})$.

After normalising the z_i 's to lie between $[-1, 1]$, we plot histograms of them to consider what distribution to give them. These can be seen in Figure 3.6. Although they appear to have a peak and tails that would suggest a Gaussian distribution, this seems less suited when we consider where the errors for the four most significant shapes lie on the plots. The four most important shapes are outliers or spread out uniformly for z_1 , z_2 and z_3 . If we took a Gaussian distribution, then these important ice sheet shapes would be treated as unlikely scenarios. We therefore set each of the z_i 's to a uniform distribution over $[-1, 1]$. This would make every ice sheet shape equally likely to occur so none of our collected shapes are disregarded.

3.6 Summary of Chapter

In this chapter, we have used a collection of ice sheet reconstructions to build a prior model describing the variation within the Antarctic ice sheets at the LGM. By developing a new form of Principal Component Analysis that allows us to prioritise one set of data over another, we have modified this method to better suit our requirements and create a model that is influenced by the most widely accepted existing ice sheet

reconstructions. The next stage of this research is to run a set of outputs from this prior model through the general circulation model HadCM3. In the next chapter, we describe the process of expert elicitation to decide on this set of ice sheet shapes.

There are some limitations to this approach. By using published reconstructions of the ice sheets to describe the vector space our shapes will lie in, we are automatically rejecting any shapes outside of this space. We therefore are assuming that these reconstructions describe all plausible shapes of the ice sheets at the LGM, and any shapes not in this subspace are implausible.

We have had to limit the number of prior model variables to five; this is to ensure we do not have too many dimensions to explore when running simulations through HadCM3. However, this means we have had to discard some variation within the data. Although we have calculated the RMSE values to check we are not throwing away too much information about the ice sheet shapes, inevitably some has been lost.




(a) Histogram of z_1 .(b) Histogram of z_2 .(c) Histogram of z_3 .(d) Histogram of z_4 .(e) Histogram of z_5 .

Figure 3.6: Histograms of the scalars, z , for the five basis vectors in our prior model. The points represent the subset of four ice sheet reconstructions that we consider most important in our analysis. They are spread out or are outliers in most of the histograms; this suggests a Gaussian distribution would not be a good fit for z .

Chapter 4

Elicitation of Ice Sheet Shapes

As described in Chapter 3, we have a prior model, built using a set of forty ice sheet reconstructions, that outputs synthetic ice sheet shapes. We now wish to use this model to create a set of synthetic ice sheet shapes to run through the global circulation model HadCM3. This will provide us with a set of input and output data with which we can build a Gaussian process emulator of HadCM3, allowing us to run MCMC in order to sample the posterior distribution of the ice sheet shape at the LGM. The shapes we choose to input are therefore important. We must consider which set of shapes would cover the space defined by our prior model evenly, to ensure important areas are not underrepresented in our HadCM3 simulations. The expense of running HadCM3 is also an important motive in considering which ice sheet shapes to input. The model is both financially and computationally expensive to run, taking twenty-four hours on the remote access super computer Archer to

simulate approximately twenty years model time. As we wish to run each of our simulations for sixty years, we are limited in how many we can input into HadCM3. Any simulations run through HadCM3 containing implausible LGM ice sheet shapes would be a waste of time and resources. We must therefore ensure that all ice sheet shapes we run are plausible for the LGM, as well as that they span the range of possible prior shapes so that they embody the uncertainty our model describes about the ice sheet at the LGM.

Mathematically, we wish to model $f : S \rightarrow D$, where S is the space the prior model describes and D is the space described by the isotope values that HadCM3 outputs. A good design of S is usually a space filling design (Pronzato and Müller, 2012).

We have already incorporated some expert judgements when we used a collection of ice sheet reconstructions to create variables for our prior models. This has allowed us to better represent uncertainty about the shape of the ice sheets at the LGM, giving us a more informed model. We therefore wish to consult with experts on which shapes should be used as input for HadCM3 and incorporate their beliefs into our decision-making process. A consultation was organised with ice modellers and LGM experts at the British Antarctic Survey, with their judgements to be used to help determine which shapes to input into the climate model.

Elicitation is an important tool when wishing to draw on expert judgements, and can play a vital role when making decisions on complex problems (O’Hagan et al., 2006, Chapter 1). Statisticians are often called on to work in a variety of different fields; often these are areas that they have no experience in. In these cases, a consultation

with experts is essential to understand what is needed. Here, we have little knowledge of what the Antarctic ice sheets looked like at the LGM so we are reliant on a consultation with experts to decide on our set of synthetic ice sheets.

An expert elicitation is of particular use when performing Bayesian analysis (O'Hagan et al., 2006, Chapter 1). When building a prior distribution, it is often useful to incorporate judgements from experts in that particular area as eliciting these opinions about parameters in the distribution can allow us to reduce uncertainty about them. This is an especially important stage when building a subjective prior. If there is limited data, the prior can be very influential on the results.

There is a wide range of literature on expert elicitation; both Dias et al. (2018) and EFSA (2014) give an overview of the process. Johnson et al. (2010) demonstrates the use of upper and lower extremes to build probability distributions describing an experts belief; however, it is reliant on the participants having sufficient knowledge of probability and statistics. Albert et al. (2012) shows the importance of building informative priors for parameters that have few observations to provide information. They also discuss how to effectively combine multiple expert judgements, modelling both the consensus and diversity of opinions. Garthwaite et al. (2005) gives a thorough overview of elicitation, motivating the use of the technique and the need to continue further research. All of these papers are good examples of the usefulness of a formal elicitation process, and consider the various biases and other challenges inherent in such a focus. In this project, a more informal procedure was necessary to engage experts and model their beliefs in an easily interpretable format.

In this chapter, we describe the process of eliciting expert opinions with reference to

O'Hagan et al. (2006). In section 4.1, we consider how to prepare for an elicitation meeting. Section 4.2 explains the process of conducting a meeting. Section 4.3 describes the follow up to an elicitation meeting and using the collected judgements. Throughout the sections, we describe our own process of expert elicitation.

4.1 Preparation

A great deal of preparation is required to have a successful elicitation process. The style of elicitation must first be decided on. This consists of two main decisions: the first is whether to consult with a single or multiple experts. Both of these options have their pros and cons (O'Hagan et al., 2006, Chapter 2). Having multiple judgements can lead to a more informed prior and it is therefore often better to consult with multiple experts. It is also rare that one person has all of the expertise required. By meeting with multiple people, their beliefs can be aggregated so that the future analysis is as informed as possible. However, meeting with multiple experts can be extremely time-consuming, particularly if you are doing so individually. Meeting with a single expert can be the most appropriate choice if there is a time constraint. If using multiple experts, we must also decide whether to gather their opinions individually or as a group. Having to consult many different experts takes a long time, but it can often lead to more detailed judgements and allows us to concentrate on their area of expertise. However, meeting in a group allows experts to discuss things amongst themselves and come to a consensus on a question. It also can lead to discussion on points not already raised, and introduce new concepts to the analysis.

The second decision is whether to conduct the elicitation through face-to-face interviews or a survey or questionnaire for the expert(s) to fill out themselves. A questionnaire is a low-cost approach to conducting an elicitation, and means each expert is provided with the exact same questions and information. An interview, however, allows the facilitator to make sure the experts understand what is being asked of them and amend later questions based on earlier answers.

When considering our own situation, we decided the best approach was a group interview with multiple experts. The people we were consulting were from a wide range of expertise (Robert Arthern¹, Richard Hindmarsh², Dominic Hodgson³, Robert Mulvaney⁴ and James Smith⁵), which we thought would benefit our research a great deal more than a single expert. The group meeting was partly due to time constraints, but also because we hoped meeting the experts at the same time would lead to a more collaborative response to our questions. Our own limited knowledge of the research area meant that sitting in on a discussion between these experts would be a valuable learning experience as well. The experts we consulted with were chosen by my supervisor, Louise Sime (BAS).

Careful thought was put in to what we wanted to achieve from the elicitation process for the current project. The amount of time available to explain the project and ask for advice was short, so it was important to decide beforehand the most effective way of explaining this work and to consider what to seek advice on. Our preparation

¹Ice sheet modeller, rart@bas.ac.uk

²Glaciologist, rcah@bas.ac.uk

³Sedimentologist, daho@bas.ac.uk

⁴Science leader of the Ice Dynamics and Palaeoclimate team, rmu@bas.ac.uk

⁵Sedimentologist, jaas@bas.ac.uk

took two forms: building an interactive version of the prior model to help convey what we were trying to do, and creating a list of specific points on which we wished to gather opinions.

4.1.1 The Interactive Model

Interactive plots and tools are an effective way of communicating methods or research to a group of people with a wide and differing range of expertise (McInerny et al., 2014). Often during the elicitation process we must explain a method that the experts have no knowledge or experience of. A visual or interactive tool can be a way of explaining methodology concisely; spending too much of the meeting explaining the research takes time away from gathering judgements.

Before we could consult the experts on which synthetic ice sheet shapes to run through HadCM3, we had to explain our project and how the prior model had been built. We were consulting with scientists with a broad range of expertise, and it was important that we present the prior model in a form that they could all understand and visualise easily. We therefore produced a visual version of the model with an interactive element rather than focusing on the statistical nature of the model. Using `matplotlib` in `python`, we built a plot of the prior model that could be manipulated with a set of sliders. A screenshot of this is shown in Figure 4.1. Each slider controls one of the five variables in the model. These sliders set the value of each z_i , determining how much of the five basis vectors contributes to the ice sheet shape

plotted above. By moving the sliders, we can see how the prior model is influenced by each variable and create new ice sheet shapes.

The z_i 's were scaled and given a uniform distribution between $[-1,1]$, as described in Section 3.5. However, for the elicitation process it was decided to expand this to $[-3,3]$ to ensure that we included as many potential ice sheet shapes as possible. Although we believed these limits would be shortened a great deal during the elicitation process, we wished to start with an overly large range of values to ensure that no feasible ice sheet shapes were being excluded before we met with our group of experts. One of the main aims of our meeting was therefore to agree on new minimum and maximum values for each of the five uniform distributions.

4.1.2 Planning the Elicitation Meeting

It is essential to plan what to cover in the elicitation meeting, as well as specifically what judgements to gather from the experts. Time is often a factor when deciding how much can be achieved; people who are well regarded in their fields are often busy. Finding a time well suited to a group of experts can prove difficult so the meeting must be used efficiently. In particular, if they come from varied backgrounds they may require different parts of your research to be explained in more detail to them. This must be factored in to the agenda of the meeting.

Our meeting was limited to an hour. To keep time spent explaining our project to a minimum, we wrote a summary of the project and shared it with the attendees the

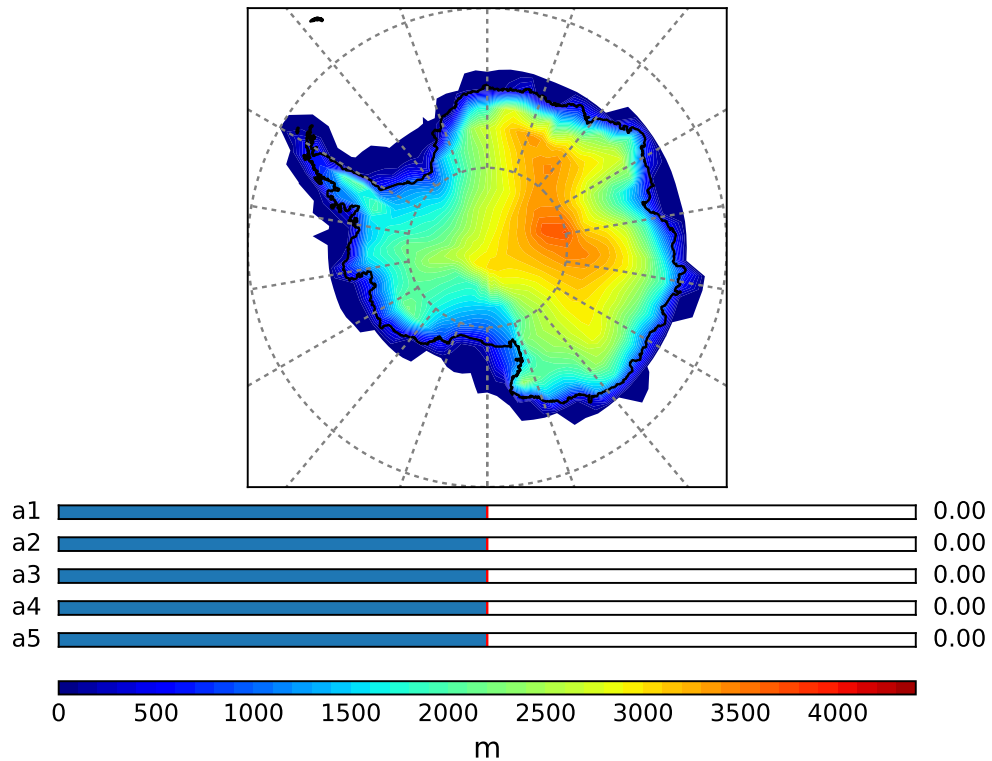


Figure 4.1: The prior model in an interactive plot. The sliders below the plot control the five variables in our model, creating synthetic ice sheet shapes that will be run through the global climate model HadCM3. The numbers on the right of the sliders show what values the z_i are currently taking. Here, they are set at 0 and only the mean shape is being plotted.

day before the meeting. This was a page long and described the motivation for the work, the shapes collected from the literature and the prior model. It ended with the following questions that we wanted them to answer

- Does our prior model capture a plausible range of ice sheet shapes?
- Are we including only ice sheet shapes that are plausible for the LGM? Are

there any that seem implausible?

- Are there shapes that are plausible but not included in our forty reconstructions?
- Where are you most uncertain about the extent and nature of the LGM ice sheet? Where are you most certain? The variation in our statistical model is plotted in Figure 4.2, does this match with your beliefs?
- We currently assume all ice sheet shapes are equally likely to occur. Could you rule out ice sheet shapes as unlikely, and suggest a set of shapes that are more likely?

These questions were written after deciding on the plan for the meeting. We would explain the aims of the research and show them the interactive plot, inviting them to play around with it themselves. This would allow them to answer the questions above by seeing which values of the sliders produced plausible and implausible shapes.

4.2 Conducting the Meeting

Although the meeting was carefully planned, we also prepared to be flexible about what was discussed. Although it can be frustrating, when experts choose to discuss




Figure 4.2: A plot of the standard deviation within our model. It shows that there is more variation in the WAIS than the EAIS.

a topic that the facilitator had not considered beforehand we can sometimes learn something important that would not otherwise come to light. This can lead to the incorporation of new ideas and techniques in to the resulting research. However, if the discussion moves too far off track it is important to bring it back round to the agenda as it is important that time is not wasted.

Several things were learnt through conducting our meeting. Having experts with different backgrounds in the room made it difficult to focus on our specified aims. In particular, explaining the statistical work behind the model to physical geographers took time as they were asking about underlying physical models they believed we should be using. Explaining the synthetic nature of the variables in our model proved difficult, and a lot of questions on the presumed underlying bedrock and other

such things had to be fielded. The differences between the original reconstructions collected and our statistical model had to be explained, all of which took up a large portion of the allotted hour. We therefore had to lower our expectations on what could be covered in the time left of our meeting. The experts often focused on areas we had not previously considered, making it awkward to ask them to move on to the areas we wanted them to focus on. A longer meeting would have given us more opportunities to discuss the areas we were interested in, but the experts we consulted with were limited in how much time they could meet with us for.

Feedback gathered from meetings can be disheartening. It is important to use elicitation methods that are as simple as possible, but it may prove in the meeting that an even simpler task is required to gather judgements effectively. The interactive model was rejected by the experts due to the limited time they had available to consider the project. Although we only had five variables in the model, this was still judged as too many for them to review in detail. The feedback we were given from one expert was that there were too many “corners” in the model space that they would need to assess before they could give an informed opinion on the output. They also often critiqued the steep edges the model sometimes created in the ice sheet, wanting to smooth them out somehow. Finding limits for the z_i 's, one of the aims of our meeting, had to be done after the meeting had finished in a one-on-one review with my supervisor, Louise, who was present.

Some useful feedback was received, however. The paper by Bentley et al. (2014) had a maximum grounding line limit, which described their estimate of the maximum ice sheet extent. The experts we consulted recommended we use this to constrain the

prior ice sheets. There was also discussion on creating a limit on maximum ice sheet thickness from the original collection of ice sheet reconstructions, which could not be explored due to the time limits in the meeting. They agreed with the plot in Figure 4.2 as correctly representing where the most uncertainty was over Antarctica during the LGM; they seemed more unsure about the characteristics of the West Antarctic ice sheet.

The meeting concluded with the experts agreeing to take part in a second elicitation process to be conducted remotely. This would involve a much simpler approach, and would focus on deciding which shapes to input to HadCM3.

4.3 The Second Elicitation Process

At the end of the elicitation meeting we agreed that I would put together a set of shapes for the experts to review; these would have the Bentley et al. (2014) grounding line limit plotted over it to act as a constraint and be within the new z_i limits agreed with Louise. It was decided to create a set of eighty shapes as this was the most we felt we could ask the experts to assess thoroughly, given that this was a voluntary exercise. We would then ask them to recommend a third of the total to be discarded before submitting the rest to HadCM3, leaving us with around fifty simulations to run.

The z_i constraints agreed with Louise were added to the prior model. These ranges were $z_1 = [-0.5, 0.5]$, $z_2 = [-0.3, 0.75]$, $z_3 = [-0.6, 0.6]$, $z_4 = [-0.55, 0.5]$ and $z_5 =$

$[-1, 1.2]$. A maxi-min Latin hypercube (Joseph and Hung, 2008) within these ranges was then used to find a design of z_i values that spanned the whole space described by this model. We chose this method because this draws a sample from a set of variables with uniform distributions, maximising the minimum distance between design points. This ensures that the design points are distributed as evenly as possible across the model space. A pairs plot of this design is shown in Figure 4.3; we can see that the points are scattered evenly in each plot, showing that no area is unrepresented in our design.

This design of z_i values was used to create eighty shapes; these can be seen in the Appendix. The shapes were shared with the five experts we had consulted with at the elicitation meeting, each of whom was given an individual spreadsheet to complete with their comments on each shape and their recommendations. It was decided to conduct this elicitation stage on an individual basis in order to collect unbiased opinions and prevent anchoring (O’Hagan et al., 2006, Chapter 3); had one person shared their opinions before the others there would be a strong possibility that the other experts would base their judgements on what had already been shared.

Two of the experts reviewed the shapes and returned the files to me; it proved difficult once we were acting remotely to encourage the experts to continue to engage in the process. The two suggested designs we did receive are shown in Figure 4.4. One of the experts who reviewed the ice sheet shapes recommended rejecting over sixty of them, stating that a recurring issue was that the shapes extended past the ice extent limit imposed by Bentley et al. (2014). The coarseness of the longitude/latitude grid used by HadCM3 was a possible cause of this. The grid used by HadCM3 is

$3.75^\circ \times 2.5^\circ$ longitude/latitude, whilst the grounding line from Bentley et al. (2014) was on a much finer scale. This meant it was difficult to create shapes for HadCM3 input that followed the grounding line exactly; although the ice sheet was usually only a few metres above sea level when it overlapped slightly with this constraint the expert systematically rejected these shapes. The other expert who engaged in the second elicitation process recommended running fifty-one of the ice sheet shapes through HadCM3.

The final design was decided on by combining these judgements. If one or both of them approved of a shape, it was run through HadCM3. If they both rejected or were unsure of it, it was discarded. This produced a set of forty-nine shapes to be run through the climate model. A pairs plot of the final design is in Figure 4.5. It covers the whole of the basis fairly well, with few parts not represented.

4.4 Summary of Chapter

Expert elicitation is a useful process for creating informative prior distributions, thus helping to determine the direction of research. It is particularly useful when performing research in an area about which the statisticians have little knowledge; consulting with experts in that field ensures research is as informed and impactful as possible.

Although often a long process that requires a great deal of thought and planning, the outcomes of an elicitation process can ensure that the research is focused on

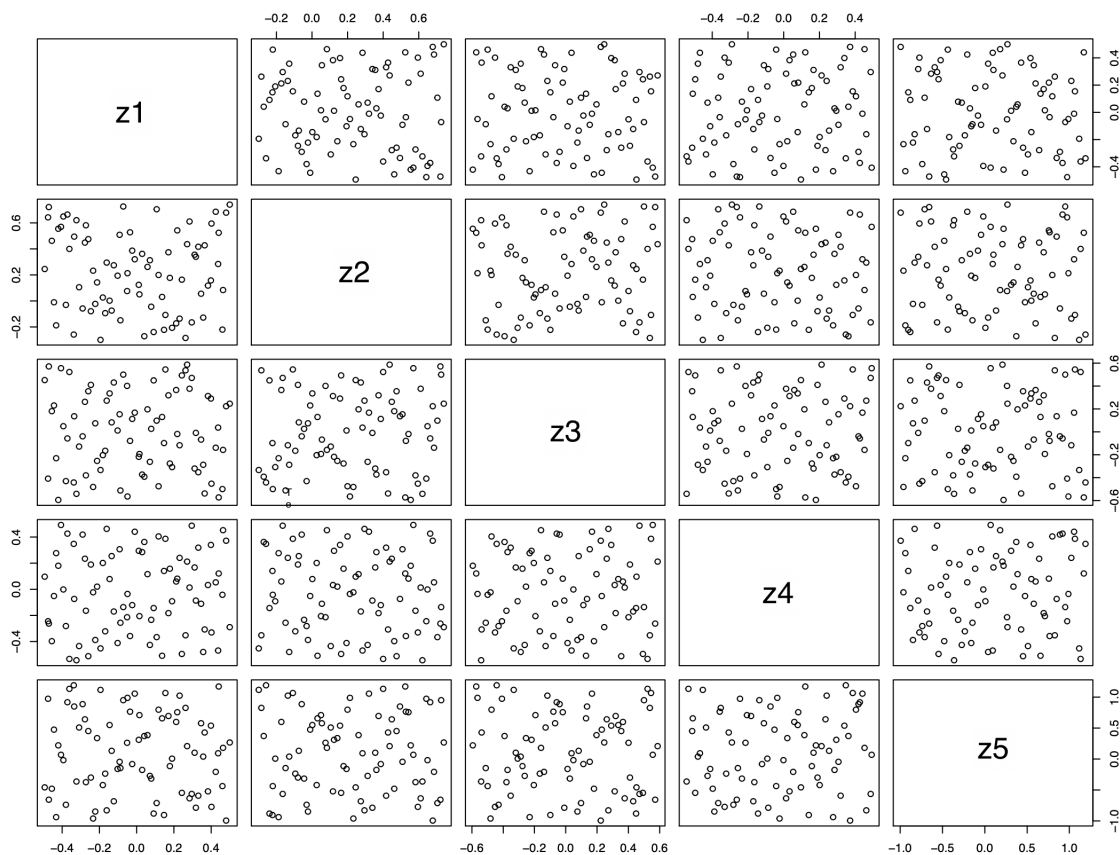


Figure 4.3: The 2D projects of the design of 80 shapes we created after the first round of expert elicitation. The points cover the entire basis evenly.

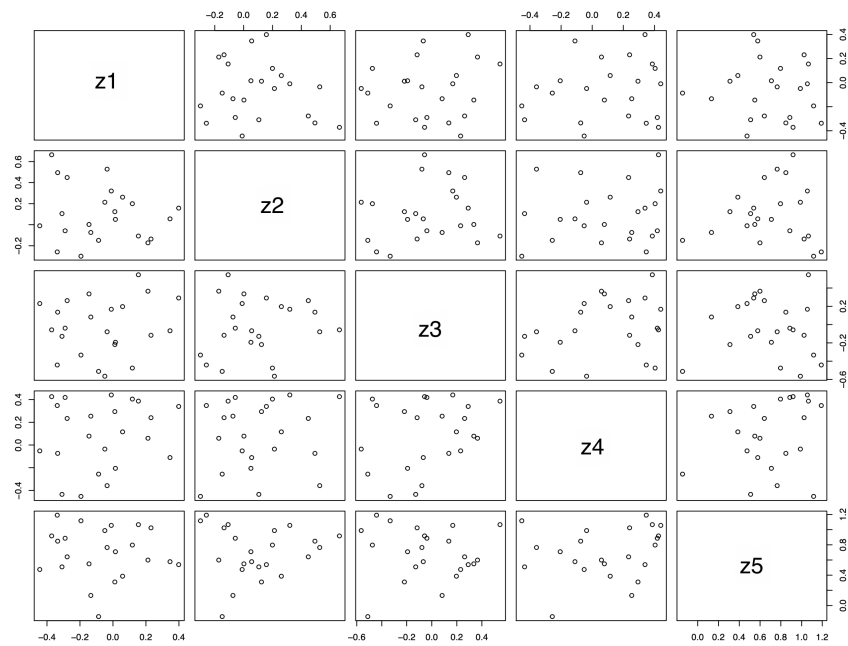
relevant and important facets of the subject area. A certain amount of flexibility is required to run an effective elicitation process as new thoughts and ideas will be suggested throughout. Whether to incorporate these or disregard them takes careful consideration.

We performed an informal elicitation of expert opinions on plausible Antarctic ice sheet shapes at the LGM. From this we set constraints on our prior variables, z_i , and

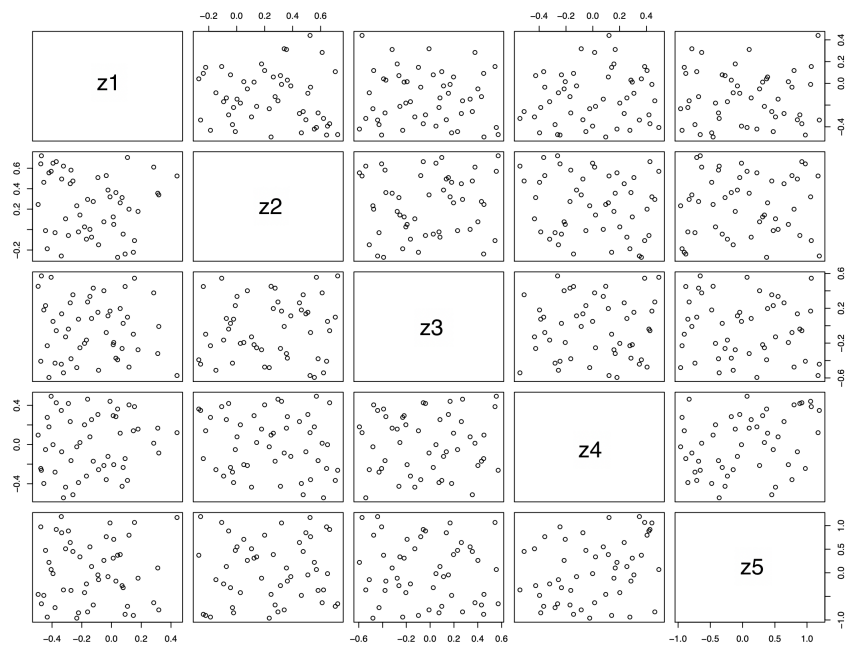
then created a design of ice sheet shapes for a second elicitation process.

Although a useful process, our elicitations could certainly have been improved. Our first elicitation process was hampered by time constraints. By only having an hour to discuss our project with the experts, we were unable to discuss every point that we had wanted to cover. It also proved difficult to encourage the experts to engage in the second elicitation process remotely, which meant that a certain amount of pragmatic compromises had to be used to incorporate judgements and to proceed with the elicitation.

Following this process we now have an informed design of shapes to run through HadCM3. In the next chapter, we discuss running HadCM3 and using the output to build a Gaussian process emulator of the climate model.



(a)



(b)

Figure 4.4: The two designs we received after the second round of expert elicitation. There is some agreement between the two experts, though the first discarded a lot more shapes and therefore has left gaps in the space we are trying to cover.

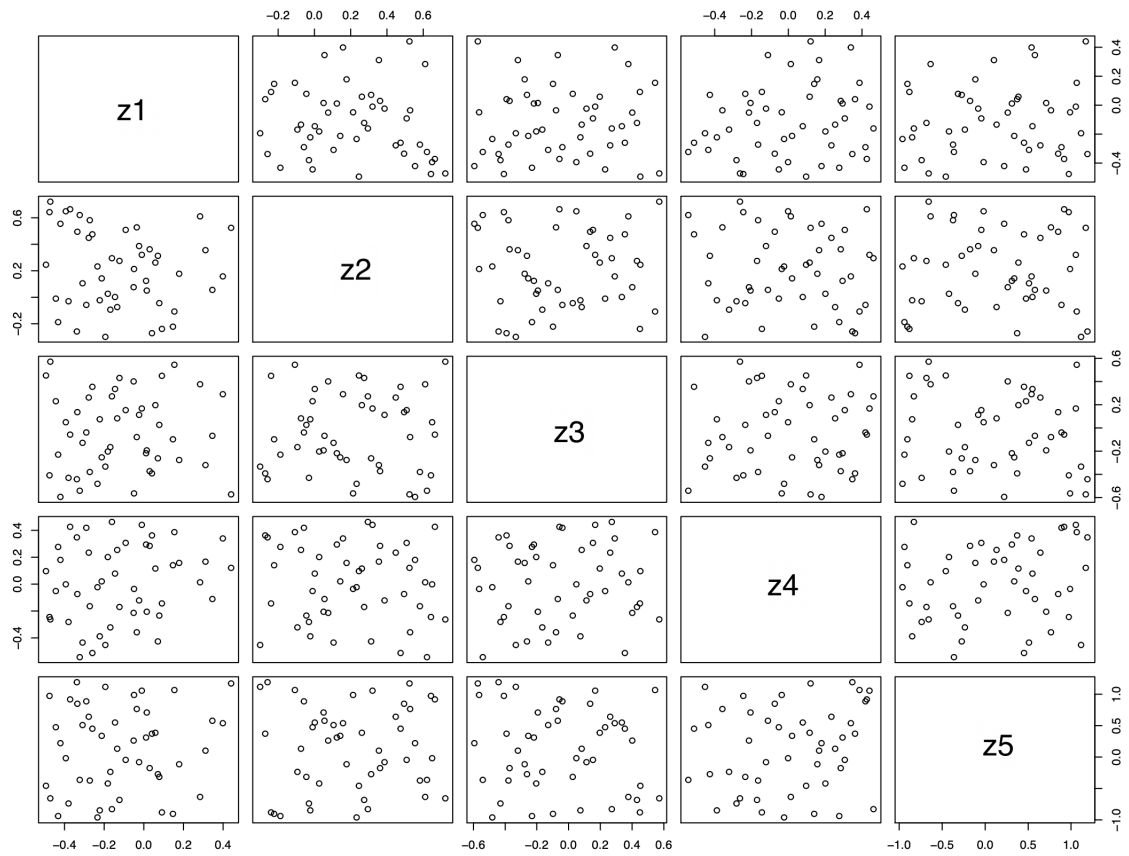


Figure 4.5: The final design we made by combining the two sets of designs that the experts recommended. The five dimensional basis that the variables describe is well covered by our design.

Chapter 5

Building The Emulator

Having agreed on a design of forty-nine ice sheet shapes in Chapter 4, we now wish to run the GCM, HadCM3. This will provide us with a set of $\delta^{18}O$ simulations paired with the synthetic orographies. From this, we can attempt to build a Gaussian process emulator of HadCM3 that will allow us to emulate the GCM and explore the relationship between ice sheet shape and $\delta^{18}O$ without incurring the costs associated with HadCM3. We can then use this emulator along with the prior model described in Chapter 3 to sample from the posterior distribution of the ice sheet shape and reduce uncertainty about the ice sheet at the LGM.

The motivations for building the emulator have already been described in Chapter 4. HadCM3 is very slow to run, simulating approximately twenty years a day. As we must allow time for the model to spin up and require several decades to simulate a

climate observation, we must run each simulation for sixty years and therefore require several days to complete each simulation. We cannot run all simulations at once, as we have a limit of running sixteen simulations simultaneously on Archer, the UK National Supercomputer Service. Running a model on an HPC is a time consuming and difficult process, adding more motivation to limit the number of simulations we create. Since Archer also costs a great deal to run; we are reliant on a limited number of HPC credits to run HadCM3 (allocated by the British Antarctic Survey) and so can only afford a small number of simulations. We therefore cannot use HadCM3 to sample the posterior distribution, as we cannot run it the thousands of times required to create a large enough sample. Instead we rely on a Gaussian process emulator which will use the five prior variables as inputs and output $\delta^{18}O$ anomalies. This will allow us to model the relationship between ice sheet shape and water isotopes more cheaply than HadCM3, so that later it can then be used as part of the calibration model; MCMC can then be used to sample from the posterior distribution of the ice sheet shape, reducing the uncertainty around it.

In this chapter, we describe the process of running HadCM3, comparing the output of the GCM to the collected ice core data and building Gaussian process emulators at each ice core site.

5.1 Running HadCM3

The global circulation model we use in this analysis is the Hadley Centre Coupled Model Version 3 (HadCM3), a UK Met Office isotope-enabled coupled atmosphere-

ocean GCM. Motivation for using this model can be found in Chapter 1. A description of how isotopes are implemented in HadCM3 can be found in Tindall et al. (2009).

From the elicitation process described in Chapter 4, we created forty-nine orography files describing plausible Antarctic ice sheet shapes at the LGM. We ran control LGM simulations with these orography files attached through the BlueCrystal supercomputer at the University of Bristol to activate the isotopes. Ten of them failed to run; we attempted to fix this by rounding all values in the orography files to the nearest integer, which fixed eight of them. Two of them continued to refuse to run, resulting in us having forty-seven simulations. We built forty-seven simulation files in PUMA with control LGM GHG values and orbital forcing, and starting conditions from the BlueCrystal simulations. These were then inputted into HadCM3 and ran for 60 years each.

5.1.1 **Issues with HadCM3**

We faced several issues when running HadCM3. Some of these we were able to solve; others had to be worked around. The first has previously been mentioned: two of the simulations would not run through the BlueCrystal supercomputer. Due to time constraints, we had to discard these and proceed with forty-seven simulations.

Jobs crashed frequently due to a fault in HadCM3 that created negative pressure values. There was no way of working around it, so jobs had to be continually resub-

mitted and spun up past the point where it crashed. This increased the amount of time it took to finish running the simulations.

Twenty-six of the simulations had corrupted November files for one of the sixty simulated years. We could not find a cause for this. It was decided to remove the whole corresponding year so as to avoid any seasonal bias caused if we had removed just the corrupted month.

5.1.2 Processing Scripts

Once the simulations were finished, we ran the output through post-processing scripts on Jasmin, a scientific data analysis environment administered by the Centre for Environmental Data Analysis (CEDA). The code for these can be seen in Appendix B. These converted the model output into climate variables including $\delta^{18}O$, the oxygen isotope estimates that we will be using in our analysis. The first ten years of the data were treated as spin-up time, the time for which it took for the variables in the model to reach a steady state, leaving us with fifty years of simulated data. We have both monthly, annual, and a fifty year average of the climate variables.

Before we performed an exploratory data analysis, we transformed our HadCM3 output into anomalies, the difference between our output and a pre-industrial control simulation. This was to remove any underlying climate biases HadCM3 may have (Domingo et al., 2020). We used a pre-industrial simulation to create $\delta^{18}O$ anomalies

from our HadCM3 output. We compared these anomalies to observations from ice core anomalies, which we created from the observed data by subtracting the average $\delta^{18}\text{O}$ value of the last one thousand years from the LGM value, the average of the values dated either side of 21Ka BP.

5.2 Exploratory Data Analysis

Before we began building our Gaussian process emulator, we compared our $\delta^{18}\text{O}$ anomalies to the ice core data to check how well HadCM3 simulated values close to the observations. The locations of the ice core sites that we used are shown in Figure 5.1. They are Byrd (Blunier and Brook, 2001), Mount Moulton (Popp, 2008), Siple Dome (WAIS Divide Project Members et al. (2013) and Brook et al. (2005)) and WDC (WAIS Divide Project Members et al. (2013) and Steig et al. (2013)) in the West Antarctic ice sheet, and 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. (1999), Steig et al. (2000)) and Vostok (Petit et al., 1999) in the East Antarctic ice sheet. The two points in grey are Berkner and Fletcher. These ice cores are yet to be published. We include them here as we intend to use them as test sites for our emulator and posterior distribution when they are in the public domain.

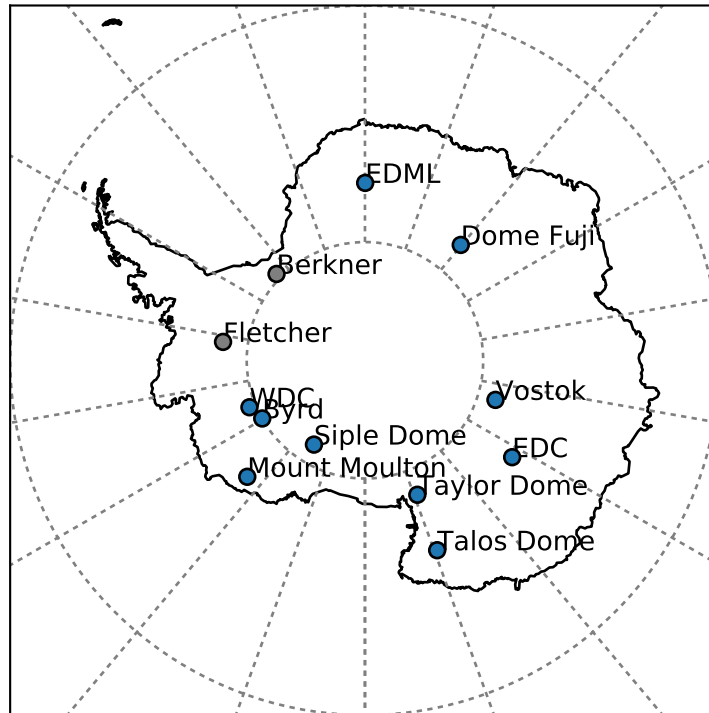


Figure 5.1: The locations of ice core drilling sites that we intend to use in this research. The two grey points, Berkner and Fletcher, are sites that have yet to publish their data. We hope to use these as test sites once we have access to the ice core data.

5.2.1 Comparing simulations of WAIS and EAIS

We compared the HadCM3 output to the ice core observations to see if we had successfully simulated the $\delta^{18}O$ anomalies, using the latitude/longitude grid cell in our HadCM3 simulations corresponding to the location of the ice core site as a comparison. We plotted the HadCM3 $\delta^{18}O$ anomalies against the change in elevation at that grid point to see if there is an obvious relationship between $\delta^{18}O$ and ice sheet size. Change in elevation is found by subtracting a pre-industrial orography from each of the ice sheet shapes we ran through HadCM3, so we can see how each ice sheet shape from our design differs to our pre-industrial control orography, and how this has affected the corresponding $\delta^{18}O$ anomaly values. We then add in a line corresponding to the ice core observation to see if our simulations are close to this value.

We can see in Figures 5.2 and 5.3 that there is a difference in how similar the HadCM3 output is to the ice core observations for the two ice sheets. The water isotope values in the West Antarctic ice sheet have been simulated well by HadCM3, with the $\delta^{18}O$ anomalies all close to the observations from the ice cores. The isotopes in the East Antarctic ice sheet have not been simulated well. The exception is Taylor Dome, suggesting that close to where the ice sheets meet, we have simulated plausible orographies. We check how well we have modelled both ice sheets by plotting pairs plots of the two groups of ice core sites. We can see again, in Figures 5.4 and 5.5, that the HadCM3 simulations have produced $\delta^{18}O$ anomalies that are close to the observations in the West Antarctic ice sheet but failed to simulate low enough values in the East Antarctic ice sheet. For this reason, we decide to focus

on building emulators for the West Antarctic ice core sites, and will incorporate the East Antarctic ice core sites later.

5.3 Gaussian Process Emulation

Gaussian process emulation is used to predict the response of a complex and expensive simulator using a limited number of inputs and outputs from the simulator (Chang et al., 2014b). Here, we wish to emulate HadCM3, focusing on the response of $\delta^{18}\text{O}$ anomalies to changes in the shape of the Antarctic ice sheets at the LGM.

A Bayesian framework for GP emulation is described in detail in Kennedy and O’Hagan (2001). A simulator is run on n design points $\mathbf{z}_1, \dots, \mathbf{z}_n \in \mathbb{R}^k$, giving outputs $y_1, \dots, y_n \in \mathbb{R}$. The emulator interpolates the y_i values, allowing us to predict the simulator output at any input value.

A Gaussian process emulator is defined as a probability distribution $f(\cdot)$ where $f(\mathbf{z})$ is normally distributed for all $\mathbf{z} \in \mathcal{Z}$. The distribution has mean function $m(\cdot)$ where $m(\mathbf{z}) = \mathbb{E}[f(\mathbf{z})]$ and covariance function $k(\cdot, \cdot)$ where $k(\mathbf{z}, \mathbf{z}') = \text{cov}[f(\mathbf{z}), f(\mathbf{z}')]$. We therefore have

$$f(\mathbf{z}) \sim \mathcal{N}(m(\mathbf{z}), k(\mathbf{z}, \mathbf{z}')).$$

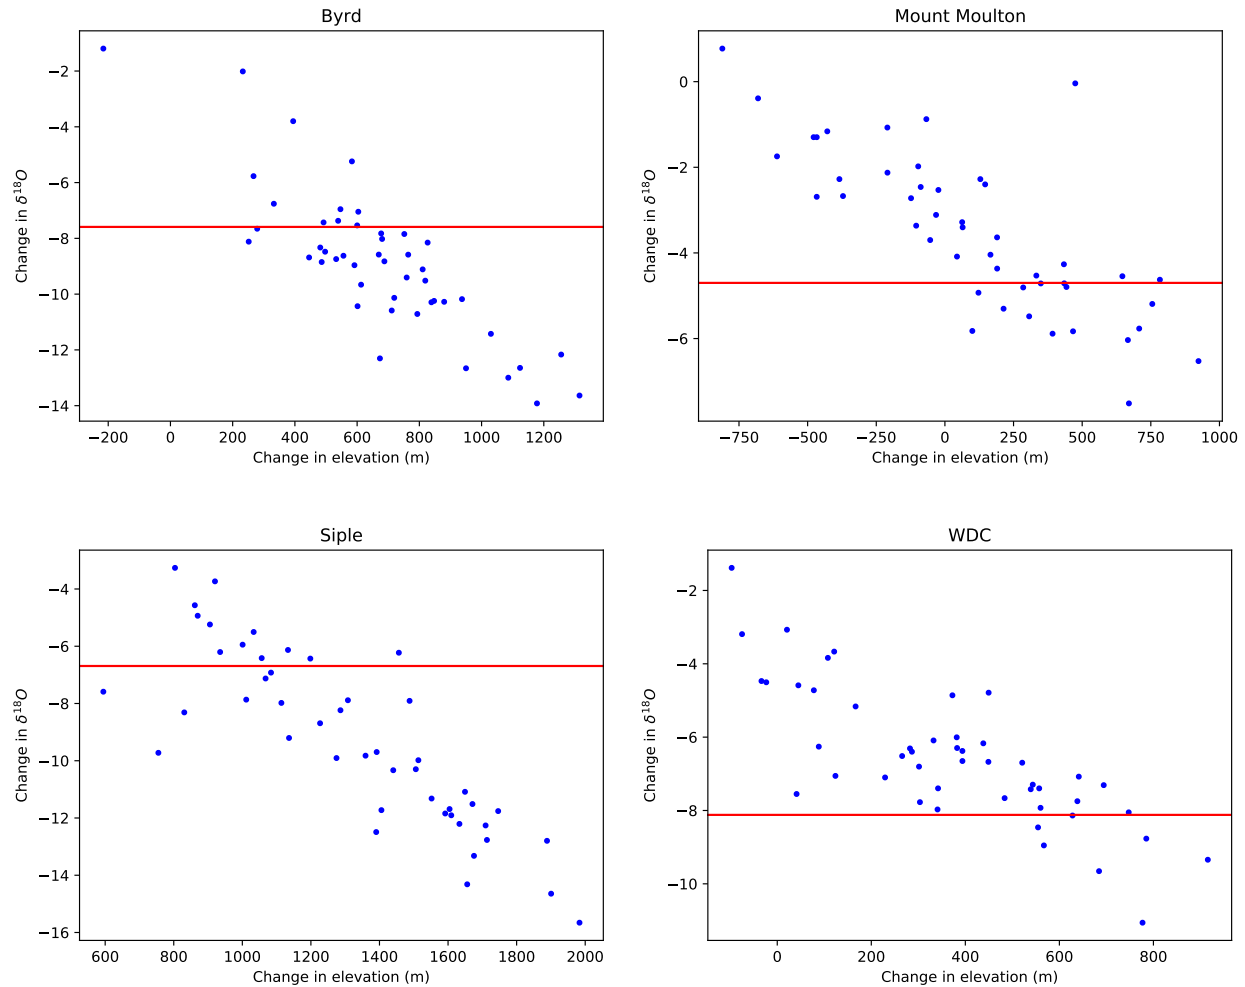


Figure 5.2: Plots of $\delta^{18}O$ anomalies against elevation anomalies for the four ice core sites from the West Antarctic ice sheet. The red lines represent the anomaly from the corresponding ice core; at all four sites HadCM3 appears to have simulated the $\delta^{18}O$ anomalies close to the observation, suggesting that the ice sheet shapes we are considering are close to the true shape at the LGM.

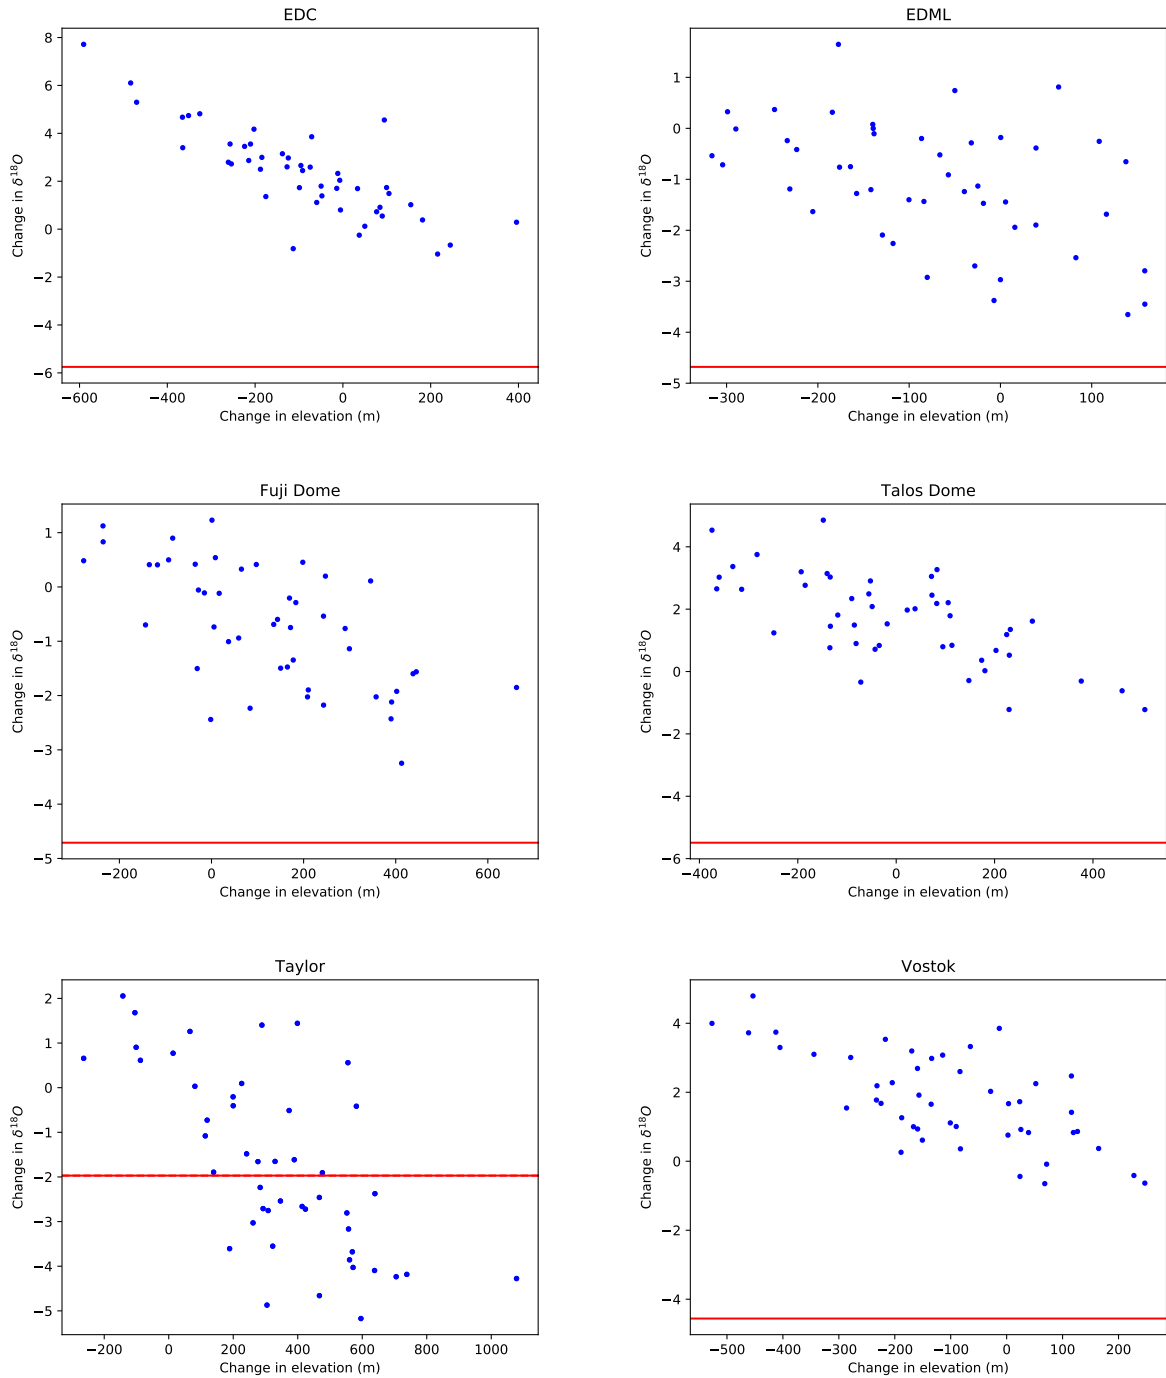


Figure 5.3: Plots of $\delta^{18}O$ anomalies against elevation anomalies for the six ice core sites from the East Antarctic ice sheet. The red lines represent the anomaly from the corresponding ice core. HadCM3 appears to have struggled to simulate $\delta^{18}O$ values close to the ice core observations. These plots suggest it may be because the ice sheet shapes we created from the prior model do not have high enough elevation across the EAIS. The only exception is the Taylor Dome ice core, suggesting that close to where the two ice sheets meet, the East Antarctic ice sheet was well modelled.

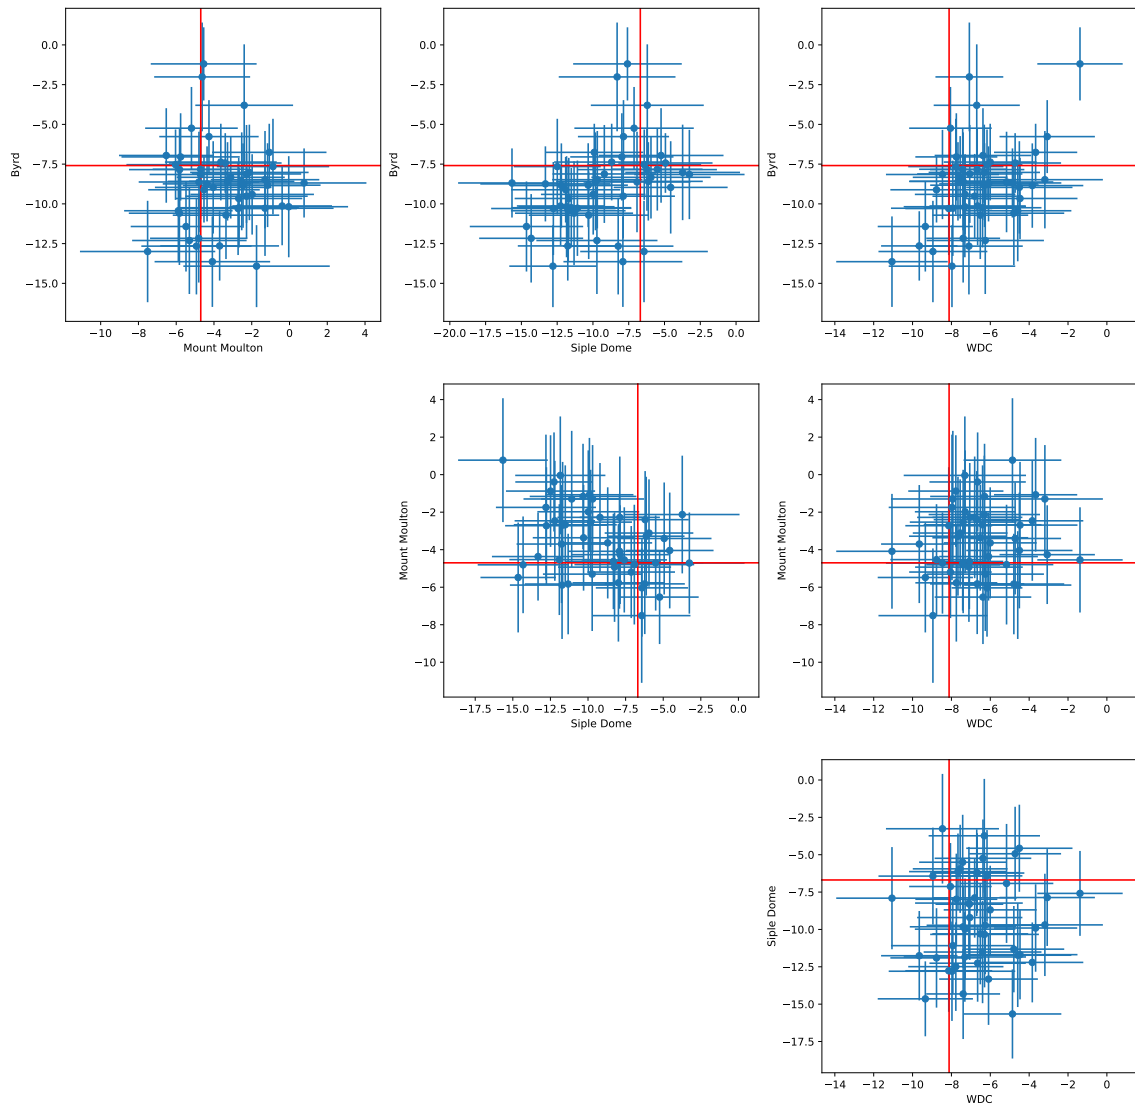


Figure 5.4: Pairs plots of $\delta^{18}\text{O}$ anomalies at the four West Antarctic ice core sites with error bars two standard deviations in length. The blue points are the output from HadCM3 simulations; the red lines are the observations from ice cores. The red lines cross in the same region as the simulations are plotted, with many of the error bars overlapping the observation, showing that HadCM3 has successfully produced $\delta^{18}\text{O}$ anomalies close to the observations at the LGM.

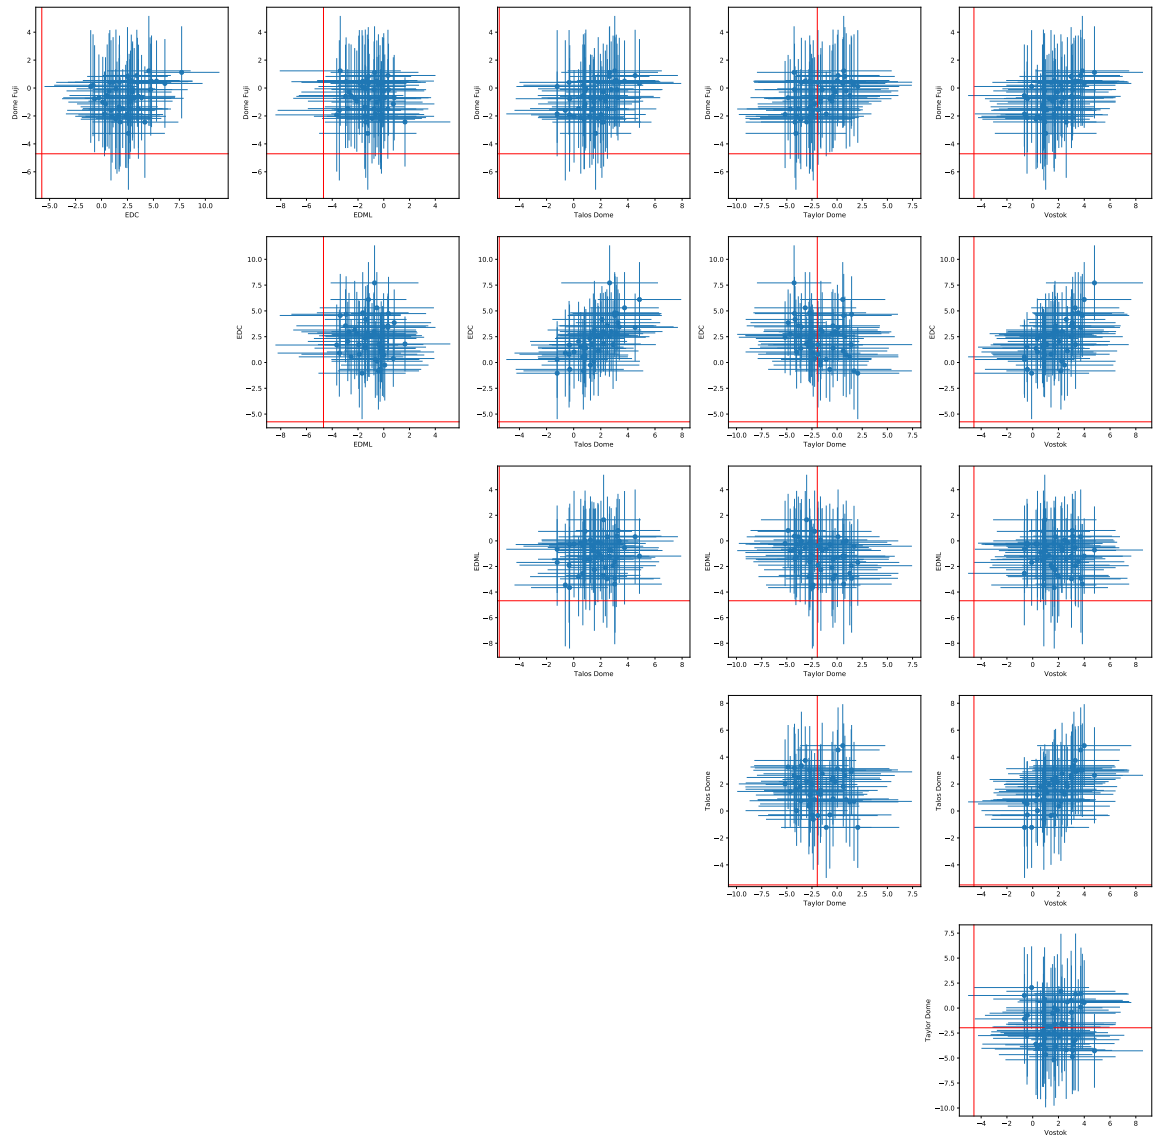


Figure 5.5: Pairs plots of $\delta^{18}O$ anomalies at the six East Antarctic ice core sites with error bars two standard deviations in length. The blue points are the output from HadCM3 simulations; the red lines are the observations from ice cores. The red lines cross at much lower points than where the simulations are plotted and few of the error bars overlap the observation; HadCM3 has not produced $\delta^{18}O$ anomalies close to the LGM observations.

5.3.1 Choosing a kernel function

Our choice of covariance function, or kernel, determines how well the GP emulates the simulator. The most common kernel choice, which we use throughout this chapter, is the radial base function (RBF) kernel, also known as the exponentiated quadratic or squared exponential kernel. The RBF kernel has a Gaussian form defined as

$$k(\mathbf{z}_i, \mathbf{z}'_i) = \alpha^2 \exp\left\{\frac{-(\mathbf{z}_i - \mathbf{z}'_i)^T(\mathbf{z}_i - \mathbf{z}'_i)}{2\rho_i^2}\right\}, \quad i = 1, \dots, d,$$

where α^2 is a variance hyperparameter, ρ_i is the lengthscale for each of the d dimensions. The lengthscale controls the decay in correlation when the distance between the input points increases.

We also experimented with other covariance functions, including linear, Matern 3/2 and Matern 5/2, but these did not give a better fit and so we chose to use the most commonly used kernel. Plots of emulators with these other kernels (for the Byrd ice core site) can be seen in Figure 5.10; we can see they do not give a better fit than the RBF. We also tried as incorporating elevation as an extra variable for our model. However, there was no improvement in the emulator so we removed this addition.

5.3.2 Building our emulator

In this chapter, we use the Python package GPy and the maximum likelihood method to find the values of our hyperparameters α , ρ from our kernel function, and σ^2 , our

nugget term, the variance term representing noise in the data. In the next chapter, we explain how to do it in a Bayesian sense, using MCMC.

1-D elevation example

As we can see in Figures 5.2 and 5.3, there is a clear correlation between change in elevation and $\delta^{18}\text{O}$ anomalies. We therefore started with a 1-D example, using the elevation anomalies as input points and $\delta^{18}\text{O}$ anomalies as output points. We run the model

$$\mathbf{y}, E \in \mathbb{R}^{47}, \quad \alpha, \rho, \sigma^2 \in \mathbb{R},$$

$$\mathbf{y} \sim \mathcal{N}(m(E), K(E|\alpha, \rho) + \sigma^2 \delta_{\mathbf{z}_i, \mathbf{z}_j})$$

for an ice core site, where E is a vector of elevation anomalies for each of the HadCM3 simulations and \mathbf{y} is a vector of the corresponding $\delta^{18}\text{O}$ anomalies, σ^2 is the nugget term and $\delta_{\mathbf{z}_i, \mathbf{z}_j}$ is defined as

$$\delta_{\mathbf{z}_i, \mathbf{z}_j} = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{if } i \neq j. \end{cases}$$

The nugget term allows the emulator to create non-deterministic predictions, even at the input points. This allows us to more closely emulate the climate system that HadCM3 is simulating; even using our 50-year averaged $\delta^{18}\text{O}$ output, simulations with equivalent inputs would give different output values. This adds more difficulty to the process, as there is irreducible noise in the data that we cannot emulate perfectly due to the temporal and spatial averaging of the data we are working with.




Figure 5.6: The elevation-only emulator for the Byrd ice core site. Our large length-scale has created a very smooth function. We can also see there is little noise within our data; the majority of the points are within the shaded confidence interval.

A discussion of the fixed variance terms that we consider due to these and other possible errors is given in Section 6.6.

The emulator for the Byrd ice core site is plotted in Figure 5.6. We have a large lengthscale ($\rho = 2311.14m$), which results in a very smooth function. Our variance term is large ($\alpha^2 = 190.27$), but the nugget term is small ($\sigma^2 = 2.02$), suggesting our training data has little noise in it.

We perform leave-one-out cross validation (LOOCV) for the elevation only emulators

for all four West Antarctic ice core sites to test how well the emulators predict the observations. Results can be seen in Figure 5.7, where the red points are the observations and the blue points are the corresponding predicted values. We also report the RMSEs for each emulator, all of which are fairly small, suggesting that our emulators are matching HadCM3 well, although the error bars in Figure 5.7, all of length two standard deviations, suggest there is a large amount of uncertainty in these emulators.

Including the prior variables

What we are interested in is the relationship between our prior model variables and the $\delta^{18}O$ anomalies. Figure 5.7 has also shown only using elevation as an input leads to emulators with large variance terms. We therefore built emulators at each of the four West Antarctic ice core sites using the prior model values as input points. We store the prior model values used to build the orographies in our HadCM3 simulations in a matrix $Z \in \mathbb{R}^{47 \times 5}$.

Each emulator is defined as

$$Z \in \mathbb{R}^{47 \times 5}, \quad \mathbf{y} \in \mathbb{R}^{47}, \quad \alpha, \rho, \sigma^2 \in \mathbb{R}$$

$$\mathbf{y} \sim \mathcal{N}(m(Z), K(Z|\alpha, \rho) + \sigma^2 \delta_{\mathbf{z}_i, \mathbf{z}_j}).$$

The values of the three hyperparameters for each of the emulators can be seen in Table 5.1. The lengthscales here are smaller than the lengthscale in our 1D example,

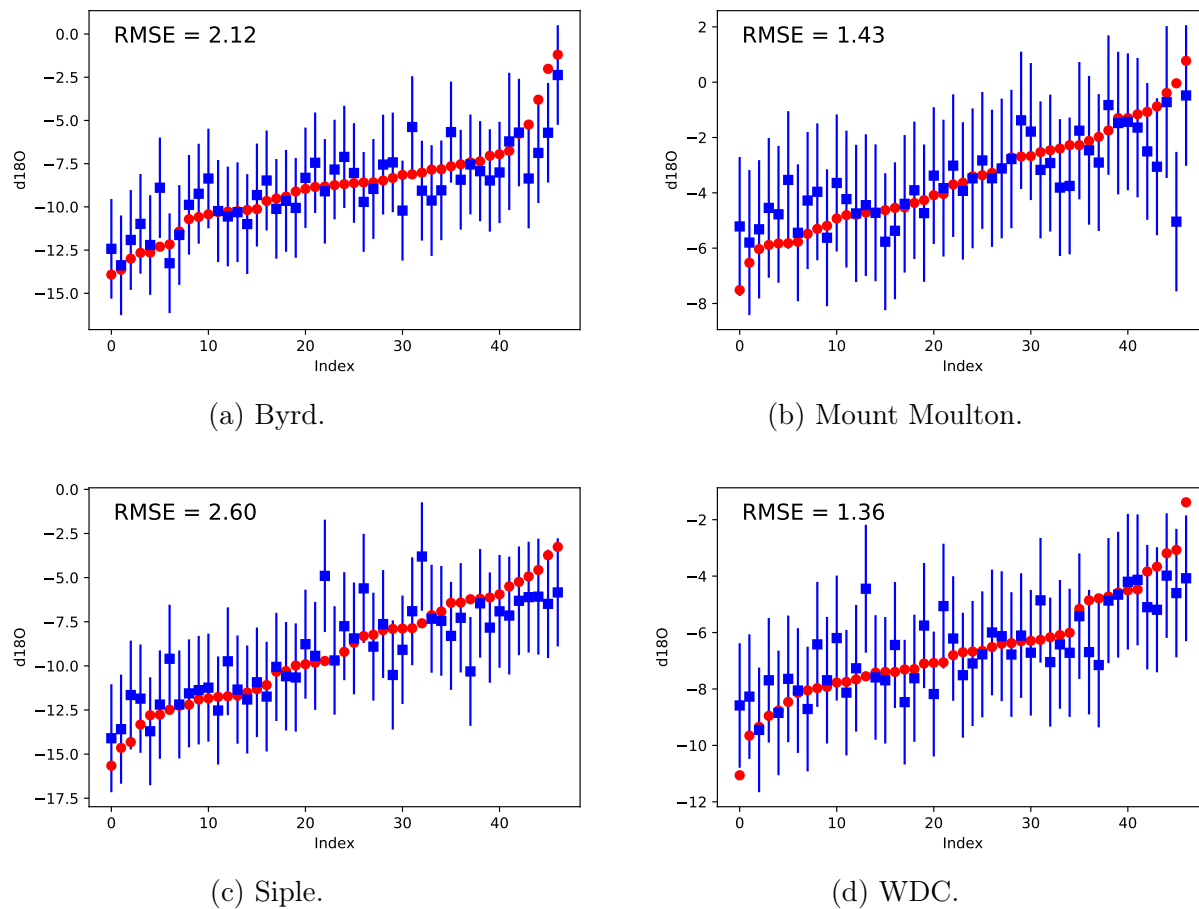


Figure 5.7: Comparison of HadCM3 simulated $\delta^{18}\text{O}$ anomalies (in red) against leave-one-out cross validation predictions of our GP emulators (in blue, with error bars of length two standard deviations included) when using elevation as our input data. We can see that although most of our predictions are within two standard deviations of the simulations, the size of the error bars shows that there is a great deal of uncertainty in our 1D emulators.

	Variance (α)	Lengthscale (ρ)	Nugget term (σ)
Byrd	113.76	3.00	0.61
Mount Moulton	249.71	7.01	1.05
Siple	222.10	3.23	1.38
WDC	245.83	6.67	0.82

Table 5.1: ML estimates for the variance, lengthscale and nugget term parameters for the emulators at the four West Antarctic ice core sites. We can see the lengthscales are much smaller than the 1D example emulator, meaning points far from each other in the sample space have large covariance terms.

creating functions with small covariance terms for points far away in the sample space. The variance terms are fairly large, allowing the function to vary a great deal.

We run leave-one-out cross validation (LOOCV) again on these emulators, the results of which are in Figure 5.8. These emulators are matching the HadCM3 output well, with predictions very close to the observations. We also report the RMSEs for each emulator, all of which are smaller than the RMSEs in Section 5.3.2, suggesting an improvement in the performance of our emulators when working in higher dimensions. The nugget term is representing a large component of the uncertainty; with values around one, that means the smallest possible 95% confidence interval is width four, assuming there is no emulator uncertainty. As the $\delta^{18}O$ anomalies from our HadCM3 simulations vary by approximately ± 10 , the nugget term is an important part of our emulator.

We also plot in Figure 5.9 predictions against observations with error bars, of both one and two standard deviations in length, and the $y = x$ line to see how closely predictions match observations. We can see that the majority of the two standard deviation error bars cross the $y = x$ line. We also calculate the percentage of our

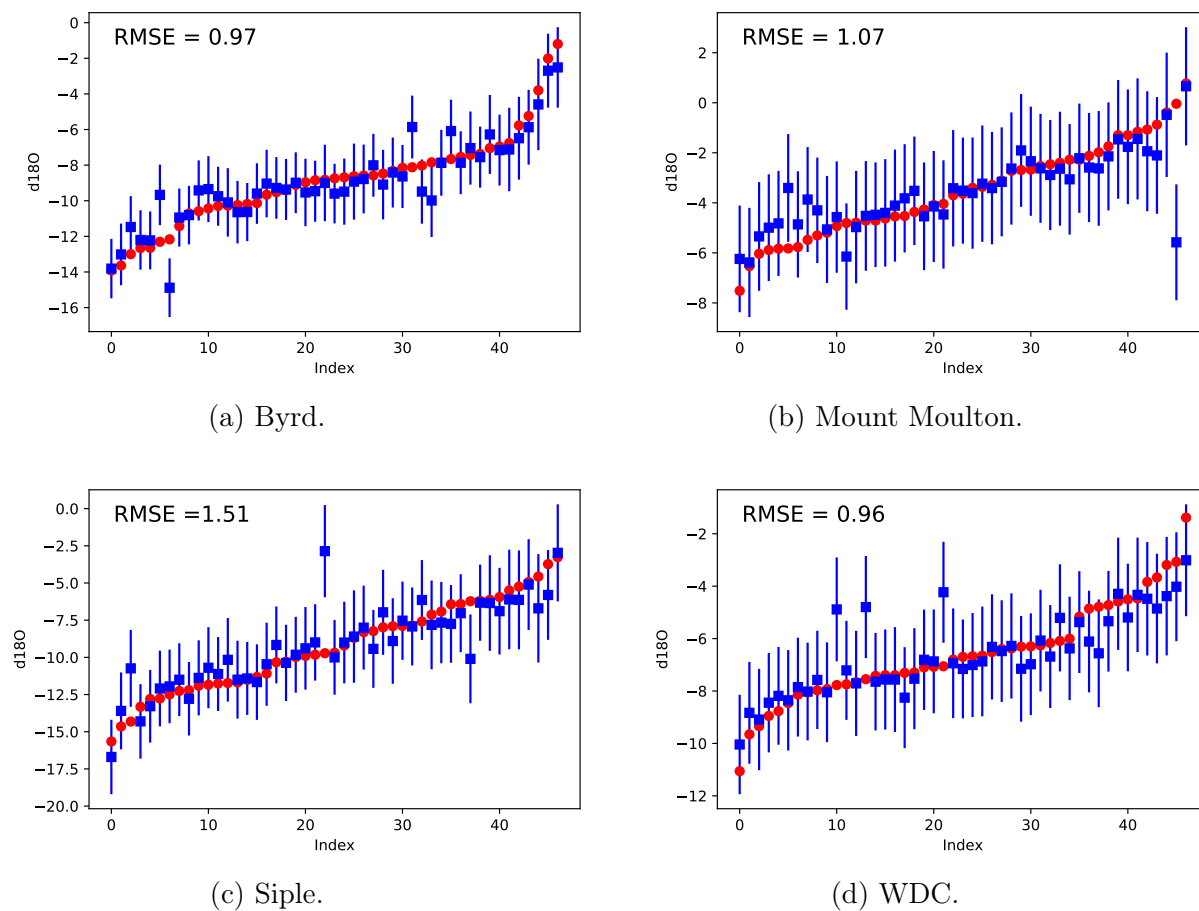


Figure 5.8: Comparison of HadCM3 simulated $\delta^{18}\text{O}$ anomalies (in red) against leave-one-out cross validation predictions of our GP emulators (in blue) when using our prior variables as input data for the emulator at the four WAIS sites. The RMSE for each site is also included as a legend. We can see that using the five prior variables rather than the elevation has produced emulators that are better at predicting HadCM3.


	Within 66% CI	Within 95% CI
Byrd	79%	91%
Mount Moulton	87%	96%
Siple	87%	94%
WDC	83%	94%

Table 5.2: A table showing the percentage of predictions from the emulators at the four WAIS ice core sites within 66% and 95% confidence intervals of the HadCM3 data.


predictions within 66% and 95% confidence intervals of our observations, shown in Table 5.2. We can see that the majority of our predictions are within two standard deviations of the observations.

5.4 Incorporating the EAIS sites


We now look at building emulators for the six ice core sites in the EAIS, using the same methods described in Section 5.3.2. Figures 5.11 and 5.12 show that our emulator is working well; the predictions are all close to the true values and the RMSEs are small, suggesting HadCM3 is being well emulated. Although the set of ice sheet shapes we submitted to HadCM3 may have had lower elevation than the true LGM shape, our emulator is able to simulate HadCM3 well enough to model the isotope anomalies in the EAIS as well as it models anomalies in the WAIS.




(a) Byrd.



(b) Mount Moulton.




(c) Siple.




(d) WDC.


Figure 5.9: Observations vs. predictions, with error bars included, for the 5D emulator at the four WAIS sites. The blue error bars are of length one standard deviation, whilst green represents two standard deviations. We can see that for all sites the majority of the green error bars cross this line, suggesting that our emulators are predicting HadCM3 output well.



(a) Matern32 kernel.




(b) Matern52 kernel.




(c) Linear kernel.


Figure 5.10: Comparison of HadCM3 simulated $\delta^{18}\text{O}$ anomalies (in red) against leave-one-out cross validation predictions of our 5D GP emulators (in blue) when using three different kernel functions for the Byrd ice core site. We can see that none of them fit as well as the RBF function, with all RMSE values larger than those in Figure 5.8.




(a) EDC.




(b) EDML.




(c) Fuji.



(d) Talos.




(e) Taylor.




(f) Vostok.


Figure 5.11: Comparison of HadCM3 simulated $\delta^{18}\text{O}$ anomalies (in red) against leave-one-out cross validation predictions of our GP emulators (in blue) when using our prior variables as input data for the emulator at the six EAIS sites. The RMSE for each site is also included as a legend. Again the emulators seem to be predicting HadCM3 well.




(a) EDC.




(b) EDML.




(c) Fuji.



(d) Talos.



(e) Taylor.



(f) Vostok.

Figure 5.12: Observations vs. predictions, with error bars included, for the emulator at the six EAIS sites. The blue error bars are of length one standard deviation, whilst green represents two standard deviations. We can see that for all sites the majority of the error bars cross this line, suggesting that our emulators are predicting HadCM3 output well.

5.4.1 Using an Anisotropic Kernel

After testing the emulator with an isotropic kernel, that is a kernel where the length-scale is fixed for each dimension, we test an anisotropic kernel to see if having different lengthscales for each dimension helps the fit. Details on this function and the reasons behind it are given in Section 6.5.3. Figures 5.13 and 5.14 show that the fit is improved when we use an anisotropic function, with the RMSE values smaller for all ice core sites.

5.5 Summary of chapter

Gaussian process emulation is a useful method when working with expensive simulators. It allows us to approximate the simulator in a fraction of the time, allowing us to more extensively explore this system.

In this chapter we have used our prior model variables and HadCM3 output to build a Gaussian process emulator of HadCM3. This has emulated the relationship that HadCM3 models between orography and $\delta^{18}O$, allowing us to see the effects that a change in ice sheet shape has on the isotope values. The tests we have run on our emulators suggests that they are modelling HadCM3 well, as well as producing $\delta^{18}O$ anomalies close to the observations we have from ice cores. In Chapter 6, we calibrate our model, combining the prior model from Chapter 3 with our emulators to sample from the posterior distribution of the ice sheet shape at the LGM.

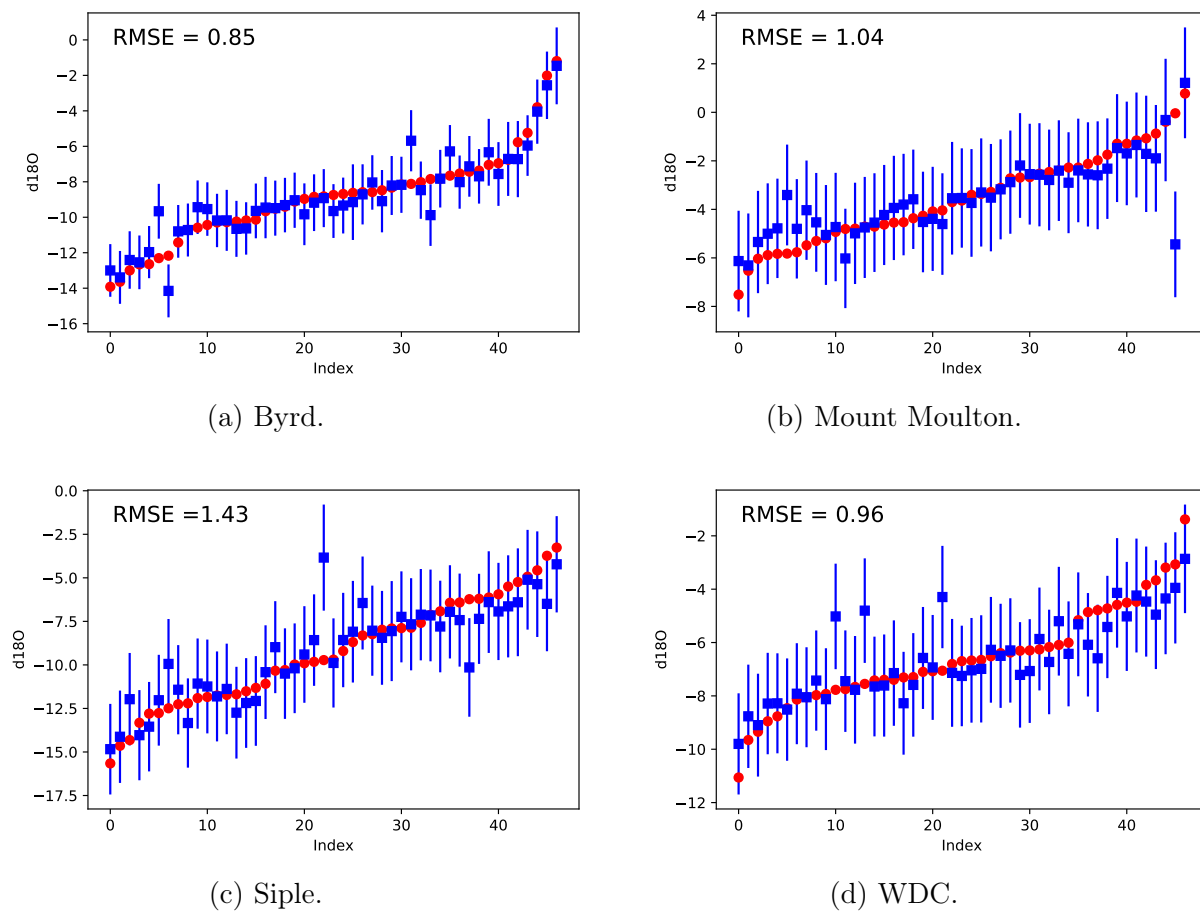





Figure 5.13: Comparison of HadCM3 simulated $\delta^{18}\text{O}$ anomalies (in red) against leave-one-out cross validation predictions of our GP emulators (in blue) when an anisotropic kernel is used in our 5D emulators for the four WAIS sites. We can see that the fit is now improved, and the RMSE are all smaller than the emulator with an isotropic kernel.




(a) EDC.




(b) EDML.




(c) Fuji.



(d) Talos.



(e) Taylor.



(f) Vostok.

Figure 5.14: Comparison of HadCM3 simulated $\delta^{18}\text{O}$ anomalies (in red) against leave-one-out cross validation predictions of our GP emulators (in blue) when an anisotropic kernel is used in our 5D emulators for the six EAIS sites. Again there is an improvement in the fit and the RMSE, although less so than for the WAIS emulators.

Chapter 6

Calibration of the Model

In the previous chapter, we successfully built a GP emulator of the GCM HadCM3 and are now able to emulate the relationship between ice sheet orographies and $\delta^{18}\text{O}$ anomalies that HadCM3 describes. In this chapter we will use this emulator, and the prior model built in Chapter 3, to sample from the posterior distributions of our five prior variables, allowing us to create ice sheet shapes based on the observations collected from ice cores. We describe the processes of calibrating our model using Markov Chain Monte Carlo (MCMC), testing our model, and performing sensitivity analyses. Geographical interpretations of our results are given in Chapter 7.

6.1 Markov Chain Monte Carlo methods

Markov Chain Monte Carlo (MCMC) is a method of sampling from a probability distribution by drawing a series of correlated samples that will converge to the target distribution (Geyer, 1992). This is done when it is too difficult to sample from the distribution directly.

We wish to simulate values $\{x_1, x_2, \dots\}$ of a random variable $x \sim p(x)$. We do this by constructing a Markov Chain (MC) $x_{t+1} \sim q(\cdot|x_t)$. By accepting or rejecting x_{t+1} according to an acceptance ratio, we simulate a MC that has $p(x)$ as its stationary distribution. There are many methods available to perform MCMC; here, we use No-U-Turn Sampler (NUTS), (Hoffman and Gelman, 2011), a variant of Hamiltonian Monte Carlo (HMC).

6.1.1 Hamiltonian Monte Carlo (HMC)

Hamiltonian Monte Carlo (HMC) is an MCMC algorithm that converges more quickly than other algorithms such as Gibbs sampling or Metropolis-Hastings. Although these other methods are simpler, the behaviour of the random walk can be affected by the sensitivity of correlated parameters. We are especially hampered when working in high dimensions, as there is a much larger space for the random walk to explore.

The auxiliary variables in HMC, such as step size and number of steps, help the chains converge more quickly to the target distribution by proposing moves to new points in the sample space that have a high probability of acceptance. This means that shorter chains are can be used to approximate the distribution.

In this chapter, all of our MCMC is performed in the computer language Stan, which can be called from R or Python and performs MCMC. Initially we used MCMC to sample from the posterior distributions of our prior variables by calibrating just our prior model before then incorporating our GP emulator, in order to check first if our model could be sampled from successfully.

6.2 1D model

Initially we used a model that was intentionally simple, with elevation as the only parameter to describe $\delta^{18}O$ anomalies. As demonstrated by Figures 5.2 and 5.3, there is a clear linear relationship between these two variables. We started with a linear emulator for one ice core site and ran MCMC assuming that the relationship between the $\delta^{18}O$ anomalies and elevation was

$$y \in \mathbb{R}^{47}, \quad E' \in \mathbb{R}^{47 \times 2}, \quad \beta \in \mathbb{R}^2, \quad \sigma_1 \in \mathbb{R}$$

where

$$E' = (\mathbf{1} \quad E),$$

contains the elevation values as well as a column of ones to model an intercept parameter,

$$\boldsymbol{\beta} \sim \mathcal{N}(0, 100), \quad \sigma_1 \sim \text{Inv_Gamma}(1, 1)$$

$$y|E', \boldsymbol{\beta} \sim \mathcal{N}(E'\boldsymbol{\beta}, \sigma_1^2)$$

$$y_{obs} \sim \mathcal{N}(E'_{obs}\boldsymbol{\beta}, \tau_1^2)$$

$$E'_{obs} = (\mathbf{1} \quad E_{obs})$$

where y represents the $\delta^{18}O$ anomalies from the HadCM3 simulations, $\boldsymbol{\beta}$ are linear model parameters and σ_1 is the standard deviation HadCM3 model error. $\boldsymbol{\beta}$ and σ_1 are given uninformative prior distributions so that we can see what we learn from our data alone. y_{obs} is the ice core observation; we generate possible elevation values, E_{obs} , for this observation using the posterior samples of $\boldsymbol{\beta}$, with a measurement error τ_1 . We set $\tau_1 = 0.5$ from the results of Keller et al. (2018). We sample four chains, each 5000 samples in length.

We appear to learn a lot about how elevation affected $\delta^{18}O$ anomalies; Figure 6.1 shows posterior predictive distribution of $\boldsymbol{\beta}$ and elevation against $\delta^{18}O$ anomalies. The horizontal blue line is the ice core $\delta^{18}O$ observation from the Byrd ice core, with measurement error $\tau_1 = 0.5\text{‰}$ plotted on each side as dashed lines. The vertical blue lines represent values of E_{obs} , generated elevation estimates corresponding to y_{obs} . The estimates of elevation are all within the range required to simulate the $\delta^{18}O$ anomaly. We can see that all generated values of elevation are within the range of what the model should predict given the uncertainty around the ice core observation.




Figure 6.1: Output from our 1D model in Section 6.2. The points are the values from our HadCM3 simulations. The black lines are some of the models made using the sample values of β from the MCMC runs, with the red line representing the mean values of the linear model. The horizontal blue line is the ice core observation from the Byrd ice core, with measurement error of 0.5‰ plotted on each side as dashed lines. The vertical blue lines are E_{obs} , simulated elevation (m) values that correspond to the ice core observation, y_{obs} .

6.3 Using all five prior variables

Although there is a clear relationship between elevation and $\delta^{18}O$ anomalies, we are interested in how the prior model variables affect the isotopic values. By sampling from the posterior distributions of the z_i , the prior variables described in Chapter 3, we hope to create plausible LGM ice sheet shapes with less uncertainty. The z_i are the five variables we created using our null space variant of PCA, as described in Section 3.3.3.

We start by modelling a single ice core site at a time to see how well our model is working. Given the discussion in 5.2.1, we focus on the WAIS sites initially.

Our prior model was of the form:

$$\boldsymbol{\theta} = \boldsymbol{\mu} + z_1 \mathbf{a}_1 + z_2 \mathbf{a}_2 + z_3 \mathbf{a}_3 + z_4 \mathbf{a}_4 + z_5 \mathbf{a}_5$$

where

- $\boldsymbol{\theta}$ is ice sheet shape.
- \mathbf{a}_j , $j = 1, \dots, 5$ are the five principal components of the collected reconstructions. \mathbf{a}_1 , \mathbf{a}_2 and \mathbf{a}_3 are made from the four most “important” shapes whilst \mathbf{a}_4 and \mathbf{a}_5 are from the other thirty-six shapes.
- $\boldsymbol{\mu}$ is the mean shape of the forty ice sheet reconstructions.

- z_j , $j = 1, \dots, 5$ are prior variables which determine how much of each of the principal components contributes to the synthetic ice sheet shape θ .

We have two data sets, the input and output data from our HadCM3 simulations. Z is a 47×5 matrix of the prior model variable values used to build the orographies used in our 47 HadCM3 simulations; more details of this can be found in Chapter 4. \mathbf{y} is a vector of length 47, containing the simulated $\delta^{18}O$ anomalies from HadCM3 at the latitude/longitude grid cell corresponding to an ice core site. We firstly transform our input data to

$$Z' = (\mathbf{1} \quad Z)$$

where Z' contains the prior variable values as well as a column of ones to model an intercept parameter. For a single ice core site we use the model

$$\mathbf{y} \in \mathbb{R}^{47}, \quad Z' \in \mathbb{R}^{47 \times 6}, \quad \boldsymbol{\beta} \in \mathbb{R}^6, \quad \sigma_1 \in \mathbb{R},$$

$$\boldsymbol{\beta} \sim \mathcal{N}(0, 100), \quad \sigma_1 \sim \text{Inv_Gamma}(1, 1),$$

$$\mathbf{y}|Z', \boldsymbol{\beta} \sim \mathcal{N}(Z'\boldsymbol{\beta}, \sigma_1^2), \tag{1}$$

$$y_{obs}|Z'_{obs}\boldsymbol{\beta} \sim \mathcal{N}(Z'_{obs}\boldsymbol{\beta}, \tau_1^2).$$

$\boldsymbol{\beta}$ and σ_1 are given uninformative prior distributions with fairly wide variances to not bias the model output. Z_{obs} are possible values of Z , generated using posterior samples of $\boldsymbol{\beta}$.

The trace plots in Figure 6.2 shows that the chains have converged for each of the five prior variables. Figure 6.3 shows that the prior and posterior distributions of

the prior variables z_{obs} are very similar, suggesting that the output from HadCM3 is not contributing much information to the posterior distribution. We are not learning much about our prior variables with this model. Learning anything about the posterior distribution is proving very difficult, as through our current method we are trying to learn about a five-dimensional space, the prior model, using a single observation, an ice core site. We therefore build a multivariate model that uses all four WAIS ice core sites at the same time. This allows us to learn about the ice sheet shape from four points instead of one, hopefully producing an updated posterior distribution.

6.4 The multivariate model for the WAIS ice core sites

Our model is similar to the one described in 6.3. The input data, Z' , remains the same, but the output \mathbf{y} is now multivariate. It describes the distribution of each \mathbf{y}_i using the values of Z as:

$$\begin{aligned} \mathbf{y}_i &\in \mathbb{R}^{47}, \quad Z' \in \mathbb{R}^{47 \times 6}, \quad \boldsymbol{\beta}_i \in \mathbb{R}^6, \quad \sigma_{1,i} \in \mathbb{R}, \\ \boldsymbol{\beta}_i &\sim \mathcal{N}(0, 100), \quad \sigma_{1,i} \sim \text{Inv_Gamma}(1, 1), \\ \mathbf{y}_i | Z', \boldsymbol{\beta}_i &\sim \mathcal{N}(Z' \boldsymbol{\beta}_i, \sigma_{1,i}), \quad i = 1, \dots, 4, \end{aligned} \tag{2}$$

generating a different $\boldsymbol{\beta}_i$ for each ice core site.




Figure 6.2: Trace plot of Z_{obs} , generated using posterior samples of β .

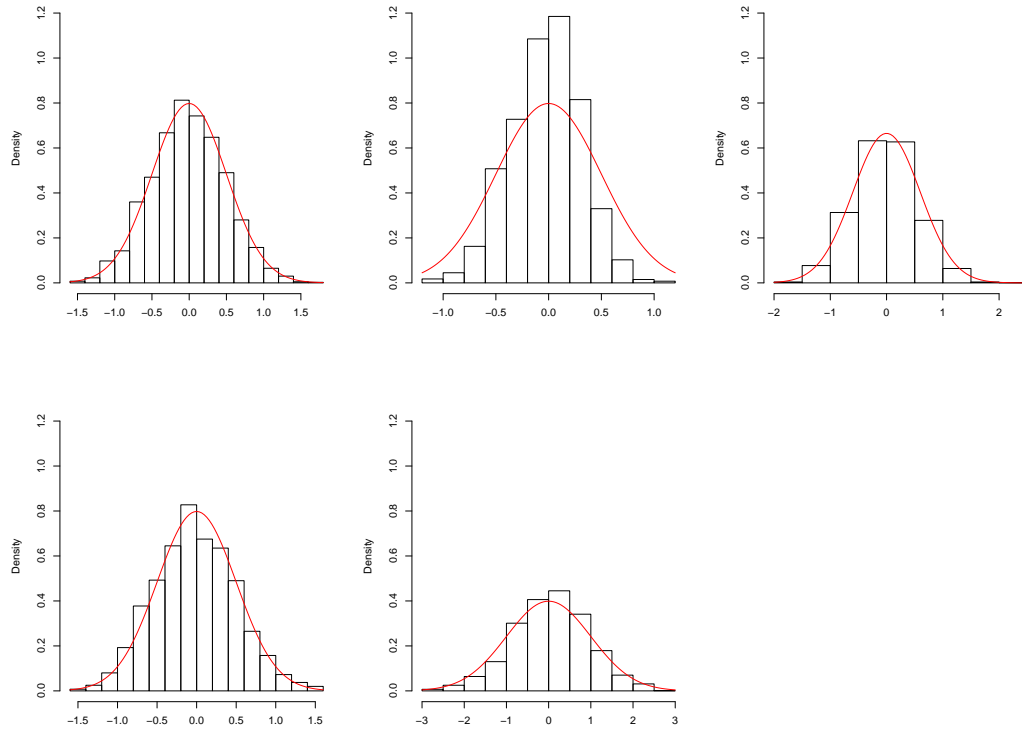


Figure 6.3: Histograms of the posterior distributions for Z_{obs} from model 1 in Section 6.3, with density curves of the prior distributions included for comparison. They are very similar, suggesting the distributions have not been updated much by the information from the HadCM3 simulations.

We use the observations of $\delta^{18}O$ anomalies from ice cores to sample from the posterior distributions of the z_j , finding values that would create ice sheet shapes corresponding to the $\delta^{18}O$ anomalies. We also include in the model

$$\mathbf{z}_{obs} \in \mathbb{R}^5, \quad \mathbf{z}_{obs} \sim \mathcal{N}(\mathbf{0}, \text{diag}(0.5, 0.5, 0.6, 0.5, 1))$$

$$\mathbf{z}'_{obs} = (1 \quad \mathbf{z}_{obs}), \quad \mathbf{z}'_{obs} \in \mathbb{R}^6$$

$$y_{obs,i} | \mathbf{z}_{obs}, \boldsymbol{\beta} \sim \mathcal{N}(\mathbf{z}'_{obs} \boldsymbol{\beta}_i, \tau), \quad i = 1, \dots, 4,$$

where $y_{obs,i}$ is the anomaly at each ice core site i and τ is a fixed error. We wish to learn more about \mathbf{z}_{obs} , the ‘true’ values of the prior variables that would create an ice sheet shape that gives $\delta^{18}O$ anomalies matching $y_{obs,i}$. \mathbf{z}_{obs} are sampled from their prior distributions, decided upon in Section 4.2 when limits were given to the prior variables, as are the $\boldsymbol{\beta}_i$, and used along with the model above to sample from their posterior distributions. For now, the standard deviation is set to just the measurement error $\tau_1 = 0.5\%$. From this we can find the marginal posterior distributions

$$\mathbf{z}_{obs}, \boldsymbol{\beta}_i \mid \mathbf{y}, Z, \mathbf{y}_{obs,i}$$

Finally, we generate predictions of $\mathbf{y}_{pred} | \mathbf{y}, Z, \mathbf{y}_{obs}$, what the $\delta^{18}O$ anomalies should be at each ice core site given the sampled posterior distributions of \mathbf{z}_{obs} , and $\boldsymbol{\beta}_i$. These are found with samples from MCMC,

$$y_{pred,i} = \mathbf{z}'_{obs} \boldsymbol{\beta}_i.$$

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.

6.4.1 Comparing model output to linear regression

Before we incorporate the GP emulator or EAIS sites, we check our model by comparing the output with a linear model of the data. This is done by building a linear model

$$\mathbf{y} = \mathbf{Z}'\mathbf{b} + \epsilon$$

with some linear coefficients, \mathbf{b} and residual error ϵ . We can then compare the values of these coefficients to the posterior distributions of β_i to see if our model is producing sensible values. The coefficients of the linear model are

$$\hat{\mathbf{b}} = \begin{bmatrix} -8.06 & -1.02 & -6.48 & -2.69 & -1.66 & 2.04 \\ -3.45 & 4.85 & 1.81 & -3.26 & 0.05 & -0.45 \\ -7.51 & -2.71 & -8.77 & 0.89 & 0.00 & -1.33 \\ -6.16 & 1.42 & -1.56 & -0.86 & -5.61 & 0.79 \end{bmatrix}. \quad (6.1)$$

We compare this to samples of the posterior mean of β

$$\beta = \begin{bmatrix} -8.04 & -1.00 & -6.46 & -2.66 & -1.66 & 2.03 \\ -3.45 & 4.81 & 1.82 & -3.22 & 0.03 & -0.44 \\ -7.52 & -2.70 & -8.72 & 0.87 & 0.01 & -1.32 \\ -6.17 & 1.41 & -1.54 & -0.87 & -5.59 & 0.79 \end{bmatrix} \quad (6.2)$$

and see that they closely match. This suggests that our model is sampling correctly.

6.5 Incorporating the Gaussian process emulator

We now incorporate the Gaussian process emulator into our model. Whilst in Chapter 5 we discussed finding values of our hyperparameters using the maximum likelihood method in the Python package GPy. Here, we describe in detail how this is done in a Bayesian sense using MCMC.

As mentioned in 6.1.1, we use the Stan language to sample from our posterior distribution using MCMC. Stan is an open-source software built to code probability models, make predictions and use posterior analysis to evaluate results, performing MCMC with the use of the No-U-Turn Sampler (Hoffman and Gelman, 2011). Here we discuss the form our Gaussian process emulator takes in Stan, using some simple examples for illustrative purposes.

The full model for a Gaussian process with normal outcome, $\mathbf{y} \in \mathbb{R}^N$, with inputs $Z \in \mathbb{R}^{N \times D}$ for finite N and D , is

$$\rho, \alpha, \sigma_1 \in \mathbb{R},$$

$$\rho \sim \text{Inv_Gamma}(5, 5),$$

$$\alpha \sim \mathcal{N}(0, 1),$$

$$\sigma_1 \sim \text{Inv_Gamma}(1, 1),$$

$$K(Z|\alpha, \rho)_{i,j} = \alpha^2 \exp \left\{ \frac{-1}{2\rho^2} \sum_{d=1}^D (z_{i,d} - z_{j,d})^2 \right\},$$

$$\mathbf{f} \sim \text{MVN}(m(Z), K(Z|\alpha, \rho)),$$

$$\mathbf{y} \sim \mathcal{N}(\mathbf{f}, \sigma_1^2),$$

where ρ is the lengthscale parameter and α is the variance parameter, σ_1 is the variance of the emulator and K is the kernel described in 5.3. f is a function $y_n = f(z_n)$ with a multivariate Gaussian distribution.

6.5.1 1-D example

We use a 1-D example to demonstrate how the model described in 6.4 works. Using the function

$$f(\mathbf{z}) = \mathbf{z} + \sin(\mathbf{z}) + \mathcal{N}(0, 1),$$

plotted in Figure 6.4, we generate data

$$\mathbf{z}_1, \mathbf{y}_1 \in \mathbb{R}^{10}$$

$$\mathbf{z}_1 \sim U[0, 10], \quad \sigma^2 \sim \mathcal{N}(0, 1),$$

$$\mathbf{y}_1 = \mathbf{z}_1 + \sin(\mathbf{z}_1) + \sigma^2, \tag{2}$$

where σ is a noise term. We also sample test data, \mathbf{z}_{obs} , taking values between 0 and 10, as data for which we want our model to simulate corresponding \mathbf{y}_{obs} .

We combine \mathbf{z}_1 and \mathbf{z}_{obs} in to one input vector, \mathbf{z} . We create a covariance function with a squared exponential kernel for \mathbf{z} , with $\alpha \sim \mathcal{N}(0, 1)$ and $\rho \sim \text{Inv_Gamma}(5, 5)$




Figure 6.4: The function we are attempting to sample from with the example in Section 6.5.1.

as hyperparameters. To be computationally efficient, we find the Cholesky decomposition of $K(\mathbf{z}|\alpha, \rho)$. This is the lower triangular matrix L_K such that $L_K L_K^T = K$. Taking a standard normal variate

$$\eta \sim \mathcal{N}(0, 1), \quad \eta \in \mathbb{R}^N,$$

where N is the length of vector \mathbf{z} . If we set

$$f = L_K \eta$$

then

$$f \sim \mathcal{N}(0, L_K L_K^T)$$

to find test points

$$\mathbf{y}_1 \sim \mathcal{N}(f[1 : N_1], \sigma^2),$$

where N_1 is the length of \mathbf{z}_1 , and generate training points

$$\mathbf{y}_{obs} = \mathcal{N}(f[(N_1 + 1) : (N_1 + N_2)], \sigma^2)$$

where N_2 is the length of \mathbf{z}_{obs} .

We plot the generated quantities, \mathbf{y}_{obs} in Figure 6.5, with the data points $(\mathbf{z}_1, \mathbf{y}_1)$ plotted over. Our model seems to have been sampled correctly; we can see that \mathbf{z}_1 and \mathbf{y}_1 constrain the variance of the model; areas of the model space with no input or output data have more uncertainty than areas with clusters of data.




Figure 6.5: Generated quantities from our 1D example described in Section 6.5.1. Samples from the model are plotted as dashed lines, with data points (z_1, y_1) represented as white dots.

6.5.2 1-D inverse example

We now run the same model, this time generating z_{obs} values for some known observation y_{obs} . This is what we will do in our full model, as we have ice core observations and want to sample prior input values that would create those isotope values.

We create a parameter z_{obs} , given a Gaussian distribution. We calculate the covariance function, $K(\mathbf{z}|\alpha, \rho)$, and find the Cholesky decomposition L_K . We sample from the model with \mathbf{y}_1 and y_{obs} combined

$$\mathbf{y} \sim \mathcal{N}(f, \sigma),$$

and calculate

$$\mathbf{y}_{pred} = \mathcal{N}(f[N_1 + 1, \sigma]),$$

values of \mathbf{y} given the estimates of z_{obs} our model has produced. Figure 6.6 illustrates what we are attempting to do here. The black curve shows the function we are attempting to sample from. The red horizontal line is a possible observation, y_{obs} , with the corresponding z_{obs} shown as a vertical red line. The posterior density of z_{obs} is plotted as a blue density curve.

We plot samples of \mathbf{y} that our model generated in Figure 6.7, with points (z_1, \mathbf{y}_1) included. We then plot \mathbf{y}_{pred} as horizontal lines with y_{obs} in red, and the corresponding z_{obs} estimates as vertical lines. We can see the model is quite accurately inferring the values z_{obs} could take that would give the value y_{obs} .




Figure 6.6: The function we are attempting to sample from with our model in Section 6.5.2, with a possible y_{obs} represented as a horizontal red line, the corresponding z_{obs} as a vertical red line, and the density of z_{obs} shown as a blue density plot.




Figure 6.7: Model output from our 1D inverse example in Section 6.5.2. Samples of the model Equation 2 in Section 6.5.1 are plotted as grey lines, with data (z_1, y_1) plotted as points. We plot samples of y_{pred} as horizontal black lines, with our observation y_{obs} represented by the red line. Samples of z_{obs} corresponding to y_{pred} are represented by the vertical black lines.

6.5.3 Automatic Relevance Determination

As we have multivariate inputs $z_i \in \mathbb{R}^D$, we fit a scale parameter ρ_d for each dimension d to generalise the covariance function

$$K(z|\alpha, \rho, \sigma)_{ij} = \alpha^2 \exp\left\{-\frac{1}{2} \sum_{d=1}^D \frac{1}{\rho_d^2} (z_{i,d} - z_{j,d})^2\right\}.$$

This allows us to model the lengthscale parameters for each of the five dimensions in our model and see the relevance of the corresponding model parameters through the values of the ρ_d (Sandhu et al., 2017), as well as creating a more flexible model that will allow us to predict better. We consider the prior and posterior distributions of ρ_d when analysing our hyperparameters in 6.6.2 to see if any of our prior variables are more relevant to the model than others.

6.6 Applying to the ice core data

Our model takes a similar multivariate approach to the examples above, in Sections 6.5.1 and 6.5.2. z_1 is our prior data, Z , a $N_1 = 47 \times 5$ matrix of prior variable values. y_1 is the corresponding $47 \times N_2$ matrix of HadCM3 $\delta^{18}O$ anomalies, where N_2 is the number of ice core sites we are modelling. y_{obs} is a vector of length N_2 containing the $\delta^{18}O$ anomalies from ice cores. We aim to find a vector of length D , z_{obs} , that with our model would create values close to y_{obs} . We then generate y_{pred} , the $\delta^{18}O$ anomalies created by inputting z_{obs} into our model.

We add in a fixed standard deviation term for \mathbf{y}_{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 an error for 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, 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.

Including the prior model in our emulator, we have for each of the N_2 ice core sites

$$\begin{aligned}
Z^* &= \begin{pmatrix} Z \\ \mathbf{z}_{obs} \end{pmatrix} \in \mathbb{R}^{48 \times 5} & \mathbf{y}_i &= \begin{pmatrix} \mathbf{y}_1 \\ y_{obs} \end{pmatrix} \in \mathbb{R}^{48} \\
\boldsymbol{\beta}_i &\in \mathbb{R}^6 & \boldsymbol{\sigma}_{1,i}^2 &\in \mathbb{R}^+ \\
\boldsymbol{\beta}_i &\sim \mathcal{N}(0, 100) & \boldsymbol{\sigma}_{1,i}^2 &\sim \text{Inv_Gamma}(1, 1) \\
\alpha^2 &\sim \mathcal{N}(0, 1) & \rho_d &\sim \text{Inv_Gamma}(5, 5) \\
f &\sim \text{MVN}(0, K(Z^*|\alpha^2, \rho_d)) & d &= 1, \dots, 5 \\
\mathbf{y}_i | \cdot &\sim \mathcal{N}(Z' \boldsymbol{\beta}_i + f, \sigma_{1,i}^2 + \tau_1^2 + \tau_2^2) & i &= 1, \dots, N_2.
\end{aligned} \tag{3}$$

We then generate

$$y_{pred,i} = \mathcal{N}(z'_{obs}\boldsymbol{\beta}_i + f[N_1 + 1], \sigma_{1,i}^2 + \tau_1^2 + \tau_2^2), \quad (4)$$

\mathbf{z}_{obs} having the same prior distributions as in Section 6.4, using the posterior sample \mathbf{z}_{obs} to create $\delta^{18}O$ values close to y_{obs} .

6.6.1 Modelling combined and separate ice sheets

Due to the discrepancies we saw between the HadCM3 output and the ice core observations in the EAIS as described in Chapter 5, we focused on the WAIS initially before incorporating the EAIS into the model. We can then compare how well we model the ice sheets, together and separately. This allows us to see whether either ice sheet is influencing results more than the other, and how modelling all ice core sites together affects the posterior distributions of both \mathbf{z}_{obs} and \mathbf{y}_{pred} . We run three versions of our model: one for the four west ice core sites, one for the six EAIS ice core sites, and one with all ten ice core sites together. We compare how modelling the ice cores in regions and as a group affects the prediction of $\delta^{18}O$ anomalies by plotting densities of y_{pred} for all three models, with the ice core observation superimposed. Figure 6.8 shows how well the model estimates $\delta^{18}O$ anomalies for models for the WAIS, EAIS and all ice core sites. We can see that when the sites are separated by ice sheet, the model generates $\delta^{18}O$ anomalies well. When all sites are combined, it struggles with some sites both in the WAIS and EAIS.

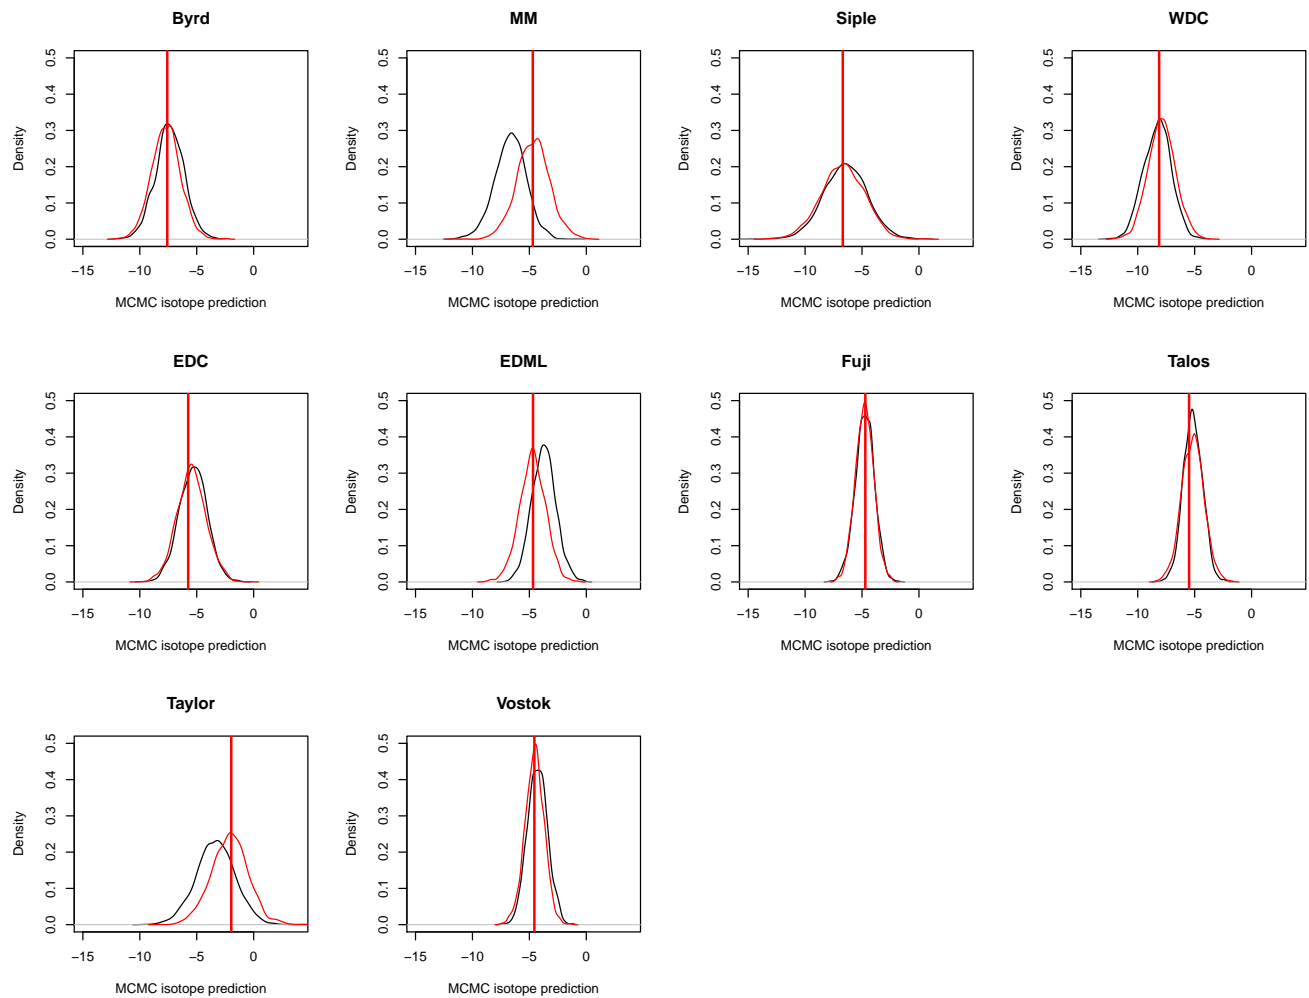


Figure 6.8: Comparing the generated densities of y_{pred} when modelling all ice core sites vs modelling them in two groups. The black lines are the densities of y_{pred} when modelling all ice core sites. The red densities are from modelling the EAIS and WAIS ice core sites separately. The vertical red lines are the observed $\delta^{18}O$ anomalies. We can see that there is a discrepancy between our model and the ice core observations when modelling the WAIS and EAIS sites together. The densities of the $\delta^{18}O$ values at the Mount Moulton, EDML and Taylor sites are all peaking further away from the observations.

6.6.2 Checking hyperparameters

We check the nugget term of the emulator, σ_1 . If this was very large then the emulator is not reducing uncertainty for the posterior distribution of the ice sheet shape, and is not useful in our model. As we can see in Figure 6.9, which shows the densities of the emulator nugget terms, the emulator variance at each ice core site is peaking at values less than one, which is reassuringly small.

For our model with a lengthscale parameter for each dimension, we check the prior and posterior distributions of them to see if anything is learned. Plots of this are shown in Figure 6.10. We can see that the posterior closely follows the prior distribution $\text{Inv_Gamma}(5,5)$, suggesting little has been learnt about the lengthscale parameters from this process. The similarity in the values of all five lengthscale parameters also suggests that the five prior variables are all equally relevant to the model.

We now see how much is learned about the five prior variables through this process. We plot histograms of the prior distribution of \mathbf{z}_{obs} with their posteriors superimposed to compare how much we have learned. These are plotted in Figure 6.11. When the ice core sites are modelled separately, there is some divergence between the prior and posterior distributions of \mathbf{z}_{obs} . When the sites are combined into one model, there is a much clearer difference between the prior and posterior distributions. We therefore choose to continue our analysis using ice core sites from both ice core sheets simultaneously. What this means in a geographical sense is discussed in Chapter 7, where all three models are again assessed.

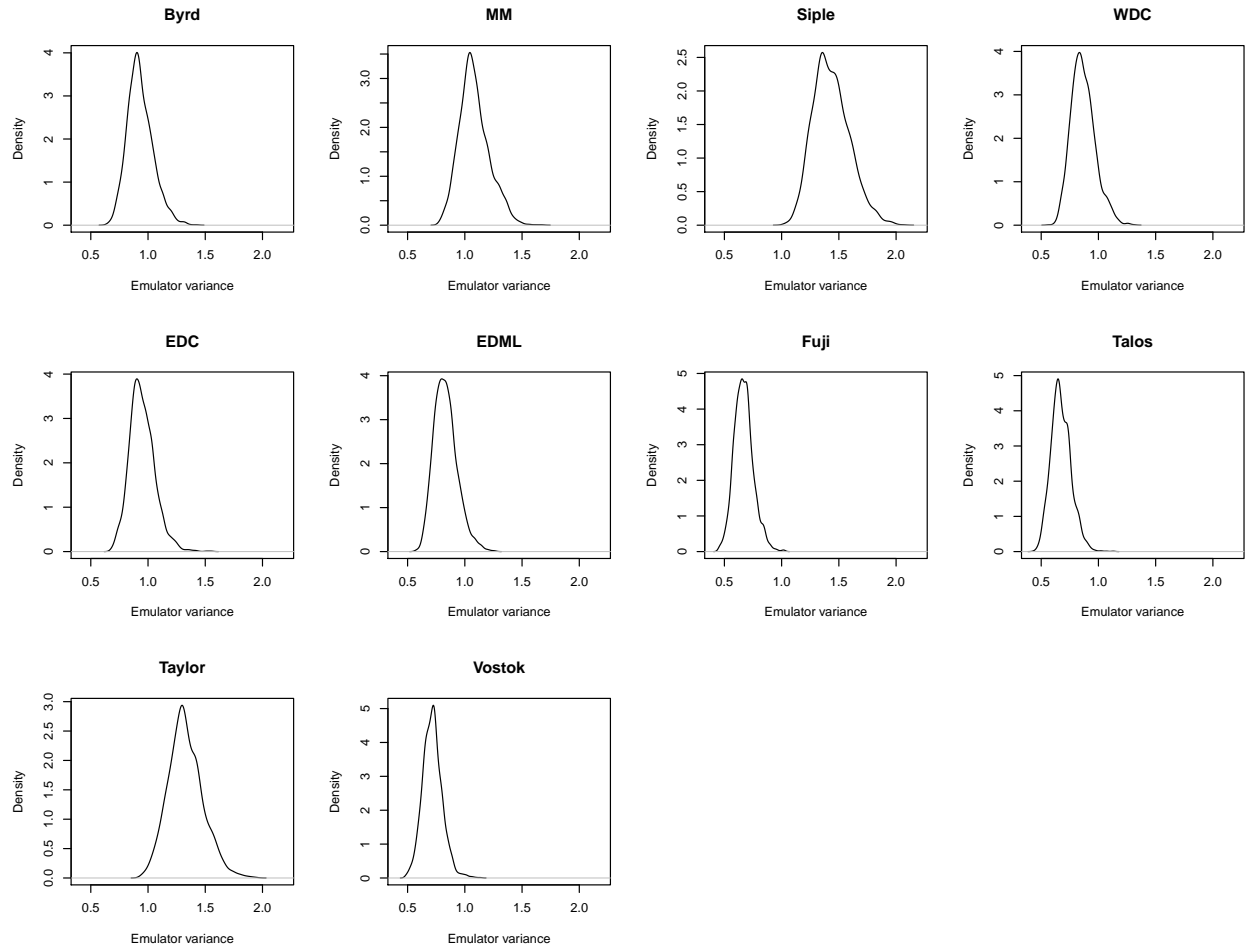


Figure 6.9: Emulator nugget term, σ_1 , at each of the ice core sites. We can see the variance at the ice core sites are small, mostly peaking at values less than one.

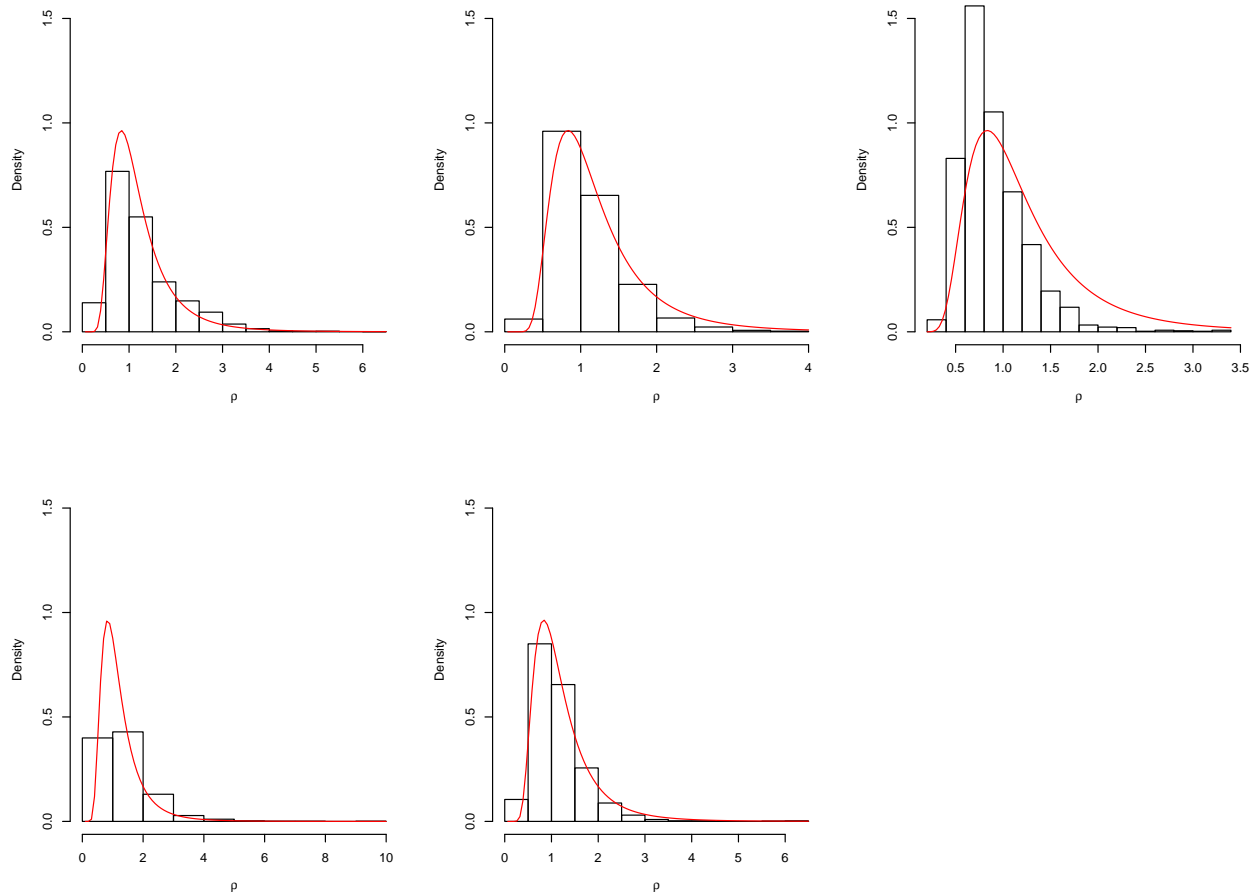


Figure 6.10: Histograms of the posterior distributions of the length scale parameters, ρ_d $d = 1, \dots, 5$, with their prior distributions superimposed in red. We can see that little is learned about the lengthscales, and none of the prior variables appears to have a significantly larger value, suggesting they are all equally relevant to the model.

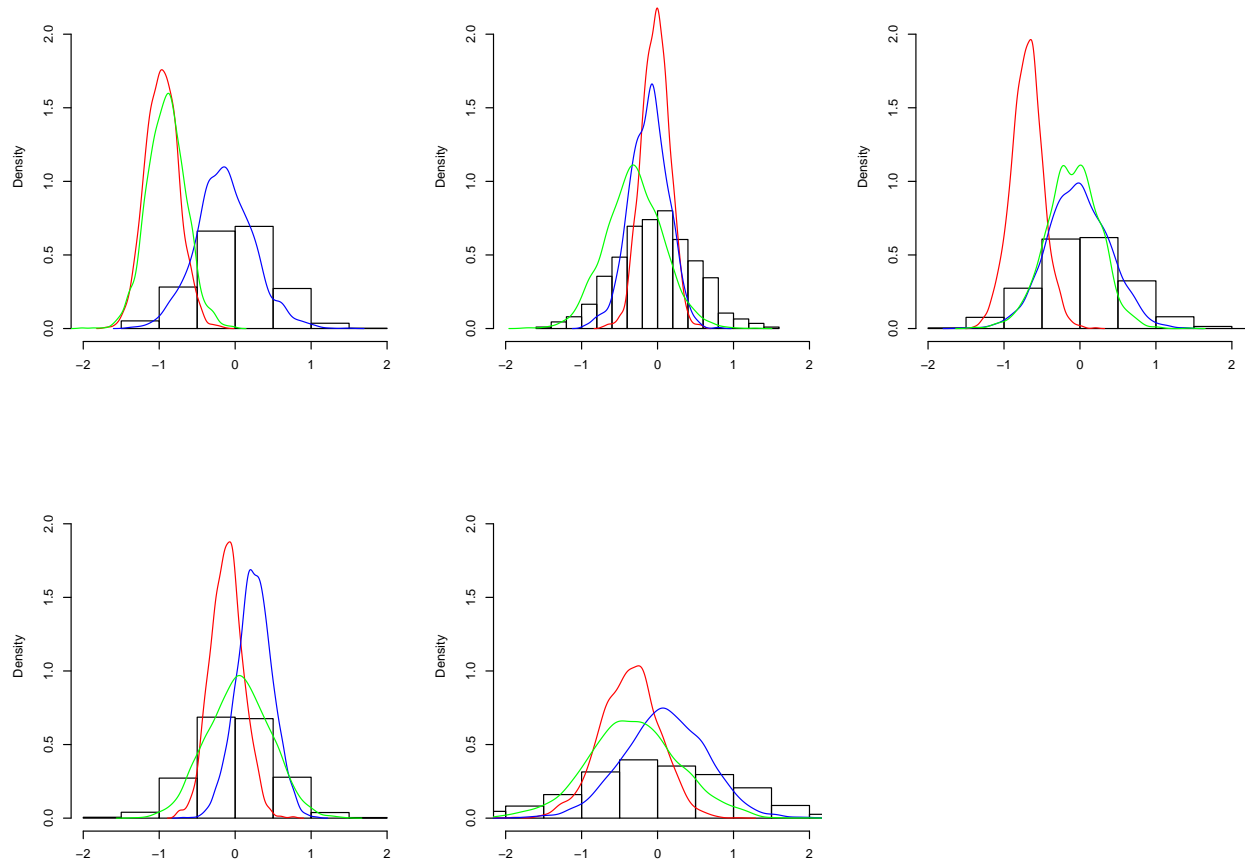


Figure 6.11: Histograms of prior distributions of z_{obs} , with the posterior distributions overlaid when modelling EAIS (green), WAIS (blue) and all (red) ice core sites. The red density curves show that when modelling all ice core sites at once there is a clearer difference between the prior and posterior distributions of z_{obs} - more so than when modelling the ice sheets separately, particularly for the WAIS.

6.7 Analysing model output

We make a pairs plot to look at the correlation between the posterior samples of the five prior variables. This can be seen in Figure 6.12. We add in the points that reconstruct the four most ‘important’ shapes (ICE5G (Peltier, 2004), ICE6G (Argus et al., 2014), W12A (Whitehouse et al., 2012) and GLAC-1D (Briggs et al., 2014)), discussed in Section 3.3, to see where they lie in the posterior sample. In some cases they are in the centre of the sample, suggesting that our model has created a sample similar to these values. For other variables they are outliers, with ICE5G at the opposite end of the axis to the other three and the posterior sample plotted in between. This suggests our model is sampling values of \mathbf{z}_i in the space in between values that would reconstruct the four ‘important’ shapes.

6.7.1 Sensitivity analysis

We now perform a sensitivity analysis of our model to analyse the source of any uncertainties.

We test our full model by increasing and decreasing the prior variance of the five variables in our prior model. \mathbf{z}_{obs} are the values we wish to simulate which would produce $\delta^{18}O$ anomalies close to the ice core observations through our model. Originally we had

$$\mathbf{z}_{obs} \sim \mathcal{MVN}(\mathbf{0}, \text{diag}(0.5, 0.5, 0.6, 0.5, 1)).$$

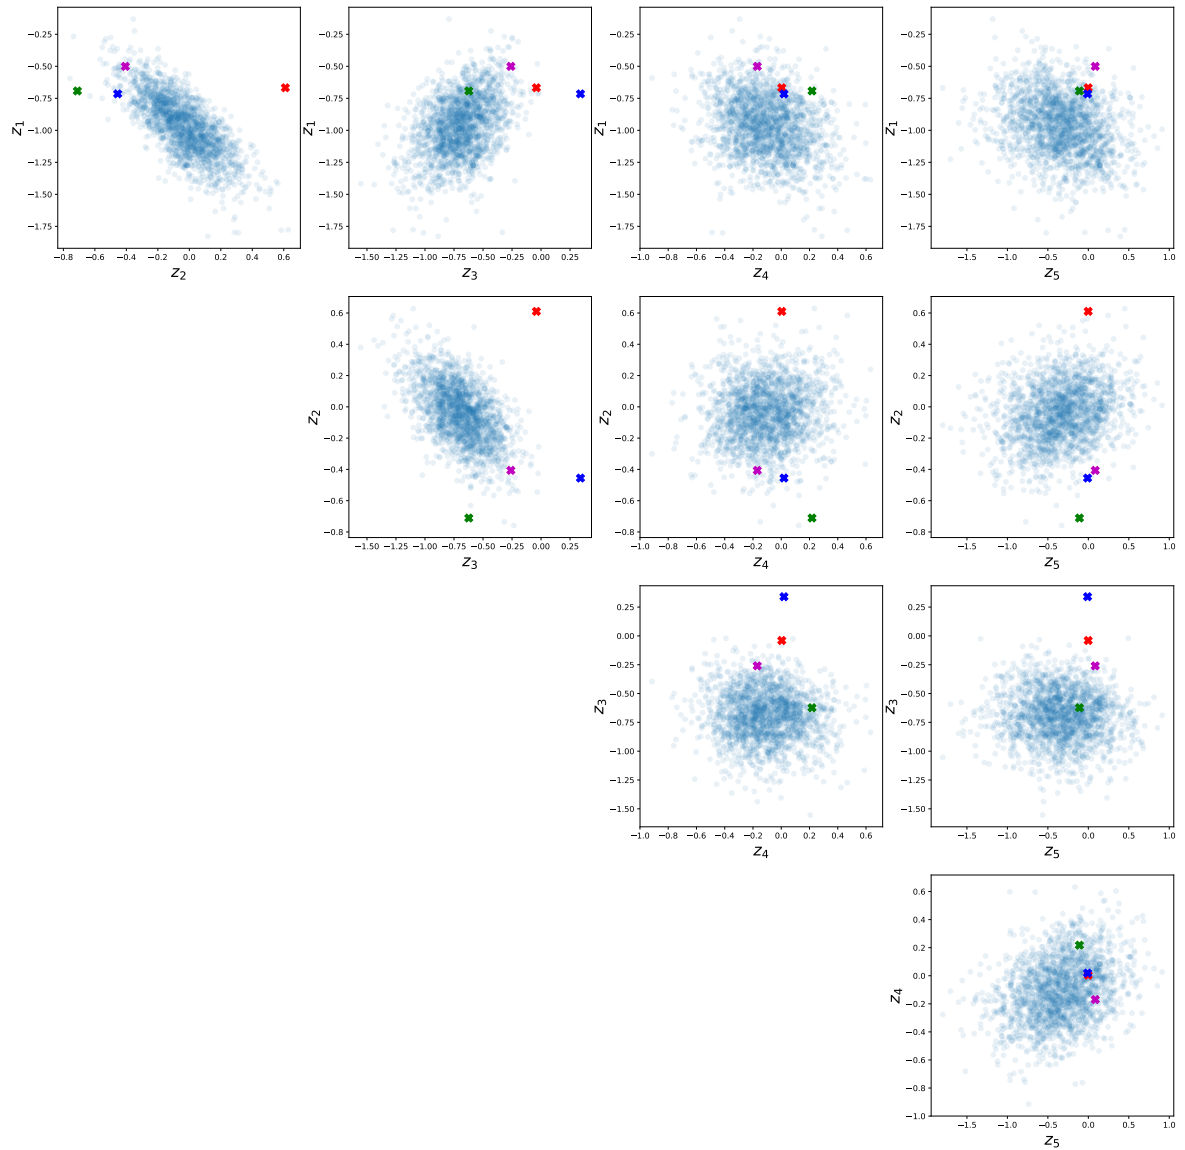


Figure 6.12: Pairs plot of posterior samples with the values for the first four ‘important’ shapes added in. ICE5G is in red, ICE6G is in blue, W12A is in green and GLAC-1D is in magenta. We can see that often they are grouped in the simulated values, or spread out. When the ‘important’ shapes are outliers, the simulated values appear in between.

We firstly shrink the standard deviation to

$$\mathbf{z}_{obs} \sim \mathcal{MVN}(\mathbf{0}, \text{diag}(0.05, 0.05, 0.06, 0.05, 0.1)).$$

We compare our posterior distributions of this version of \mathbf{z}_{obs} and the corresponding \mathbf{y}_{pred} to the distributions of \mathbf{z}_{obs} and \mathbf{y}_{pred} in Section 6.6.1. These are shown in Figures 6.13 and 6.14. The posteriors of \mathbf{z}_{obs} peak at the same points but the spread of the sample is clearly affected by the prior standard deviation. $\delta^{18}\text{O}$ anomalies are generally generated close to the observations, although shrinking \mathbf{z}_{obs} has affected some ice core sites.

We then increase the prior standard deviation of \mathbf{z}_{obs} to

$$\mathbf{z}_{obs} \sim \mathcal{N}(\mathbf{0}, \text{diag}(5, 5, 6, 5, 10)),$$

We compare our posterior distributions of this version of \mathbf{z}_{obs} and the corresponding \mathbf{y}_{pred} . These are in Figures 6.15 and 6.16. The posterior distributions of \mathbf{z}_{obs} are very different to the priors, as well as the previous posteriors. The \mathbf{y}_{pred} distributions however are now being sampled much closer to the observed values.

We vary the size of the fixed standard deviation to see how much this is affecting the predictions of the ice core sites. Figure 6.17 shows the densities of \mathbf{y}_{pred} when we reduce the fixed standard deviation, $\tau = \tau_1 + \tau_2$, where τ_1 is measurement error for the isotopes and τ_2 is spatial error for HadCM3. We tested with $\tau = 1$ and $\tau = 5$, compared to our chosen value of $\tau = 3$. We can see that when the fixed variance is changed, there is little difference in the \mathbf{y}_{pred} densities, suggesting that the fixed

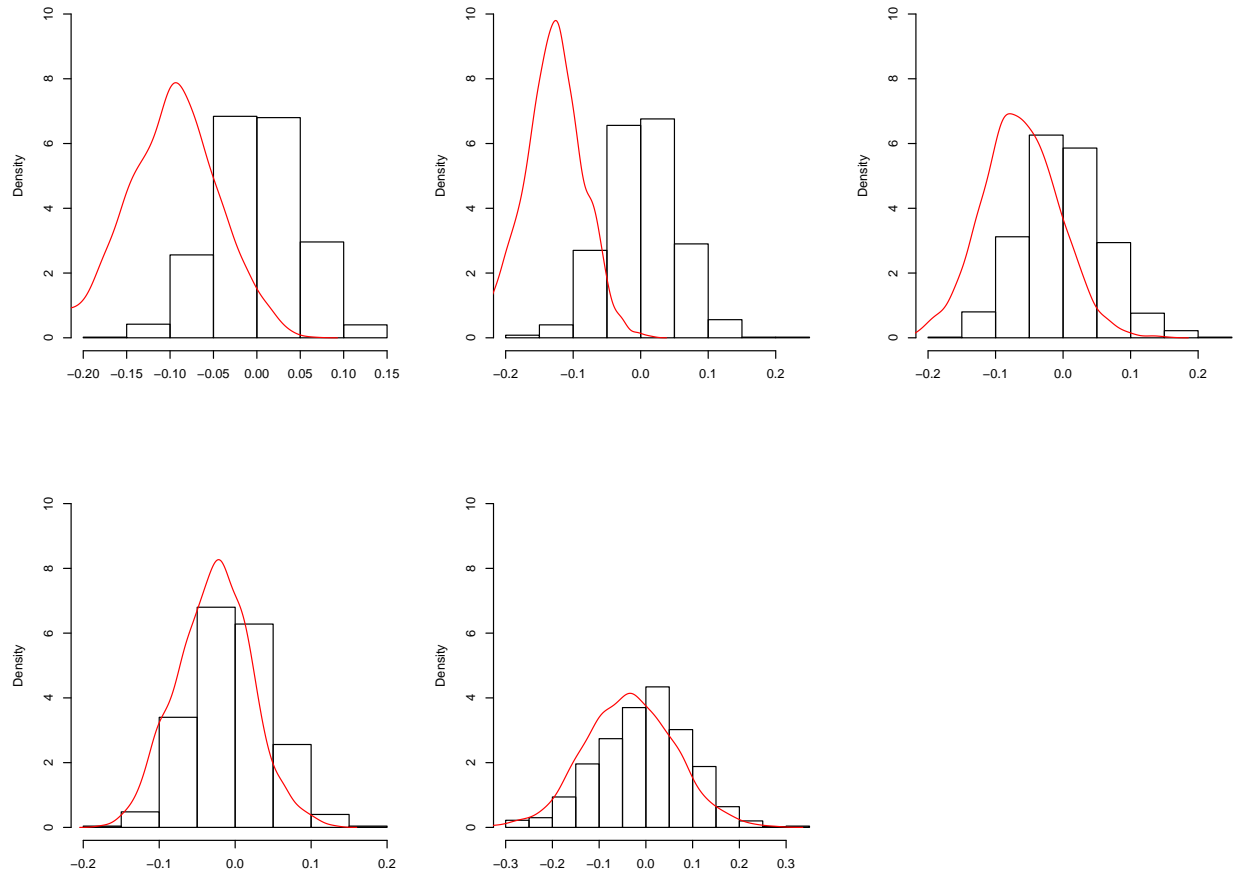


Figure 6.13: Histograms of the prior distributions of z_{obs} when the prior standard deviation is shrunk by a scale of 10, with the density curve of the posterior distributions overlaid in red. The densities of the posteriors are peaking at a similar point to the distributions in Figure 6.11, but the spread is much smaller.

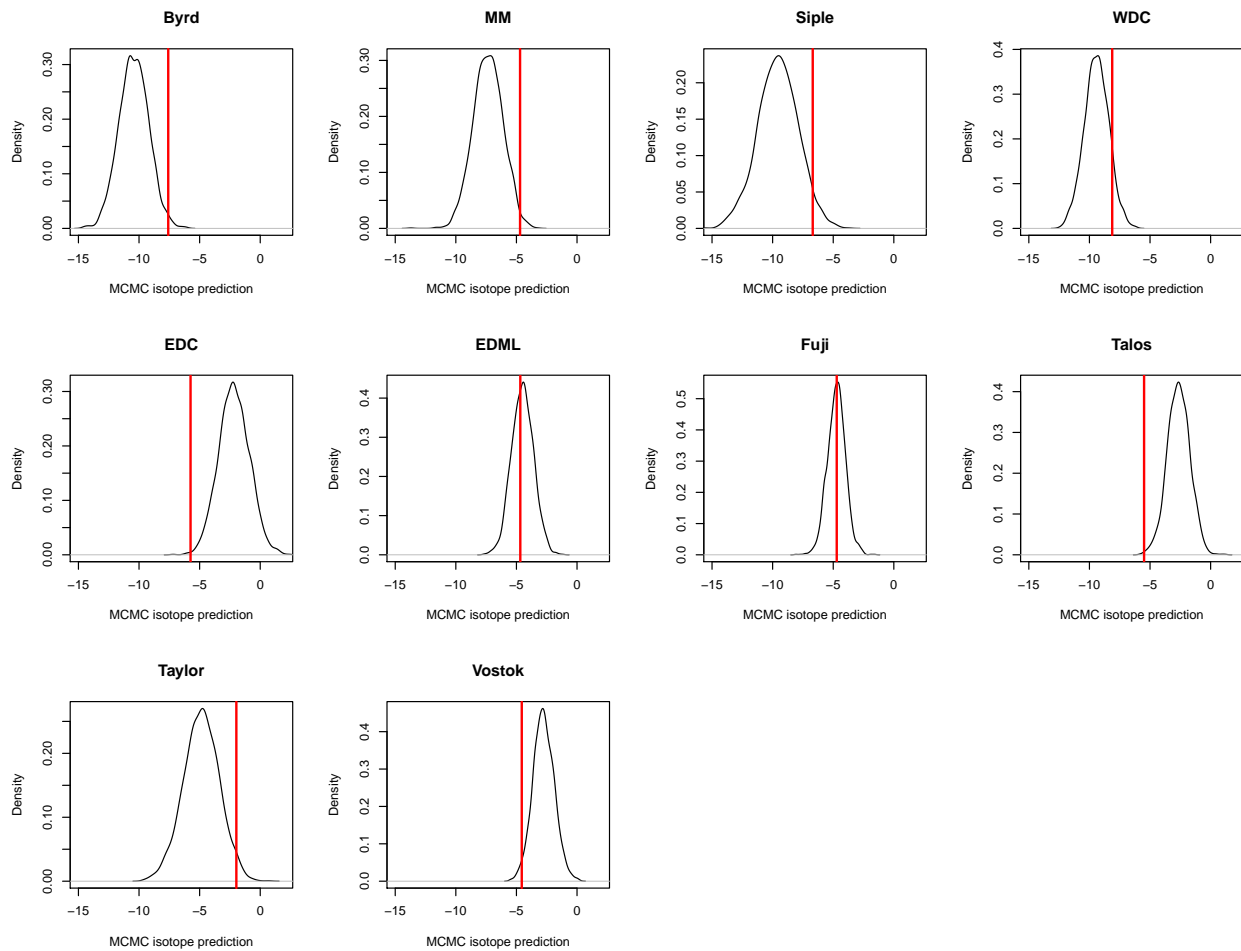


Figure 6.14: The distributions of y_{pred} (black curves), with the ice core observations shown in red, when the prior standard deviation of z_{obs} is shrunk by a scale of 10. Some of the ice core sites are clearly affected by the change, comparing this to Figure 6.8. We can see that the Byrd, EDC, Talos and Vostok sites now have densities of $\delta^{18}O$ values peaking further away from the ice core observations.

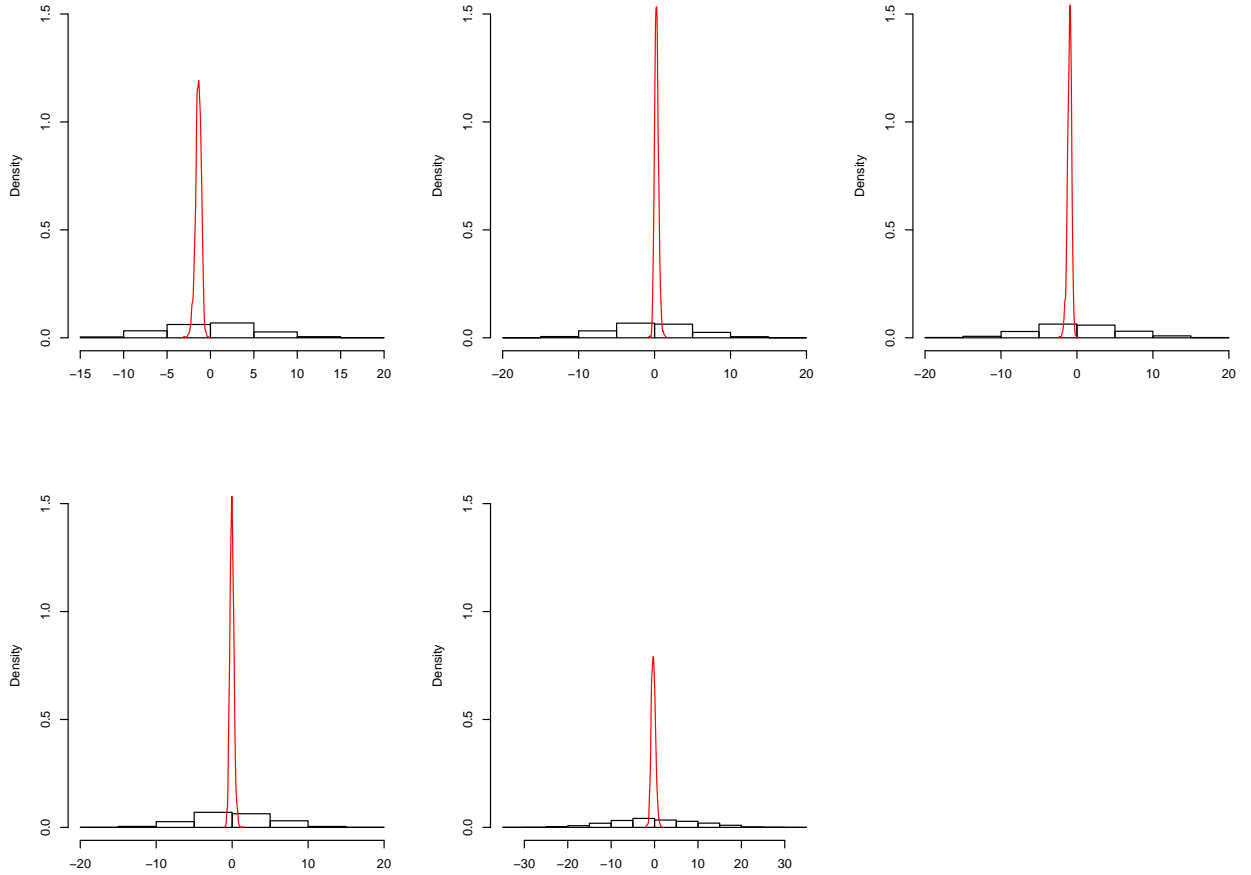


Figure 6.15: Histograms of the prior distributions of z_{obs} when the prior standard deviation is increased by a scale of 10, with the density curve of the posterior distributions overlaid in red. The posterior distributions have been updated a lot from the priors, displaying a much different mean and variance.

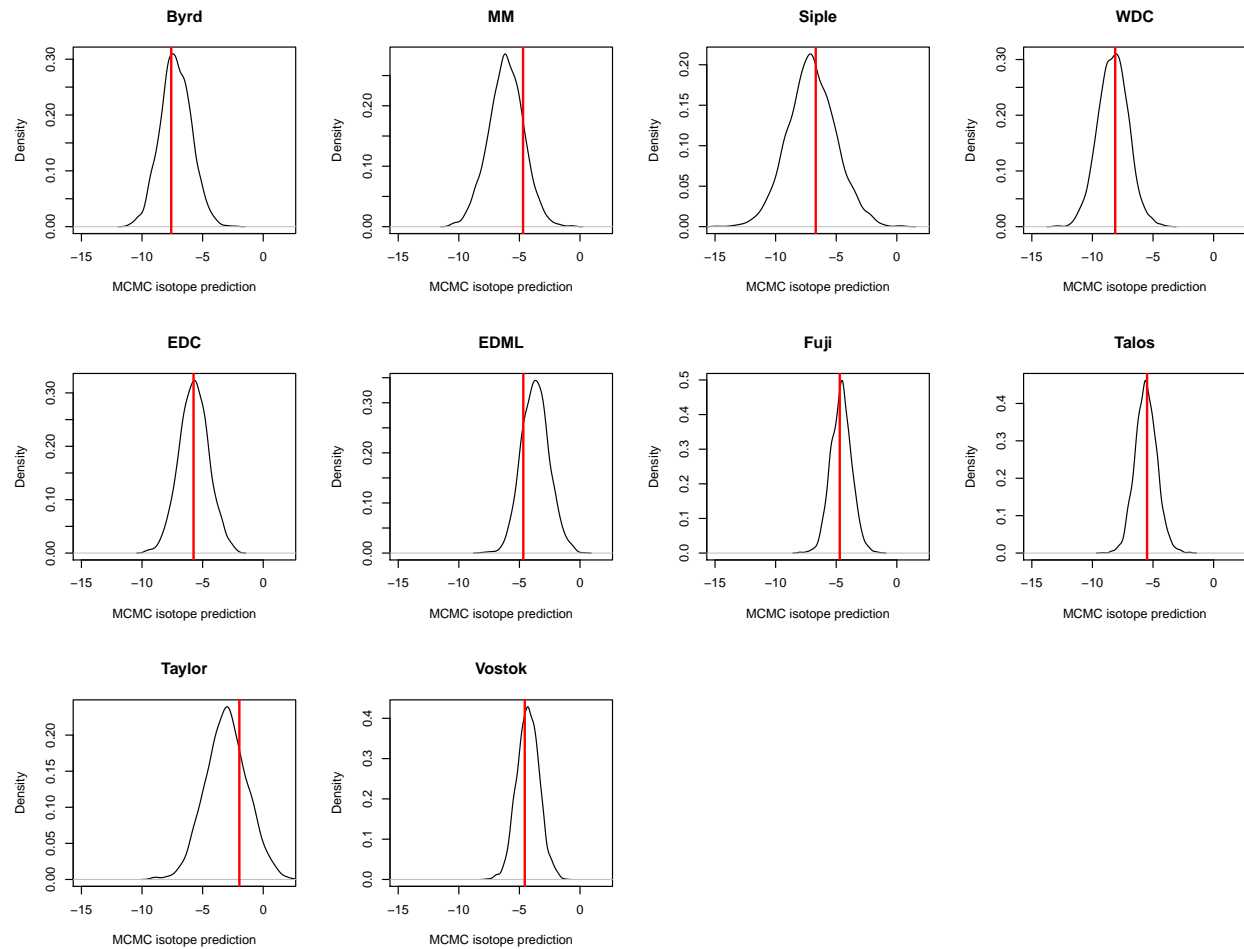


Figure 6.16: The distributions of y_{pred} (black curves), with the ice core observations shown in red, when the standard deviation of z_{obs} is increased by a scale of 10. We can see the Mount Moulton, EDML and Taylor sites all have discrepancies compared to the distributions in Figure 6.8.

errors are not affecting our models' capabilities of generating ice core anomalies close to our observations.

6.8 Testing the model

We test the ability of the model incorporating all ice core sites to recover prior variable values from $\delta^{18}O$ anomalies. We do this in a series of ways: simulating pseudo data and using this as input data, output data and observations, setting one of our HadCM3 simulations as our observations and attempting to recover the input values, and removing the data for an ice core site and seeing how well the model can predict its observations.

6.8.1 Pseudo data

We simulate synthetic data with which to test our model by sampling from the distributions of the z_{obs} s and running these values through the HadCM3 emulators built for the ten ice core sites, providing us with synthetic $\delta^{18}O$ anomalies. We then use this as input for our model, setting one of the synthetic sets of $\delta^{18}O$ anomalies as y_{obs} . We want to test how well the model predicts values of z_{obs} when we know what the 'true' values are. We plot the results in Figures 6.18 and 6.19, and can see that the posterior distributions of z_{obs} have updated a lot from the priors and are close

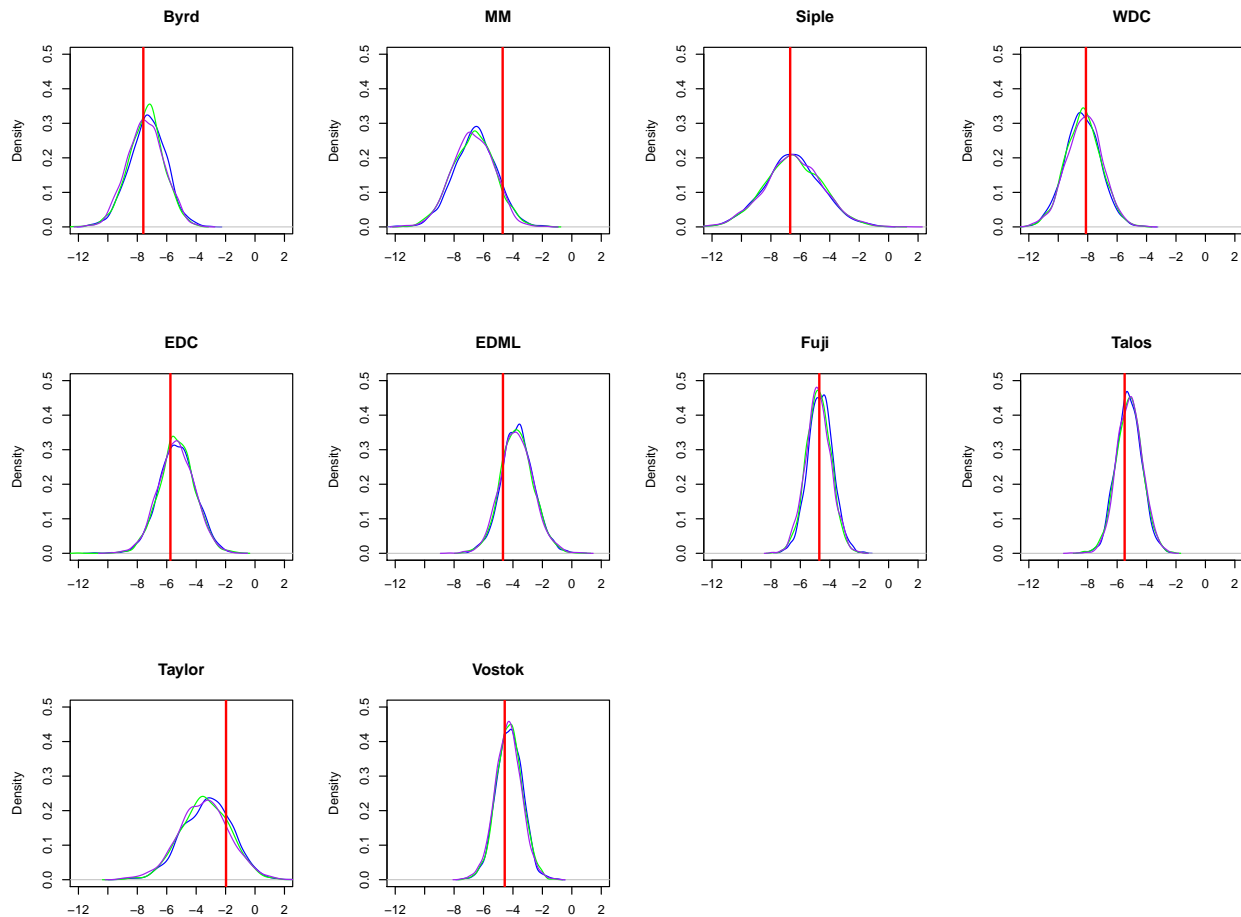


Figure 6.17: The densities of y_{pred} at the ice core sites when the fixed variance is varied. The blue line is when $\tau = 1$, green is the standard $\tau = 3$ and purple is $\tau = 5$. The red vertical lines represent the ice core observations.

to the actual values. The model is also producing generated quantities for \mathbf{y}_{pred} that are close to the synthetic $\delta^{18}O$ anomalies we created. This suggests that our model is capable of recovering prior variable values needed to calculate $\delta^{18}O$ close to known values.

6.8.2 Recovering a HadCM3 simulation

We remove one of the HadCM3 climate simulations, so that we now have $N_1 = 46$. We then set this simulation as our y_{obs} , and generate values of \mathbf{z}_{obs} and compare them to the known prior values. The posterior distributions of \mathbf{z}_{obs} and generated distributions of \mathbf{y}_{pred} can be seen in Figures 6.20 and 6.21. 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 6.20 are all peaking around the \mathbf{z} values used to create the orography used in the missing HadCM3 simulation. Similarly, Figure 6.21 shows that the $\delta^{18}O$ are all peaking around the values from the missing HadCM3 simulation.

We check with a second simulation to see if our model is capable of recovering more simulations. We can see in Figures 6.22 and 6.23 that our model is doing a good job at recovering the ‘true’ values of \mathbf{z}_{obs} and y_{obs} , with the posterior distributions of the variables all peaking close to the observations. We also compare the ice sheet shape that our model recovers to the one we submitted as part of a HadCM3 simulation. Figures 6.24 and 6.25 show the differences between the orographies used in the HadCM3 simulations and the shapes our model recovered. We can see

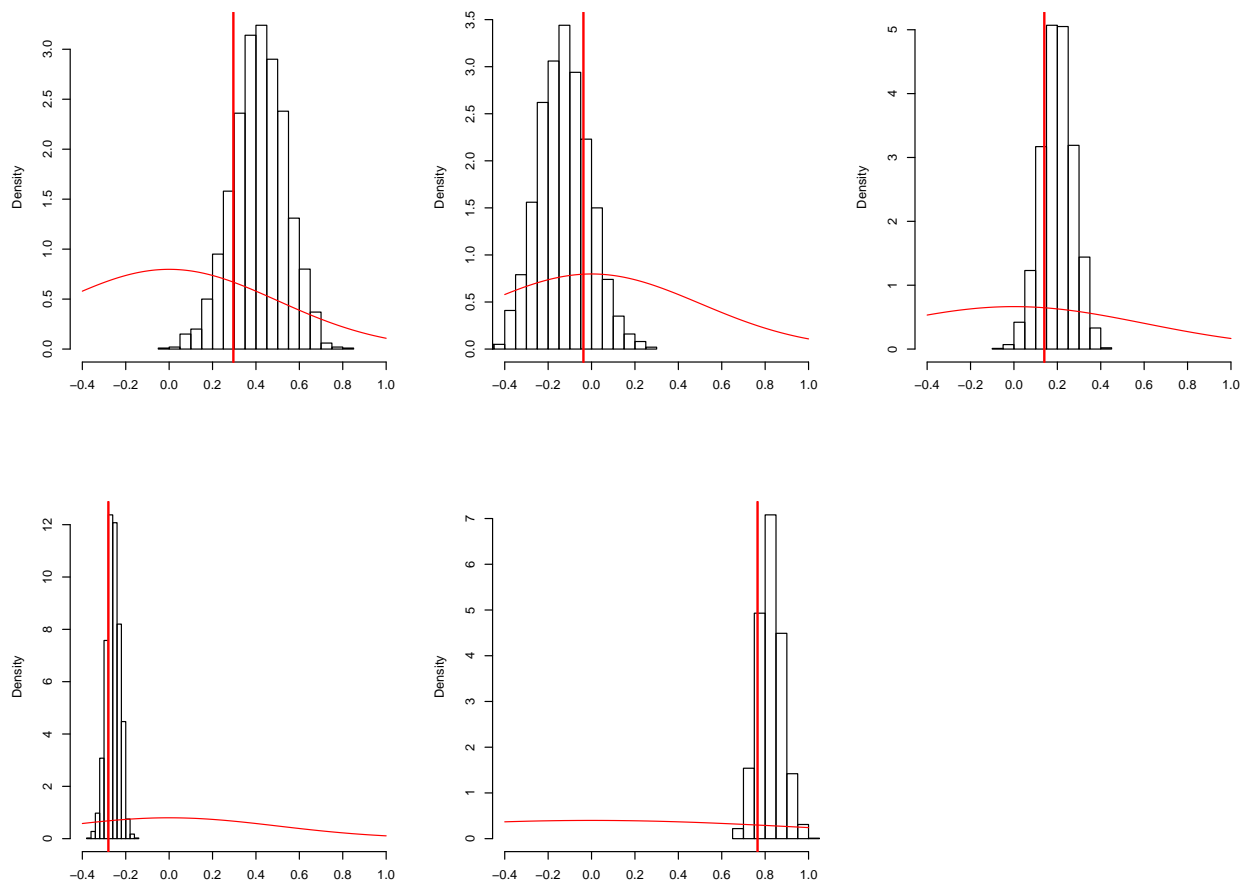


Figure 6.18: Histograms of the posterior distributions of z_{obs} , with density curves of the prior distributions superimposed over them, when using synthetic data. The true values of the synthetic z_i 's are added as red vertical lines.

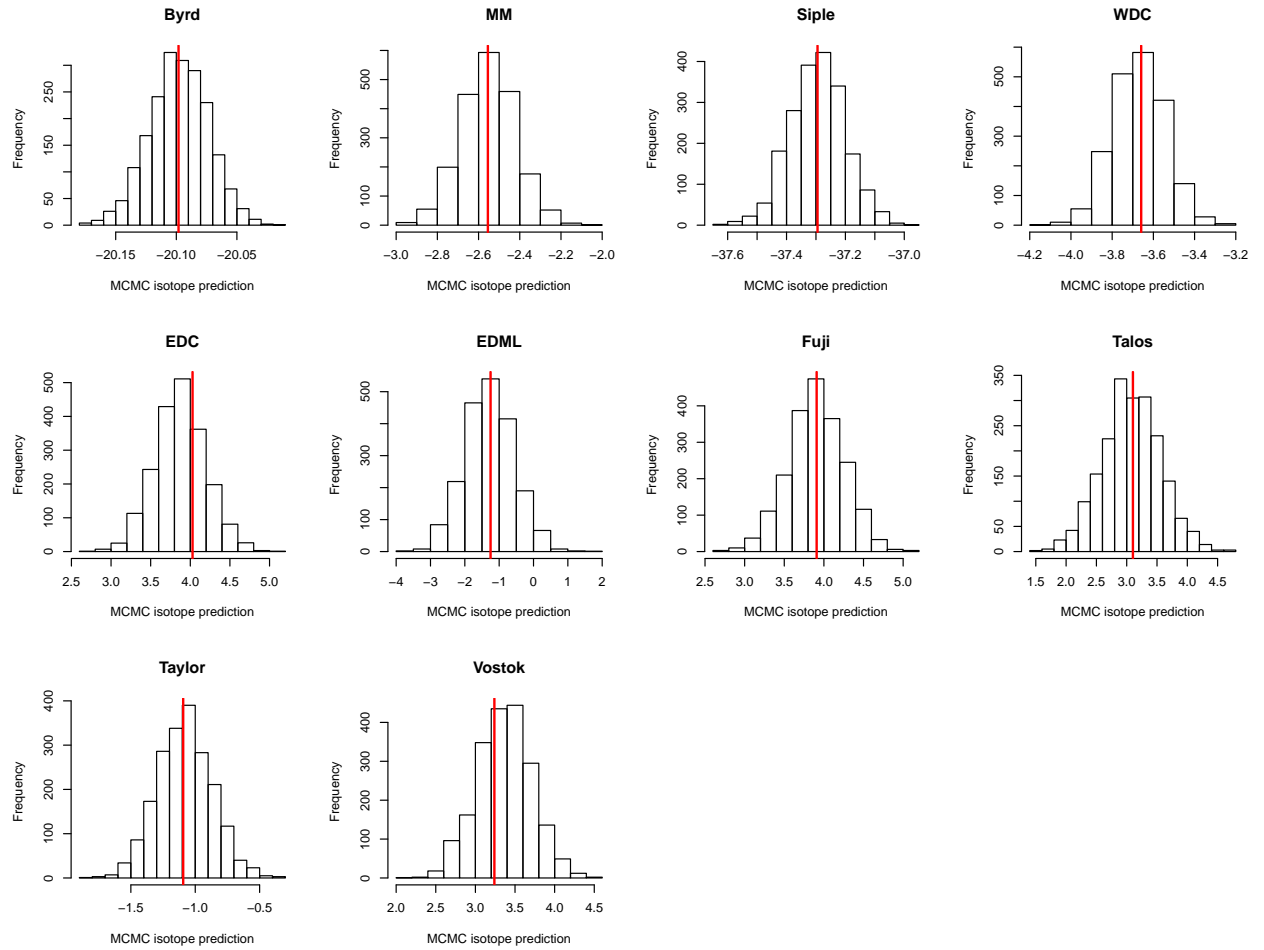


Figure 6.19: The distributions of y_{pred} when using synthetic data, with the synthetic observations added as red vertical lines.

our model has produced an accurate first simulation; the second shape recovered is slightly larger than the original orography in places.

6.8.3 Recovering an ice core site

We further test the model incorporating all ice core sites by trying to recover an ice core site. We wish to see if, by modelling only nine of the sites, we can create $\delta^{18}O$ anomalies close to the ice core observation at the missing ice core site. We do this by removing the observation from Equation 3, so that we are not learning anything about it from the input and output data. We then use Equation 4 to generate $\delta^{18}O$ anomalies for the missing ice core site and compare it to the observed value. A histogram of the missing ice core site, Vostok, is in Figure 6.26; we can see that the distribution of y_{pred} for the missing ice core observation is peaking close to the true value. However, the range of values is extremely large; as the β_i 's were given vague priors with standard deviation of 100, we can see that the model has not updated much from these. Further tests using other ice core sites showed that this applied to sites in both ice sheets; the model struggled to recover an unknown site based on the sampling of the other nine.

6.9 Summary of Chapter

We have calibrated our model and successfully sampled from the posterior distributions of our prior variables. Generated distributions of the ice core observations, as

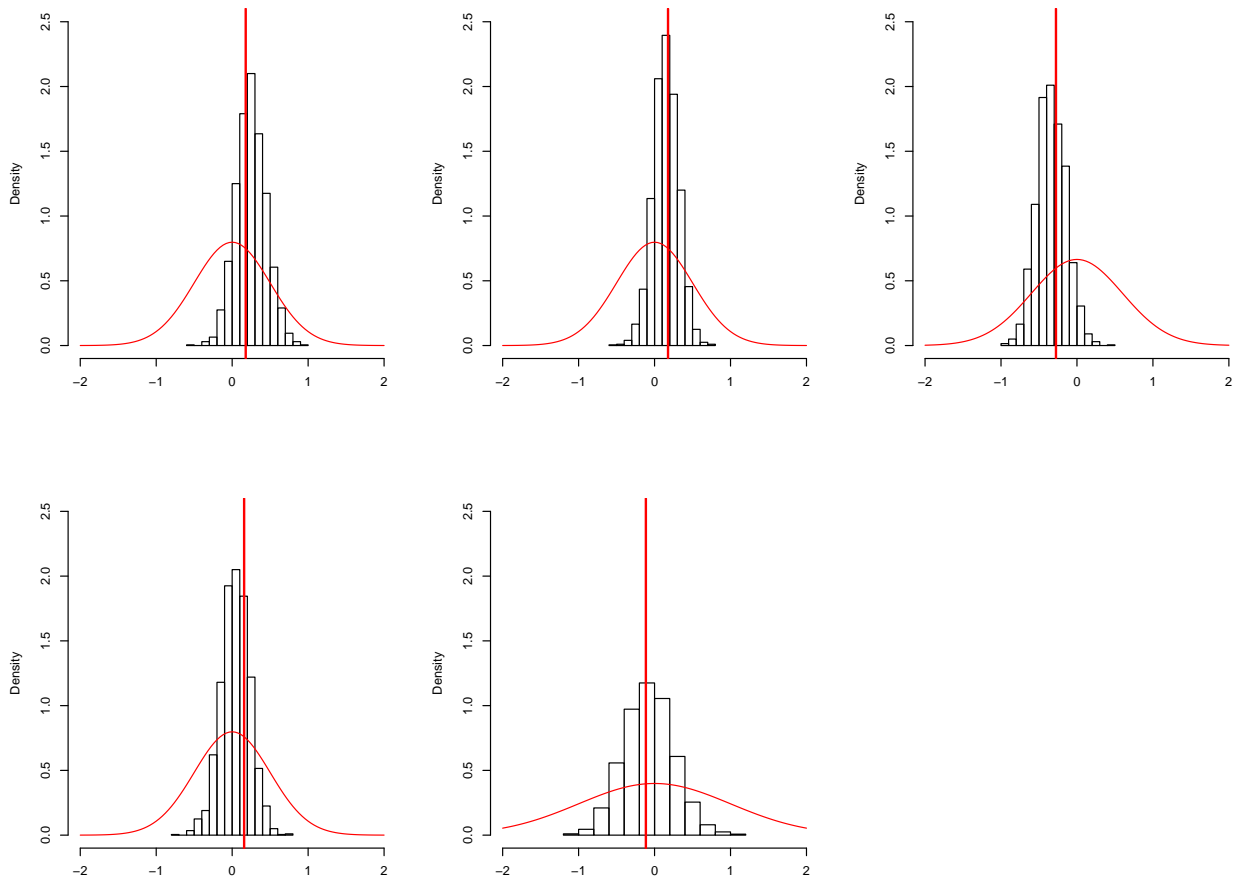


Figure 6.20: Histograms of the posterior distributions of z_{obs} , with density curves of the prior distributions superimposed over them, when trying to recover a HadCM3 simulation. The true values of the synthetic z_i 's are added on as red vertical lines. We can see the posterior distributions are very different to the priors, and are parking close to the observations.

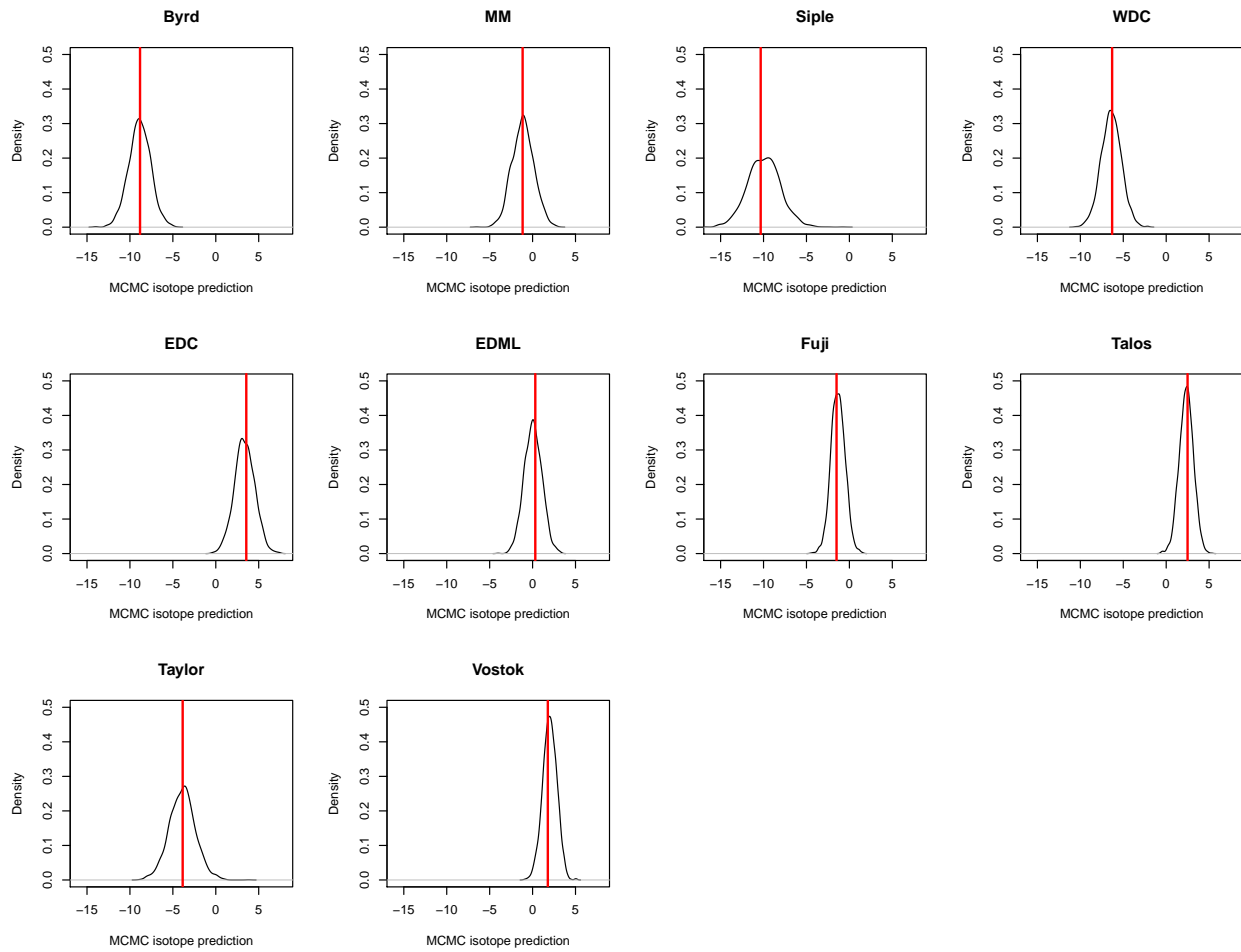


Figure 6.21: The distributions of y_{pred} when recovering a HadCM3 simulation with the synthetic observations added on as red vertical lines.

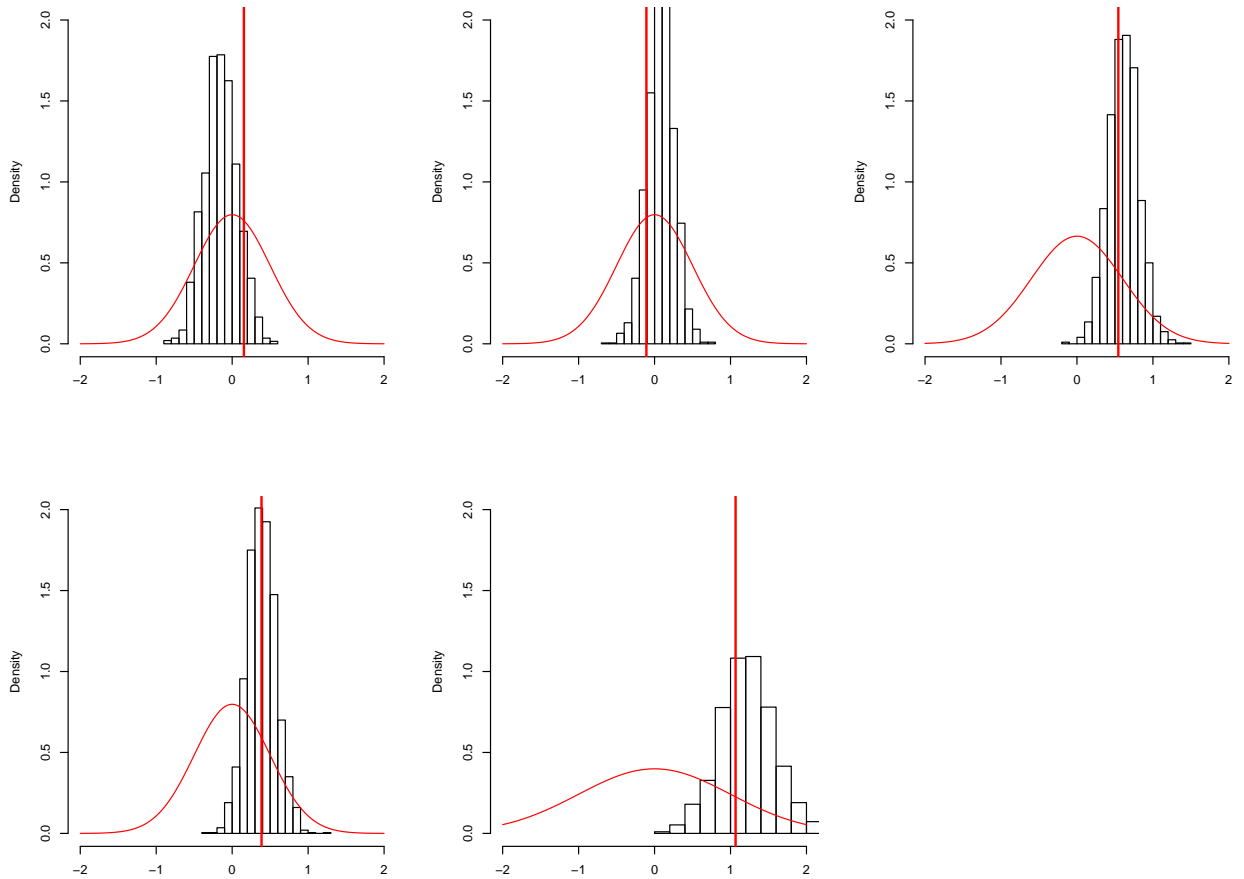


Figure 6.22: Histograms of the posterior distributions of z_{obs} , with density curves of the prior distributions superimposed over them, when trying to recover a second HadCM3 simulation. The true values of the synthetic z_i 's are added on as red vertical lines.

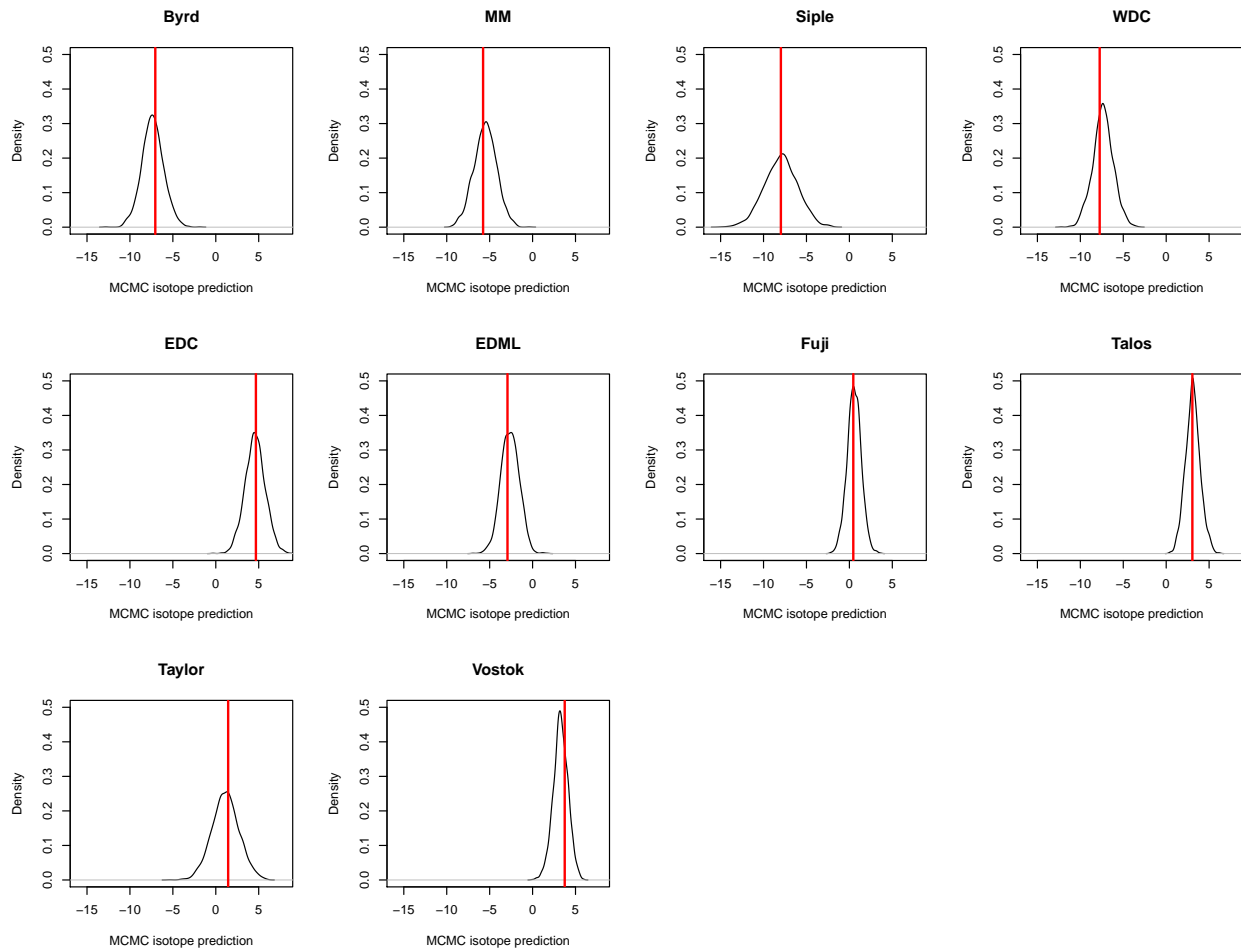


Figure 6.23: The distributions of y_{pred} when recovering a second HadCM3 simulation with the synthetic observations added on as red vertical lines.




Figure 6.24: The difference between an orography used in a HadCM3 simulation and the shape our model recovered. We can see the model has created a shape slightly smaller than the orography we were attempting to recover.




Figure 6.25: The difference between a second orography used in HadCM3 and the shape our model recovered. We can see that the model has been less successful this time, creating an ice sheet shape thicker than the original one we attempted to recover.




Figure 6.26: The histogram of generated values of the recovered ice core site. The ice core observation is plotted over as a red vertical line. The distribution of y_{pred} does not appear to have been adjusted much from the prior distributions of the hyperparameters, suggesting our model struggles to recover an ice core site from the other nine.

well as the sensitivity analysis and model tests we have run, suggest that our model is working well, with the distributions of y_{pred} centred around the observations, and inflated variances little affecting results.

In the next chapter, we discuss the geographical interpretations of our results, and analyse the ice sheet orographies our model has produced and how they differ from the reconstructions used to construct our prior model.

Chapter 7

Geographical Interpretations

In this chapter, we look into the geographical interpretations of the results discussed in Chapter 6. The sampled posterior distributions of the prior variables are used in our prior ice sheet shape model to create sets of LGM AIS orographies, we compare these shapes as well as the variance in each set. We consider how different our estimates are to the ice sheet reconstructions used to build our prior model.

7.1 Posterior shapes

We use the samples of the posterior distributions of z_{obs} to create posterior ice sheet shapes and consider how plausible the shapes are, both when modelling the ice core

sites as two separate regions or combined. Figures 7.1 and 7.2 shows the mean posterior shape from the three models. Modelling just the WAIS ice core sites produces an ice sheet reconstruction with a large extent but low elevation. When modelling just the EAIS ice core sites, the ice sheet reconstruction is noticeably thicker, but drops off much more quickly over the Ronne and Ross ice shelves (see Figure 1.1 for the location of these). When all AIS sites are modelled simultaneously, we can see the heightened elevation in the EAIS remains, along with the discrepancies over the ice shelves, whilst the WAIS is smoothed out.

Figure 7.3 shows the difference between the mean posterior shape when modelling all sites, and the mean posterior shapes when modelling sites as two sets. We can see that the reconstructions for all sites and EAIS sites are very similar, whilst the posterior shape created from just WAIS sites has more obvious differences in both negative and positive differences in elevation. This would suggest the ice core sites in the EAIS are heavily influencing the AIS model; given the issues with our HadCM3 simulations concerning the EAIS, discussed in Section 5.2.1 and Figure 5.3 this may raise questions over the plausibility of the ice sheet shapes that our models create. Due to the difficulties matching with proxy data, our emulators for the EAIS ice core sites are extrapolating a long way from our simulations. This makes it harder to trust in the results for the EAIS; the sparsity of data and the size of the ice sheet makes it difficult to model the EAIS as a whole.


Figures 7.4 and 7.5 show the standard deviation of the posterior reconstructions. Modelling all ice core sites simultaneously greatly reduces the standard deviation in our reconstructions; modelling sites in each ice sheet separately appears to result in a

great deal of uncertainty in the other ice sheet. One of our aims in this thesis was to reduce uncertainty around the AIS at the LGM; comparing the standard deviation of the posterior ice sheet shapes from our models to Figure 4.2, we can see there is a lot less variance in our posterior shapes compared to the prior model.


Sets of ice sheet reconstructions from the three models are in Figures 7.6, 7.7 and 7.8. We can see the reconstructions produced when modelling just the WAIS sites have lower elevation, although the ice sheets are smoother. When modelling the EAIS sites, or all sites combined, there are discrepancies in the ice sheet reconstructions, particularly over the Ronne and Ross ice shelves. Our model is statistical and has discarded the underlying physics of previous reconstructions; this means that it is difficult to prevent small implausibilities in our posterior shapes such as holes in the ice shelves. As such, the posterior ice sheet shapes that our models produce would be best used as suggestions for the presence of ice at the LGM, and demonstrations of where there is most uncertainty in predictions, rather than as a precise orography estimate.

7.2 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 GLAC-1D (Briggs et al., 2014). These are the four



(a) The mean posterior ice sheet reconstruction when modelling the WAIS ice core sites.



(b) The mean posterior ice sheet reconstruction when modelling the EAIS ice core sites.

Figure 7.1: The mean posterior ice sheet reconstructions when modelling WAIS and EAIS ice core sites separately. We can see that using the EAIS ice core sites produces a mean reconstruction with more extreme high and low values. The mean reconstruction from the WAIS ice core sites is smoother.




Figure 7.2: The mean posterior ice sheet reconstruction when modelling all ice core sites.

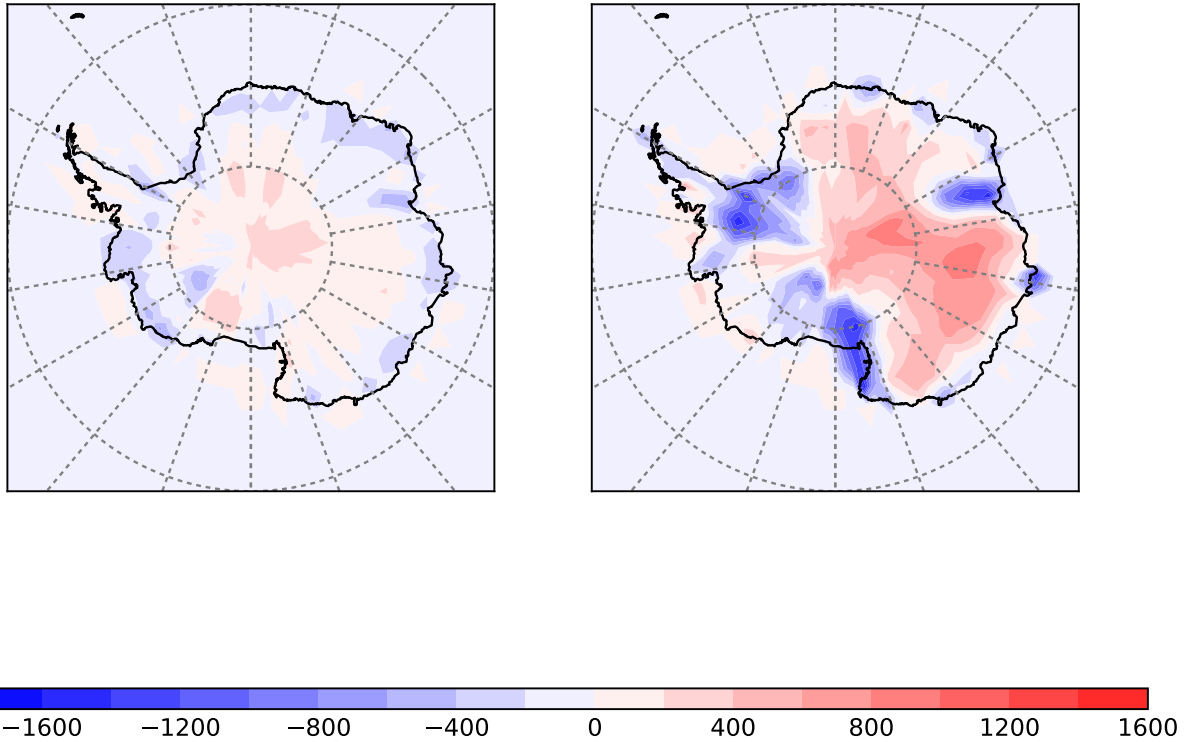




Figure 7.3: The difference between the mean posterior reconstruction of the model using all sites, and the models separating sites into EAIS (left plot) and WAIS (right plot). We can see that the mean reconstruction created from just WAIS ice core sites leads to an ice sheet shape of more extreme values than using all sites together.



(a) The standard deviation of the posterior ice sheet reconstructions when modelling the WAIS ice core sites.



(b) The standard deviation of the posterior ice sheet reconstructions when modelling the EAIS ice core sites.

Figure 7.4: The standard deviation of the posterior ice sheet reconstructions when modelling WAIS and EAIS ice core sites separately.




Figure 7.5: The standard deviation of the posterior ice sheet reconstructions when modelling all ice core sites.

7.2. COMPARING OUR OROGRAPHIES TO PREVIOUS RECONSTRUCTIONS 173

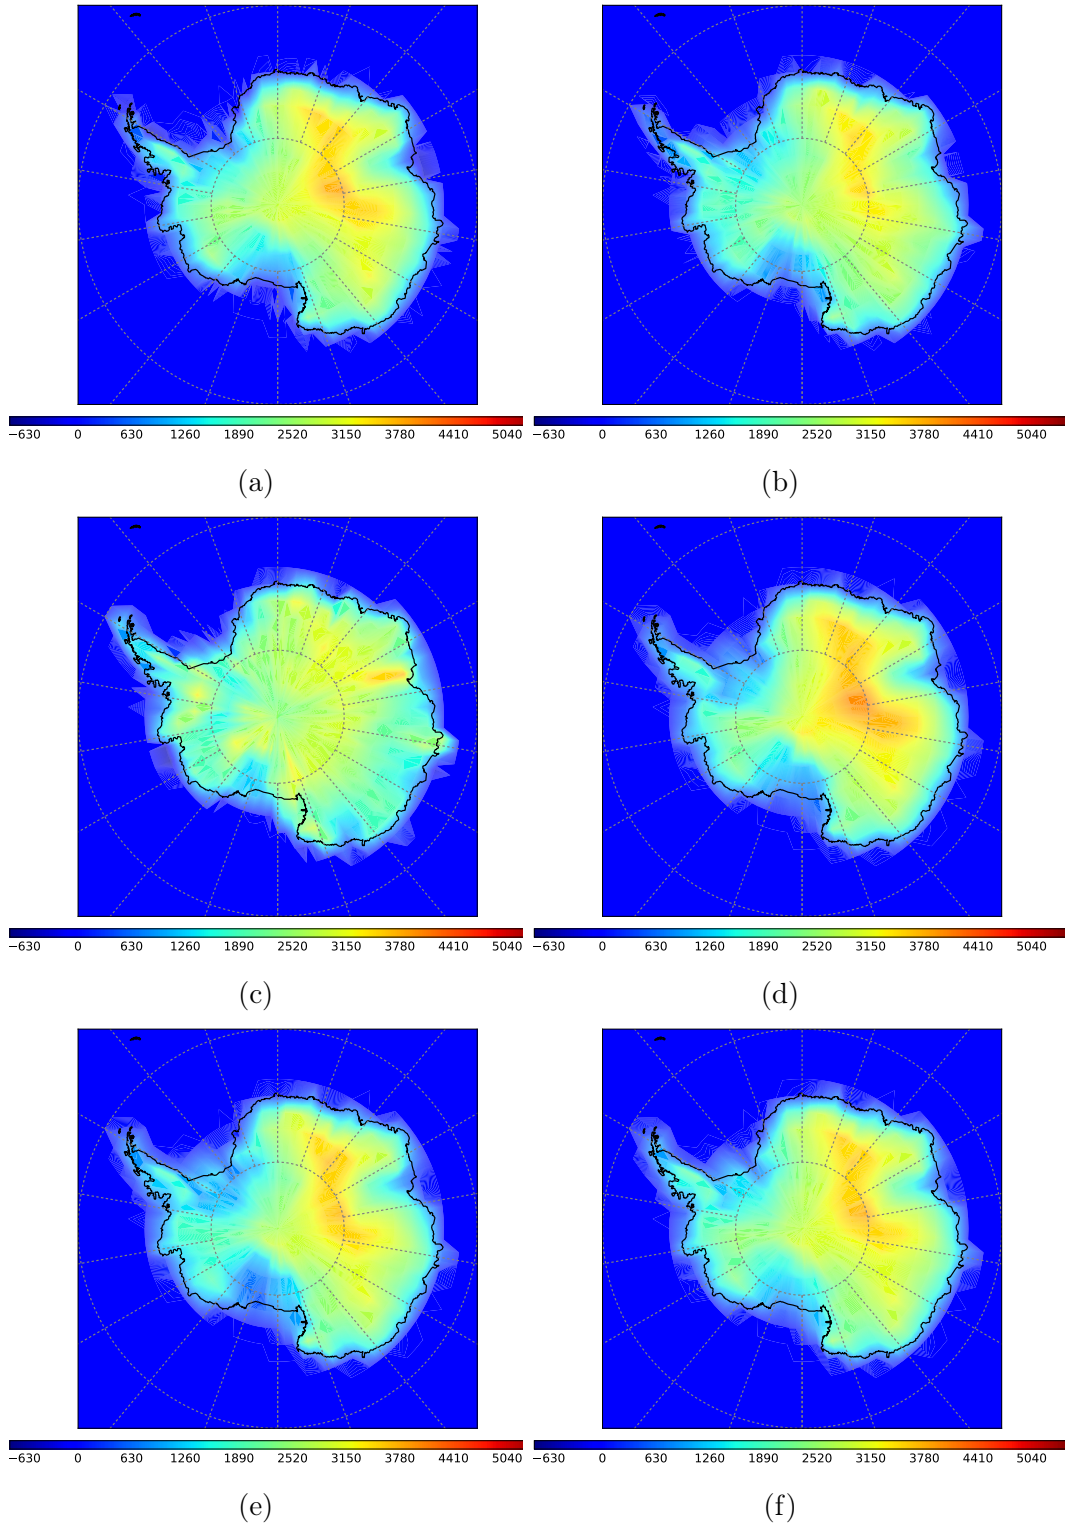


Figure 7.6: Posterior shapes from modelling the WAIS ice core sites.

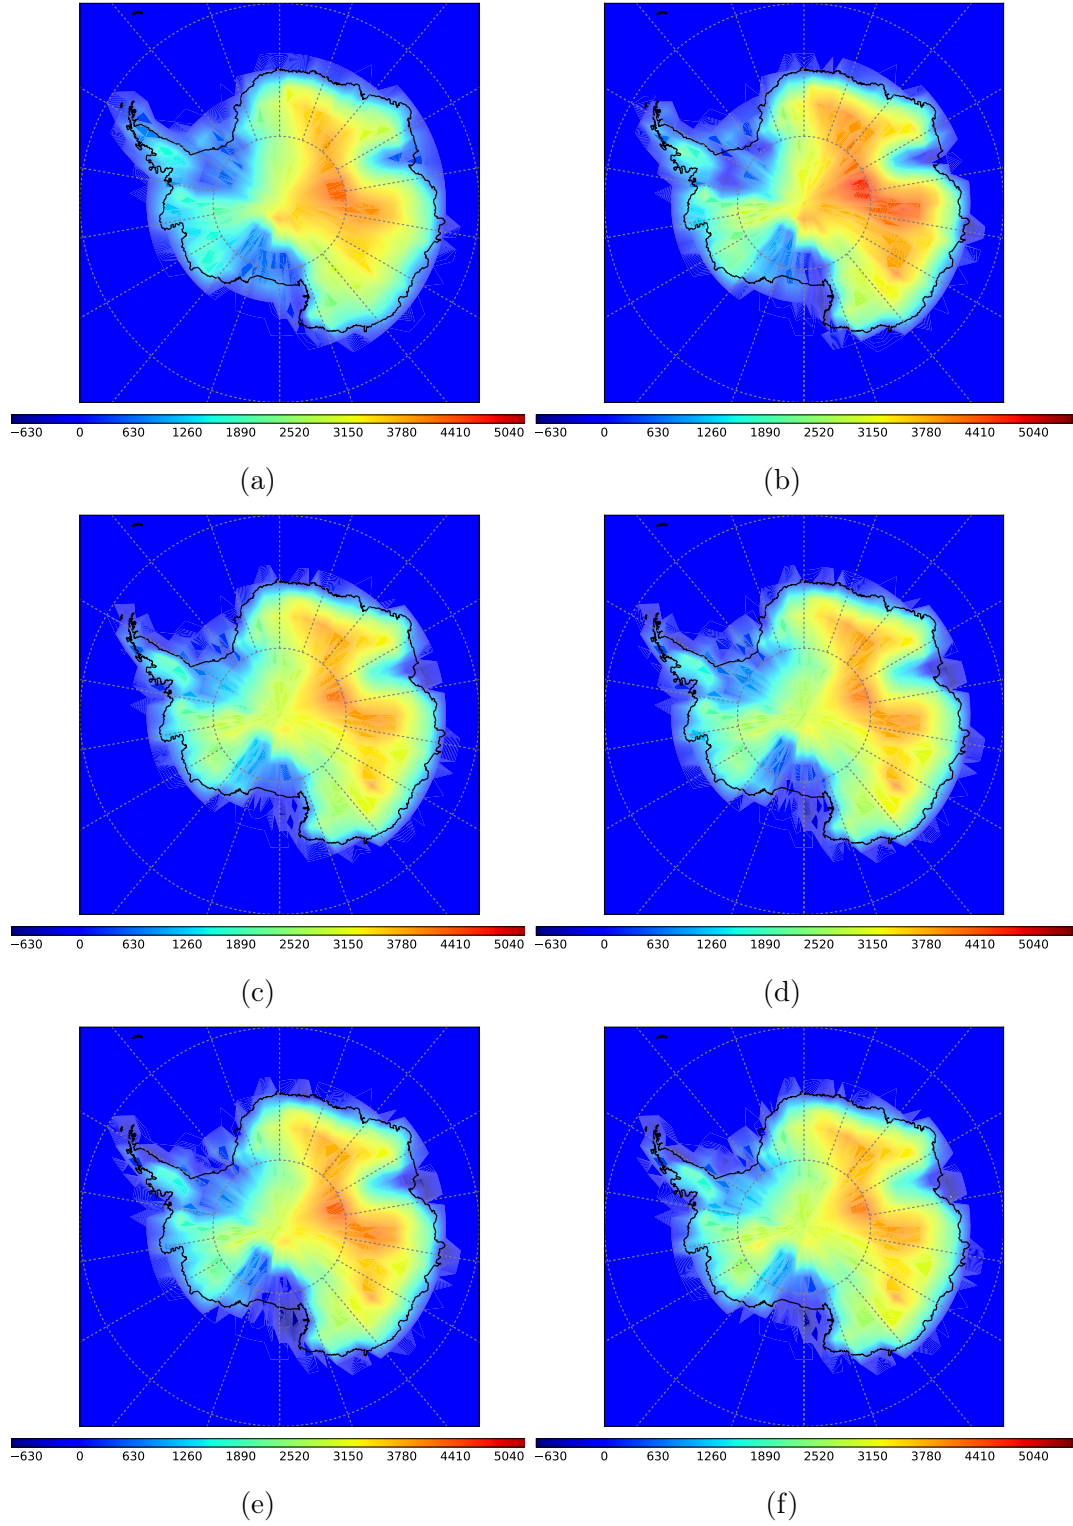


Figure 7.7: Posterior shapes from modelling the EAIS ice core sites.

7.2. COMPARING OUR OROGRAPHIES TO PREVIOUS RECONSTRUCTIONS 175

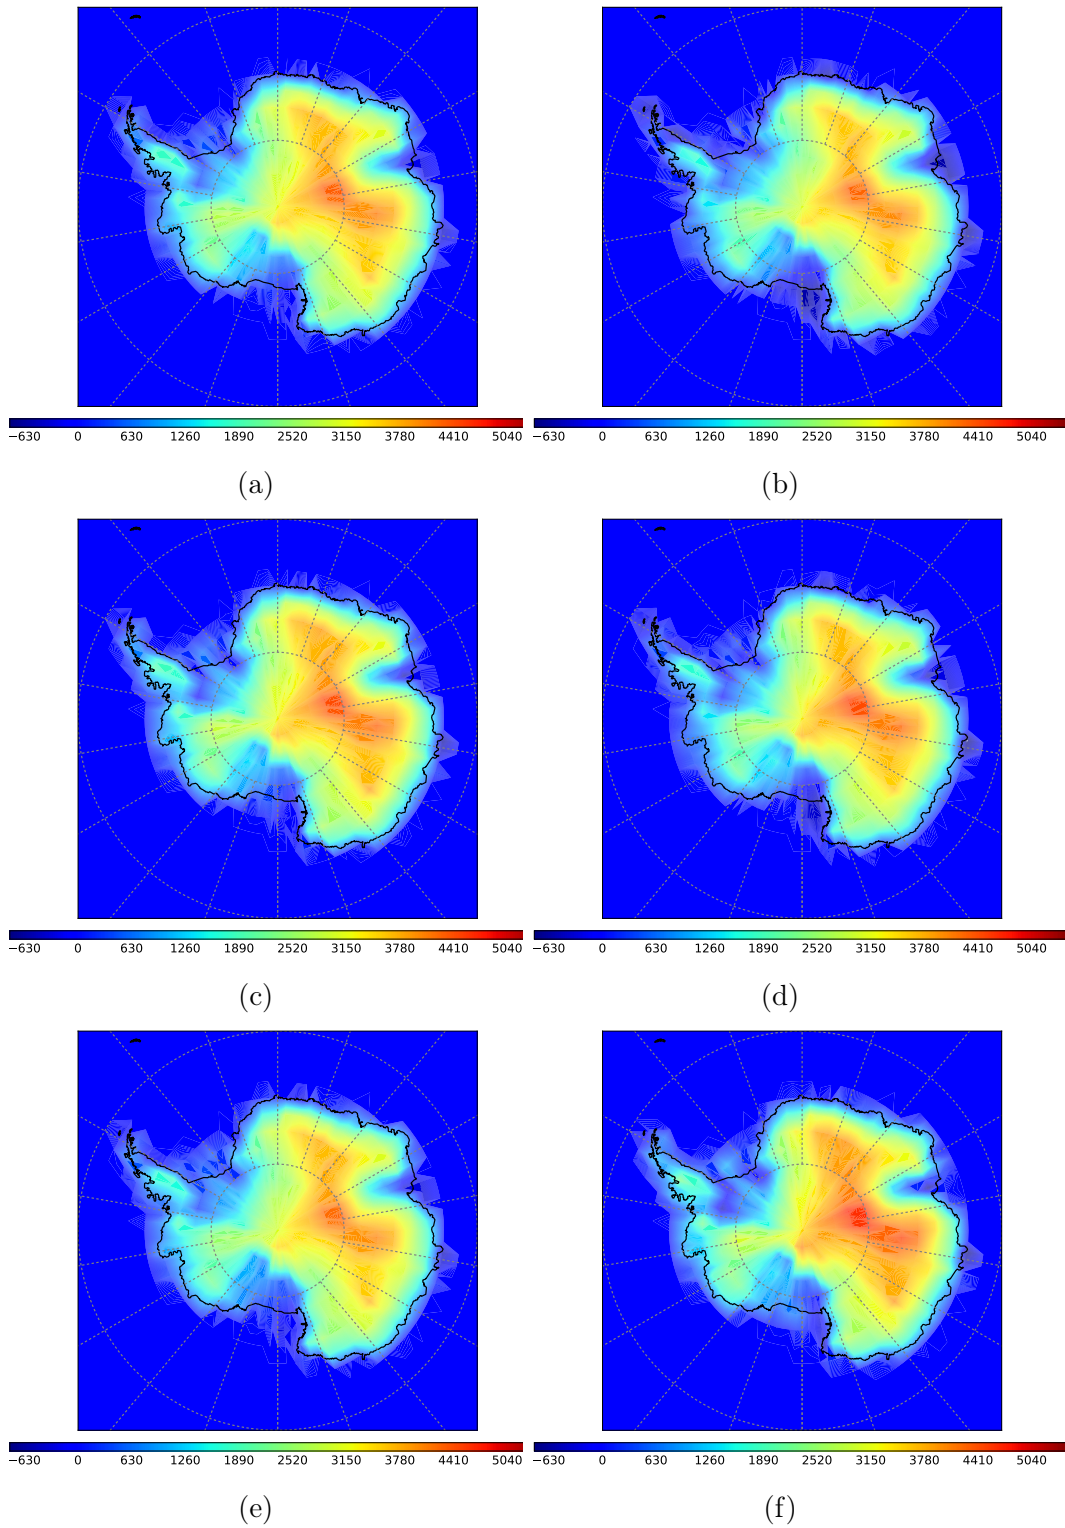


Figure 7.8: Posterior shapes from modelling all of the ice core sites.

reconstructions used in the first round of PCA, as described in Section 3.3. This allows us to see where the posterior reconstructions disagree with the original ice sheet reconstructions. We cannot compare our posterior ice sheet shapes to all proxy evidence or reconstructions, as that is beyond the scope of this work. We compare our results to the four reconstructions we originally started with to give some examples of how our method differs to previous work.

We can see in Figures 7.9, 7.10 and 7.11 that all models, whether incorporating only WAIS, EAIS, or all ice core sites, produce ice sheets with a lower elevation than ICE-5G. This is not unexpected, as ICE-5G has a much thicker Antarctic ice sheet than later models due to updated methodologies, as described in Argus et al. (2014).

Our model also has a lower coastal region than ICE-6G, and when modelling the EAIS ice core sites, there are obvious differences over the Ronne and Ross ice shelves. Despite this, our model using all sites is producing shapes with a higher elevation over the EAIS than the reconstructions used in our prior model, suggesting this is where some of the "missing ice", described in 2.1.1 and previously unaccounted for in other ice sheet reconstructions, may be found.

Looking at Table 7.1 we can see that modelling all ice core sites at once produces the ice sheets with highest mean elevation. Although still smaller than ICE-5G, the average shape from our model is a lot larger than the other three reconstructions. Due to the lack of physical modelling underlying our ice sheet shapes, we are unable to give an estimate of the ice volume or area that our posterior shapes hold. This

	Modelling WAIS sites	Modelling EAIS sites	Modelling all sites
ICE-5G	-214673.99	-134678.46	-107646.68
ICE-6G	-16260.46	63735.07	90766.84
W12A	46272.65	126268.18	153299.95
GLAC-1D	40381.59	120377.12	147408.89

Table 7.1: Total difference in the mean elevation (m) of the posterior reconstructions when modelling WAIS, EAIS or all ice core sites, and four of the original ice sheet reconstructions.

means we cannot judge the sea level contribution of our posterior ice sheet shapes, as that would be beyond the remit of this thesis.

There is much to consider when comparing the two methods of finding posterior ice sheet shapes, modelling the ice core sites altogether or in two separate groups. Modelling all sites combined results in a mean posterior shape with higher elevation and less uncertainty than the models using WAIS and EAIS ice core sites separately. We can also see in Figure 6.11 that this model gives posterior distributions that differ more from the prior distributions of z_{obs} , suggesting that we are learning more about the distributions of z_{obs} and the AIS at the LGM when modelling all ice core sites together. Although there is a slight discrepancy in the generated distributions of y_{pred} compared to the distributions when the ice core sites are modelled separately as EAIS and WAIS sites, as can be seen in Figure 6.8, the differences are small. The model incorporating all ice core sites therefore seems to give the better estimate of the AIS at the LGM, giving an approximate ice sheet shape and elevation with reduced uncertainty.

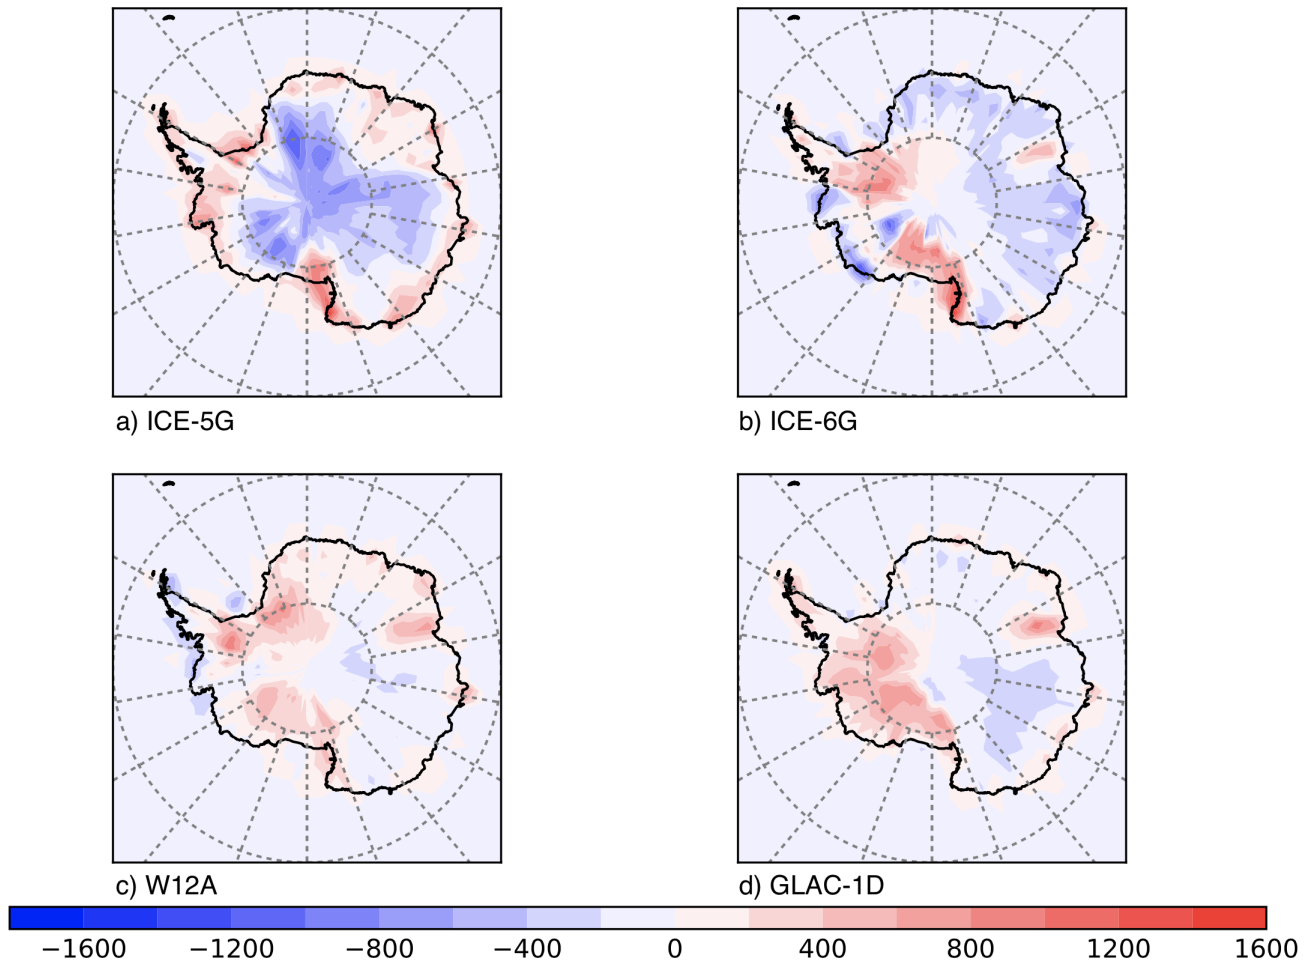


Figure 7.9: The difference between the posterior mean shape our model creates when using the WAIS ice core sites and the four ice sheet reconstructions used in the first PCA. Clockwise they are ICE-5G (Peltier, 2004), ICE-6G (Argus et al., 2014), GLAC-1D (Briggs et al., 2014) and W12A (Whitehouse et al., 2012).

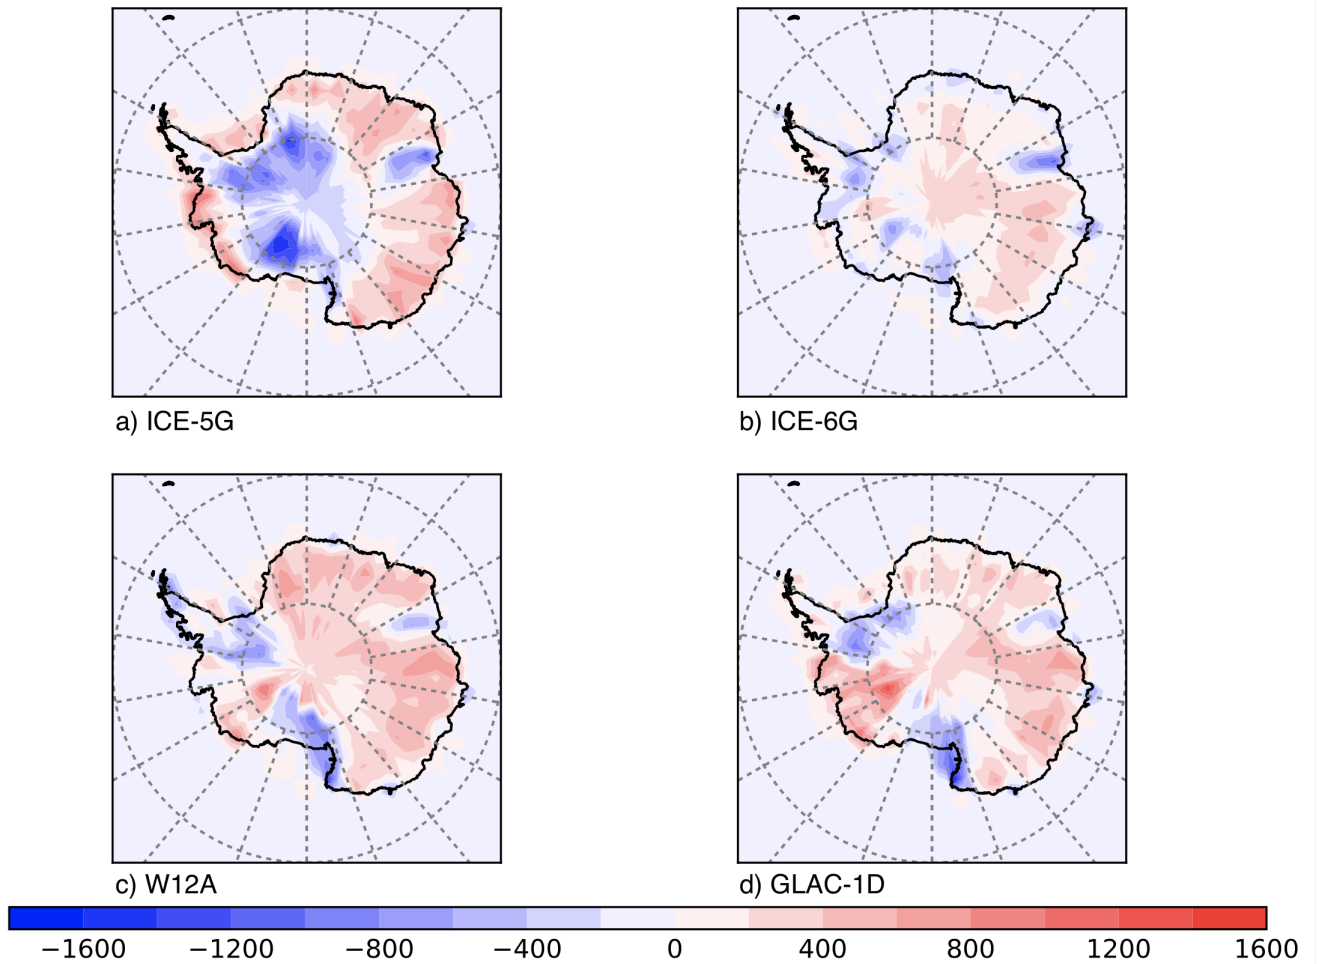


Figure 7.10: The difference between the posterior mean shape our model creates when using the EAIS ice core sites and the four ice sheet reconstructions used in the first PCA. Clockwise they are ICE-5G (Peltier, 2004), ICE-6G (Argus et al., 2014), GLAC-1D (Briggs et al., 2014) and W12A (Whitehouse et al., 2012).

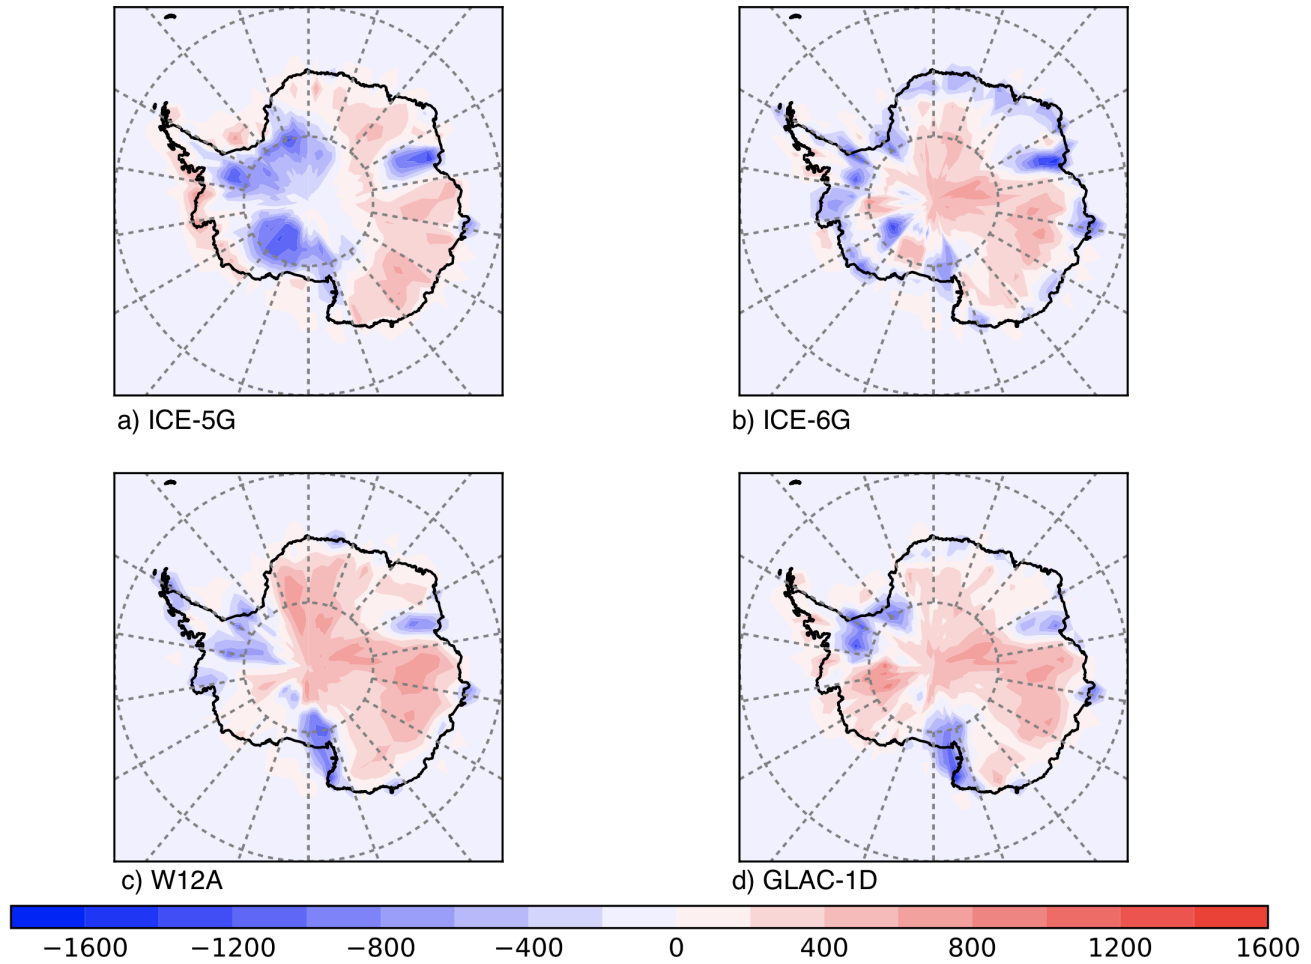


Figure 7.11: The difference between the posterior mean shape our model creates when using all ice core sites and the four ice sheet reconstructions used in the first PCA. CClockwise they are ICE-5G (Peltier, 2004), ICE-6G (Argus et al., 2014), GLAC-1D (Briggs et al., 2014) and W12A (Whitehouse et al., 2012).

7.3 Summary of Chapter

We have shown that our model produces ice sheet shapes with less uncertainty than previous reconstructions by considering Figures 7.4, 7.5 and 4.2. We can see that our posterior ice sheet shapes have less uncertainty than the set of ice sheet reconstructions we used to build the prior model. Although not as large as ICE-5G, our mean posterior shapes are larger than previous reconstructions, accounting for some of the "missing ice" at the LGM. Due to the statistical nature of our model, we are unable to give an estimate of the contribution our posterior reconstructions make to global sea level rise; this would be an area of interest in future work. Despite issues modelling the EAIS resulting from our HadCM3 simulations, we believe a model using all ice core sites together gives the best reconstruction of the AIS at the LGM, estimating the shape with reduced uncertainty.

Chapter 8

Conclusions

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. Several items of interest are beyond the remit of this thesis but would be of interest for future research.

8.1 Discussion

8.1.1 Scientific Findings

Our methods have produced ice sheets with less variance than our prior model and reduced the uncertainty around the AIS at the LGM. Our method of course relies on the previous reconstructions all being plausible for the LGM, and the model space they describe containing all possible shapes. We are therefore reliant on the belief that this model space does not exclude the "true" ice sheet shape, as that would make it impossible for our model to recover it. Posterior samples (Figure 6.12) show that we are taking values of the prior variables in the model space around our four "important" reconstructions ICE-5G (Peltier, 2004), ICE-6G (Argus et al., 2014), W12A (Whitehouse et al., 2012), and GLAC-1D (Briggs et al., 2014). These are the four reconstructions used in the first round of PCA, as described in Section 3.3. Our model is built on the variance within and between these reconstructions; we are dependent on the prior beliefs of the modellers who created these reconstructions being accurate.

Our posterior ice sheet shapes are larger than the majority of the previous reconstructions. An exception is ICE-5G, which is recognised by the palaeoclimatological community to have a larger ice mass at Antarctica than is currently accepted (Peltier et al., 2015). This increase in size in our estimates confirms the widely held belief that at least some of the "missing ice" at the LGM can be found in the AIS (Simms et al., 2019). Although we are not able to give an estimate of the ice volume and

corresponding sea level contribution from our posterior shapes, unlike previous reconstructions, we have provided estimates of where this ice may be, and the standard deviation within our model in Figure 7.5. Such uncertainty statements may point to directions for future modelling and/or observational work that could allow further refinement of the estimates.

8.1.2 Statistical Findings

We have created a novel approach to weighted PCA, as described in Section 3.3.3. 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 our are not exclusive to the climatology community, and this weighted PCA method is applicable to many other areas.

We have demonstrated a way of combining prior knowledge, expert elicitation, GP emulation and Bayesian inference to model the uncertainty around aspects of the palaeoclimate. Building on the previous work of Domingo et al. (2020), 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 a very broad event; conditions changed very slowly, allowing us to disregard

any dating errors in the isotopes as described in 6.6. 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.

8.2 Future Work

The orographies we chose to input into HadCM3 appear to have been too low to create isotope values similar to ice core observations for the ice cores from the EAIS, as can be seen in Figure 5.3. Whether this was due to a model discrepancy within HadCM3, or an error in our elicitation process, is difficult to tell. Any future attempt to replicate this method should check the choices of orographies before running simulations on HadCM3, whether through a more detailed or rigorous expert elicitation process, or by staggering the simulations so that there can be a preliminary exploratory data analysis of an initial set of simulations, with adjustments made according to early results.

There appear to be discrepancies over the Ronne and Ross ice shelves in our posterior shapes, as discussed in Chapter 7. We did not have the time or resources to investigate the causes of this. It occurred when the EAIS ice core sites were incorporated into our calibration model, so may be a consequence of the $\delta^{18}O$ simulations from HadCM3 not matching the ice core observations.

We have created a statistical model, having discarded the underlying physical models of the ice sheet reconstructions. As a result, our model can output implausible ice sheet shapes; a future collaboration with ice sheet modellers and climatologists could provide further constraints on our model to improve output. 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. This means our posterior ice sheet shapes can only be viewed as guides for the location of ice, rather than a precise map of the AIS at the LGM. Our conclusions in Chapter 7 have therefore been on elevation rather than exact orographies.

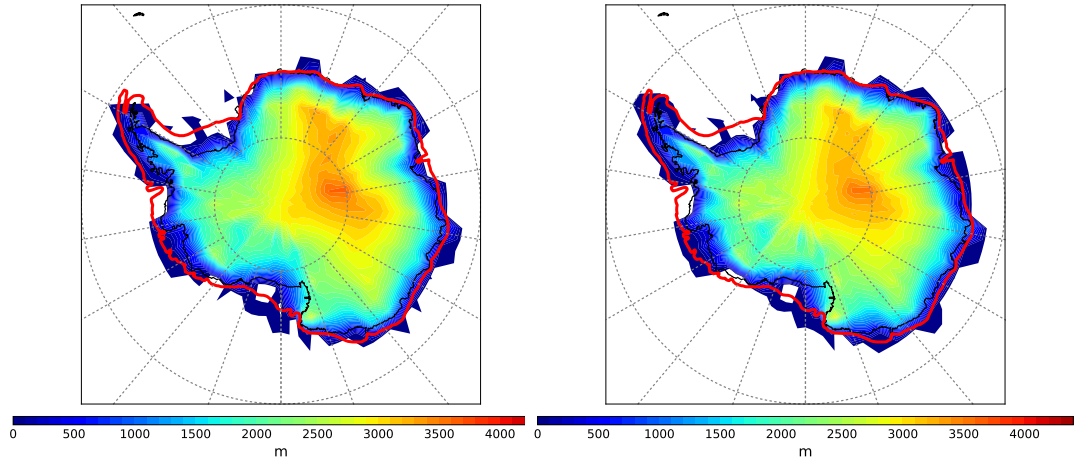
Further testing with another GCM could allow us to understand better any biases inherent in HadCM3 and that we may not have been able to remove from our model (Domingo et al., 2020). The possible model discrepancy causing our difficulties with modelling the EAIS could then be further investigated, along with any other errors our choice of GCM has created. As discussed in Section 7.1, the EAIS ice core sites appear to heavily influence our model despite the difficulties we have encountered. If carrying out a similar approach in the future, carefully constraining the GCM simulations or choosing a time period with more available proxy data could avoid this issue. More data, whether from ice cores or otherwise, would mean less extrapolation over the ice sheet would take place in our model. However, we are hampered by the nature of proxy data collection, as we cannot control when, how or where such data is collected. Particularly for the EAIS, the rate of accumulation of data is slow, and collection is difficult due to the inhospitable conditions.

As already stated, it is beyond the remit of this work to give an estimate of sea level contribution from our posterior ice sheet shapes. A conversion from our elevation estimates to ice volume would be useful to see how much of the "missing ice" at the LGM is placed in Antarctica by our method, and would contribute to the understanding of how the AIS has reacted to the changing climate.

In summary, the methods proposed in this thesis could be applied easily 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.

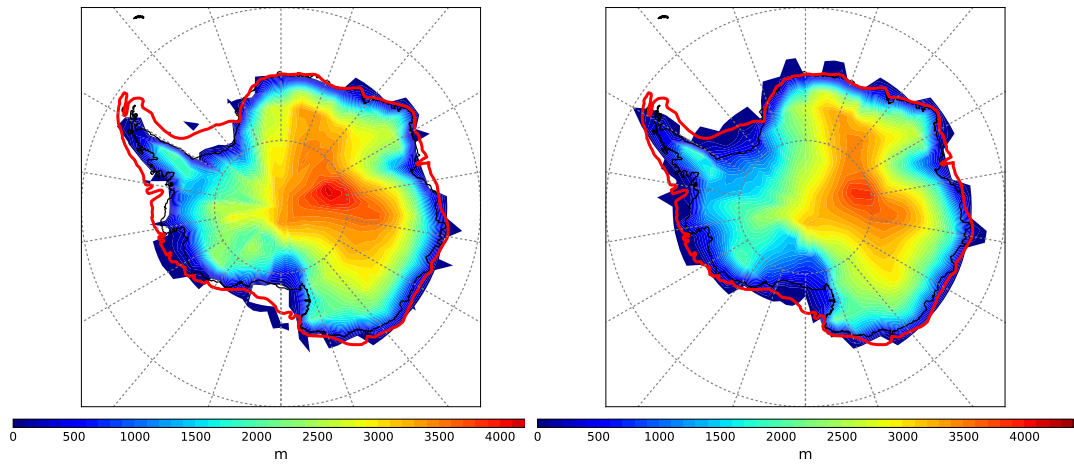
Appendices

A Design of eighty shapes from elicitation.



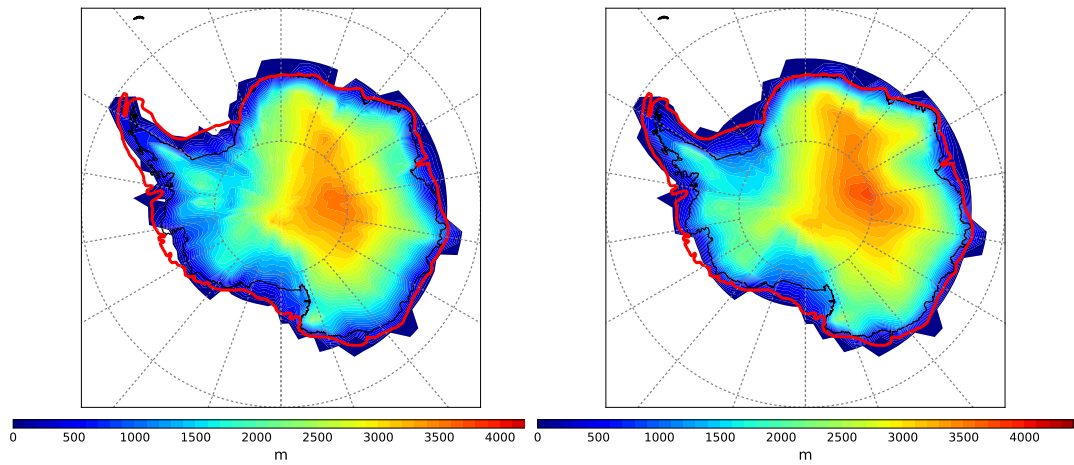
(a)

(b)



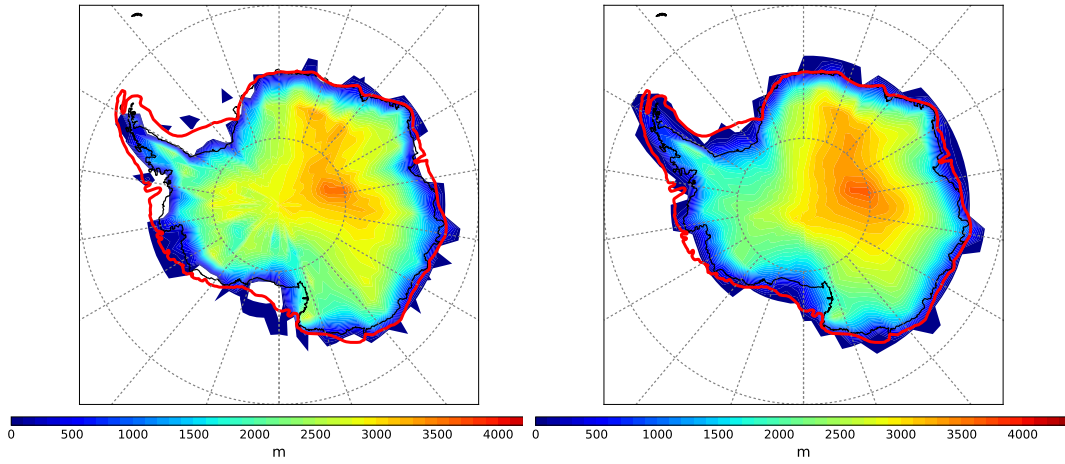
(c)

(d)



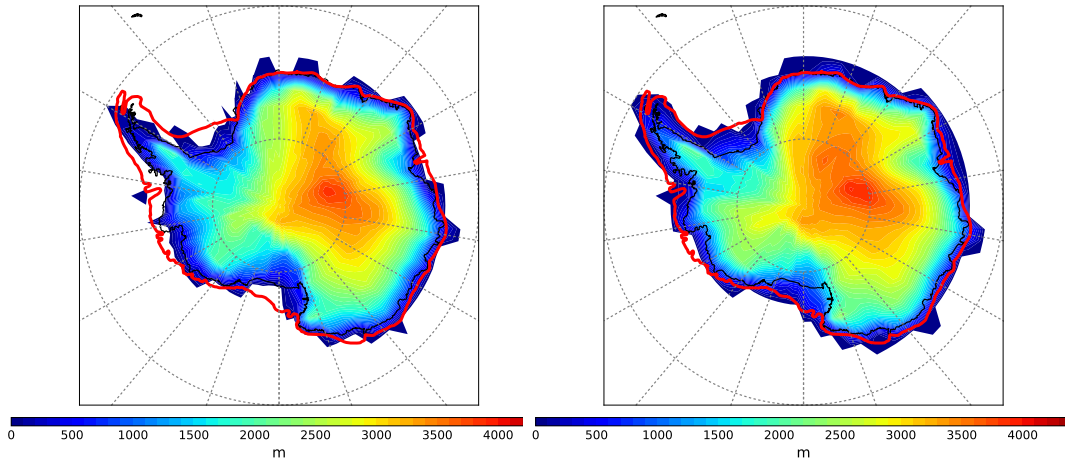
(e)

(f)



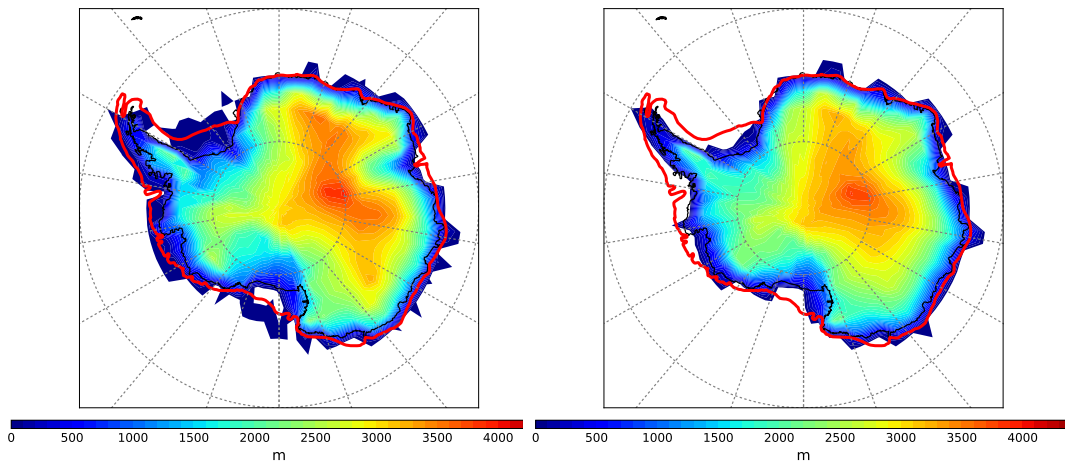
(g)

(h)



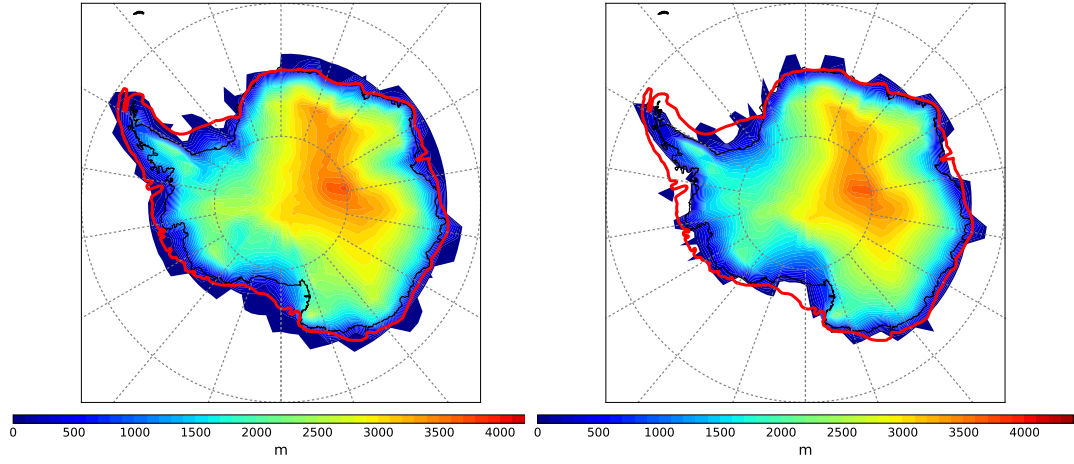
(i)

(j)



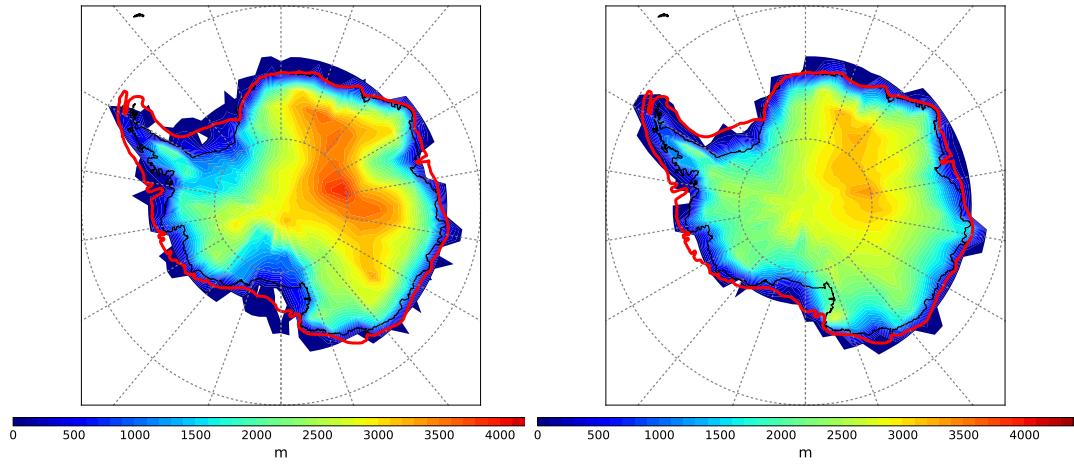
(k)

(l)



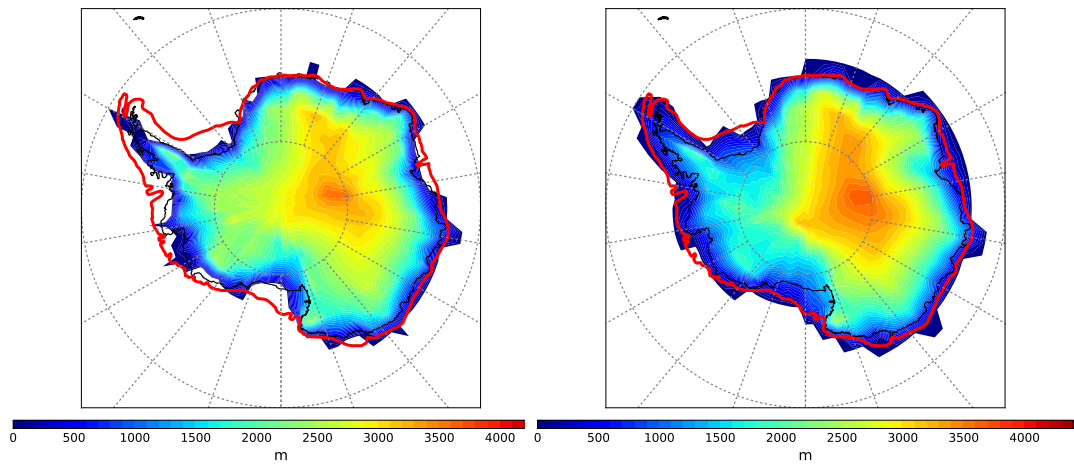
(m)

(n)



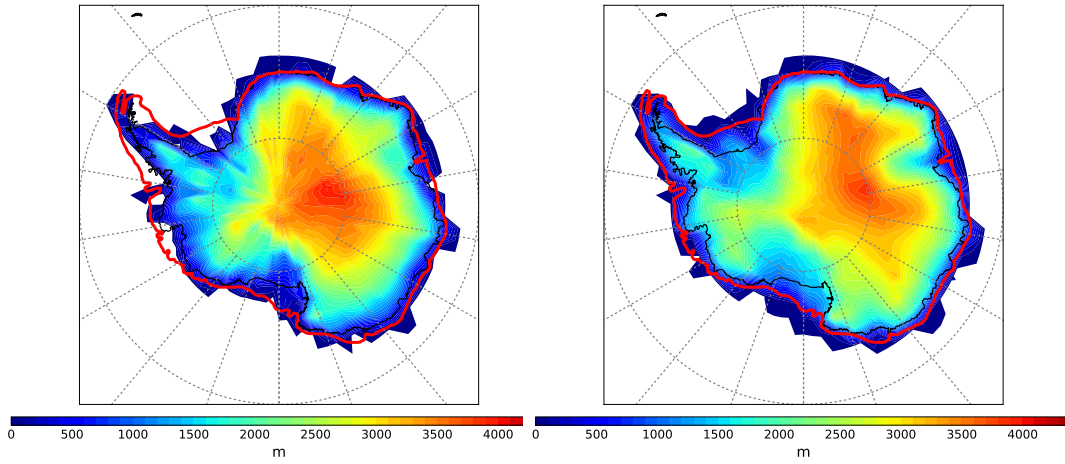
(o)

(p)



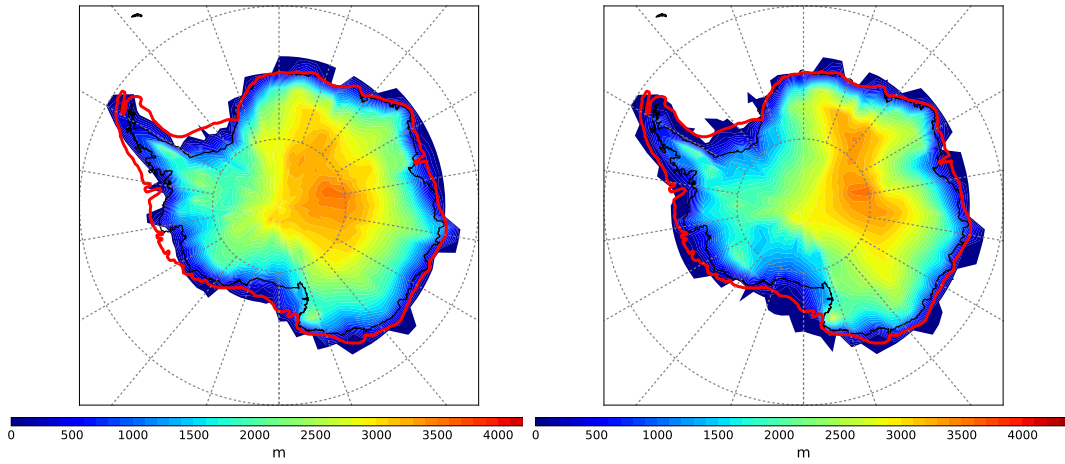
(q)

(r)



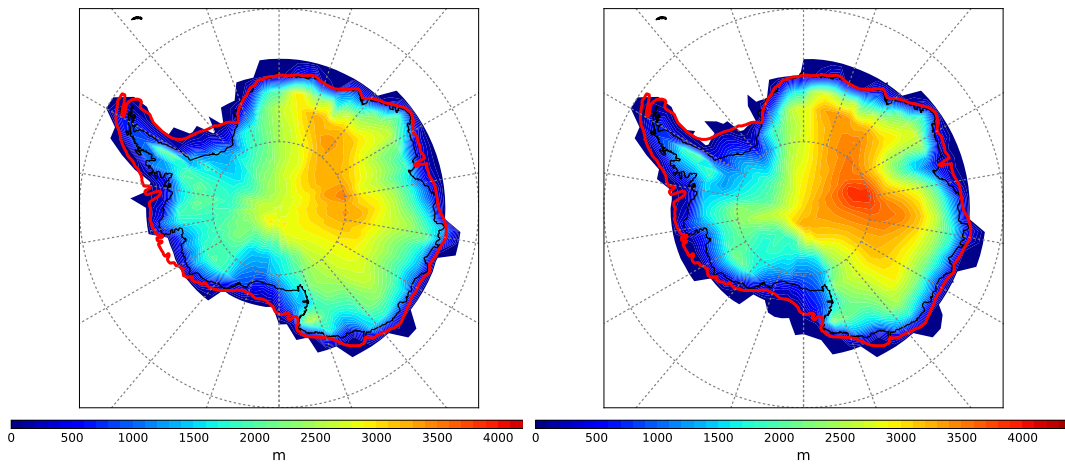
(s)

(t)



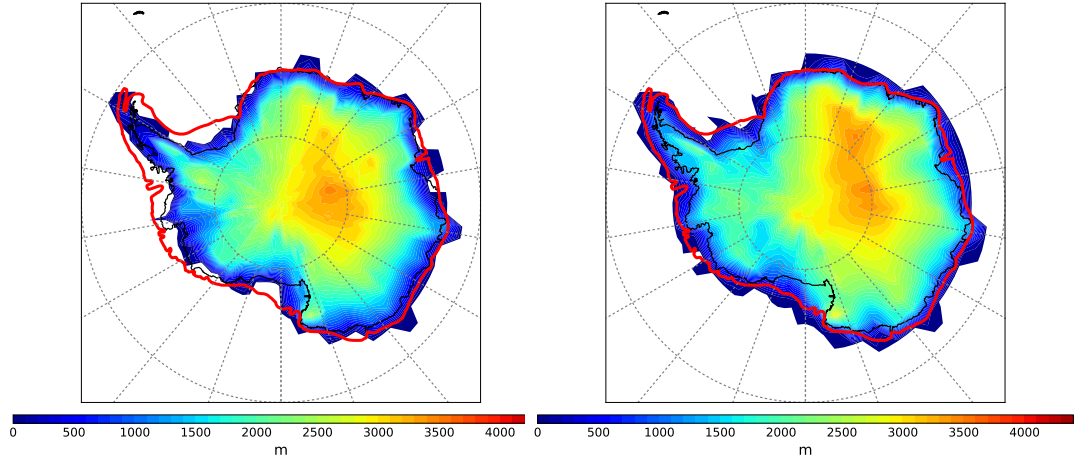
(u)

(v)



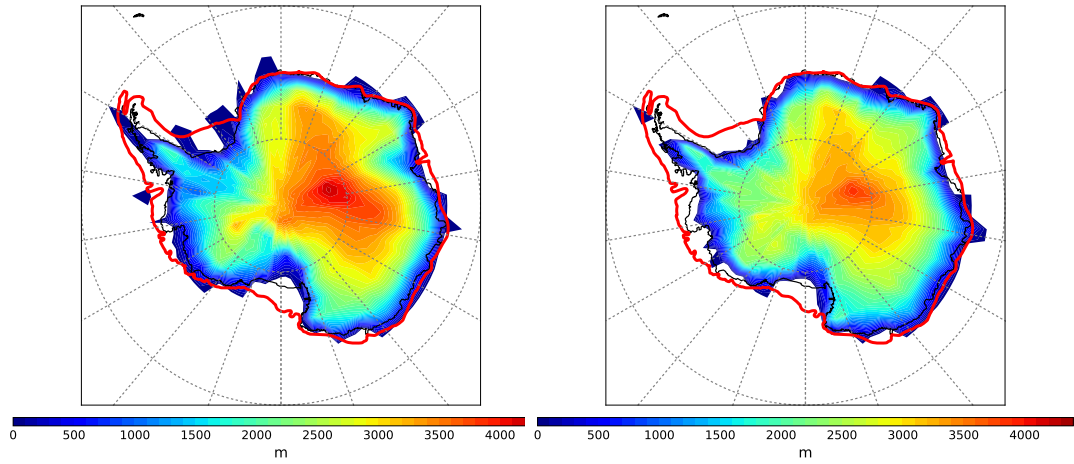
(w)

(x)



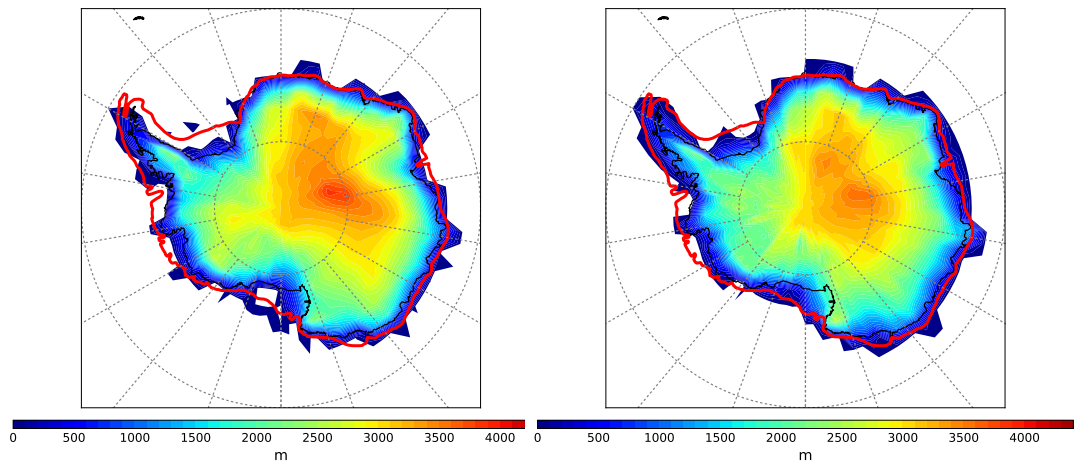
(y)

(z)



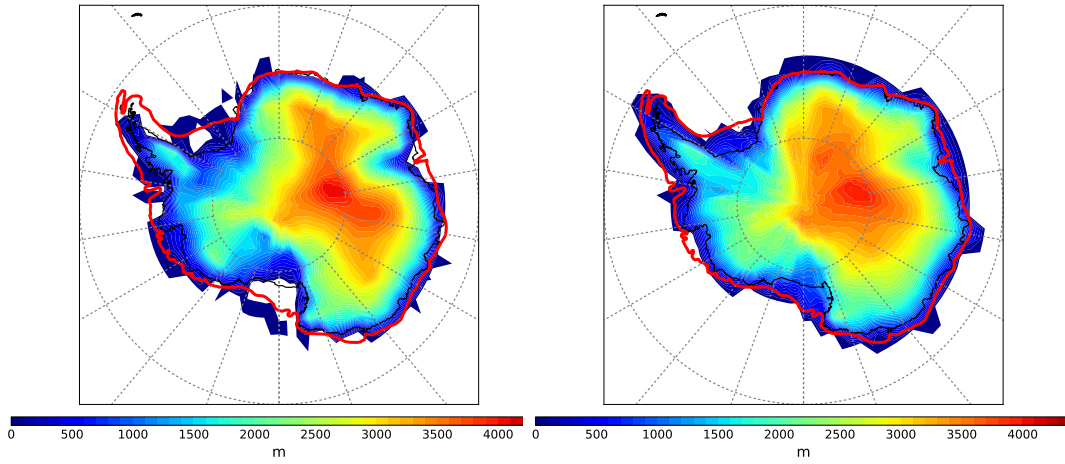
(aa)

(ab)



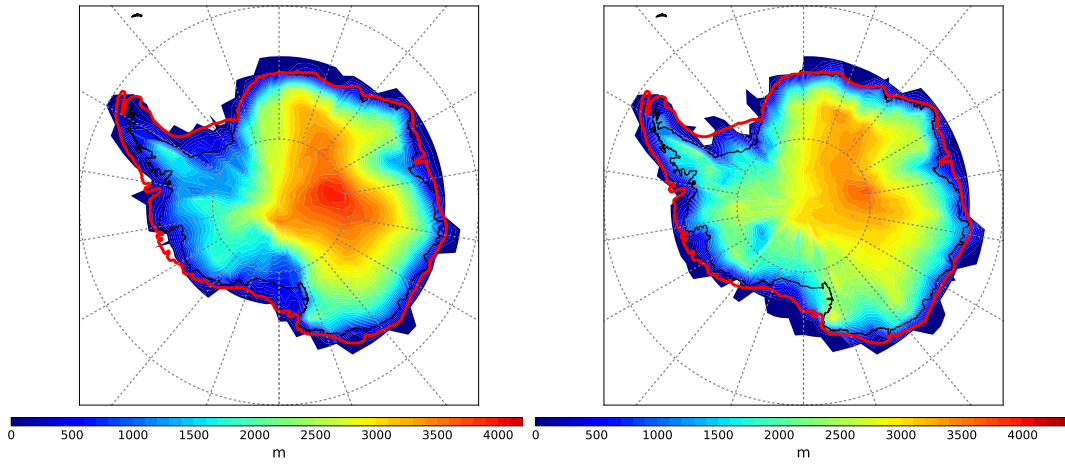
(ac)

(ad)



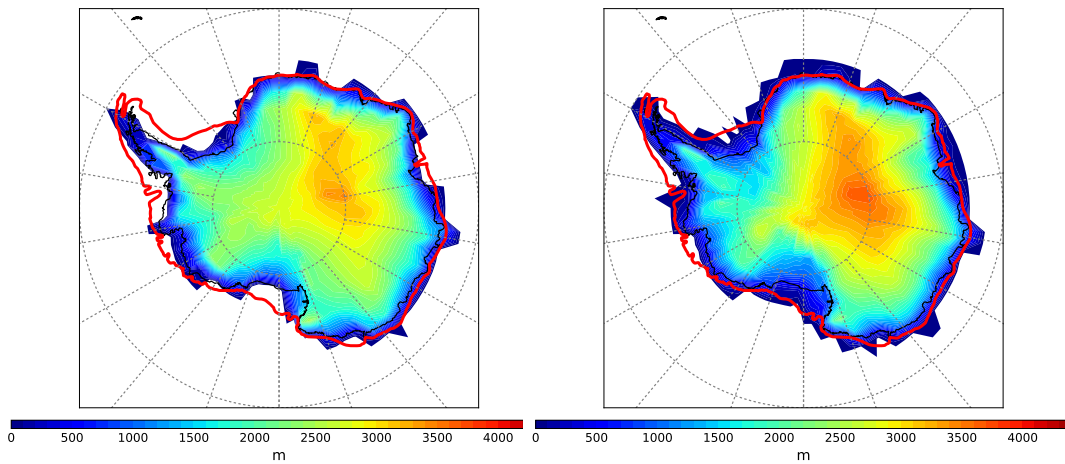
(ae)

(af)



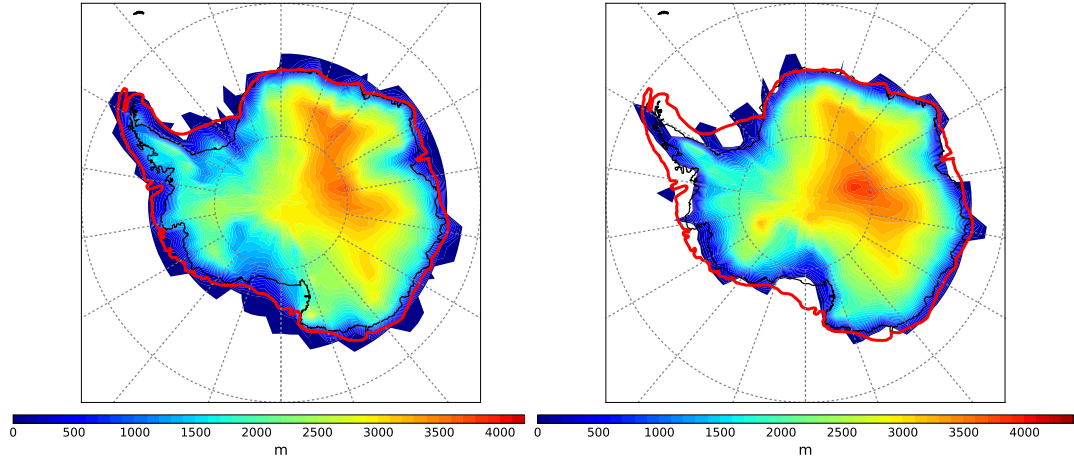
(ag)

(ah)



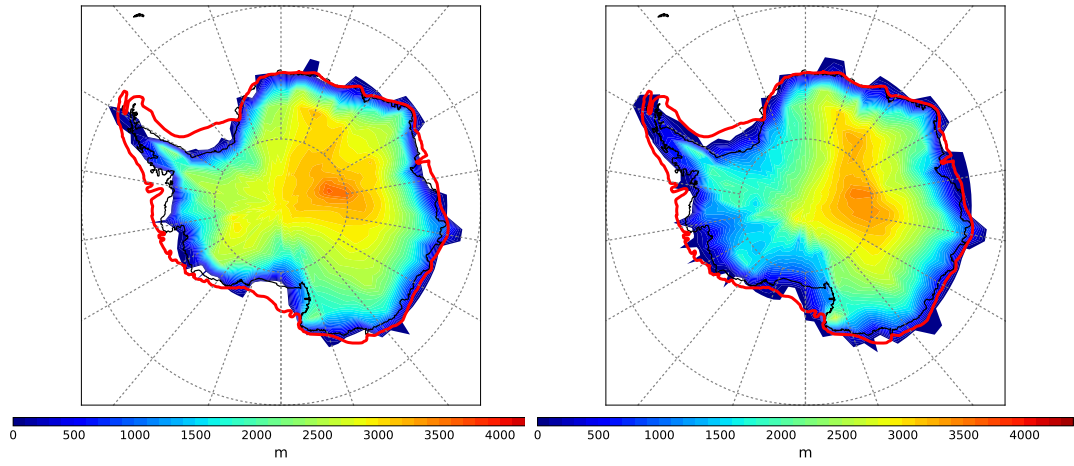
(ai)

(aj)



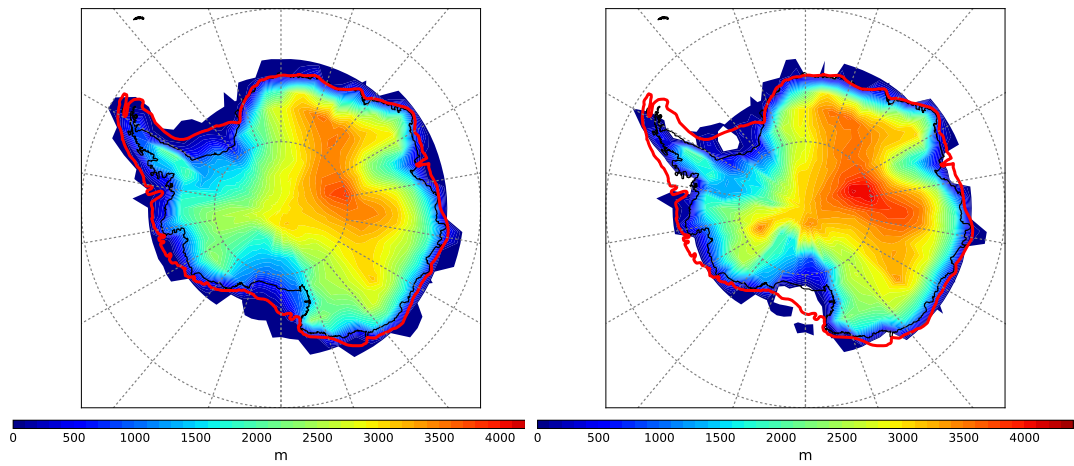
(ak)

(al)



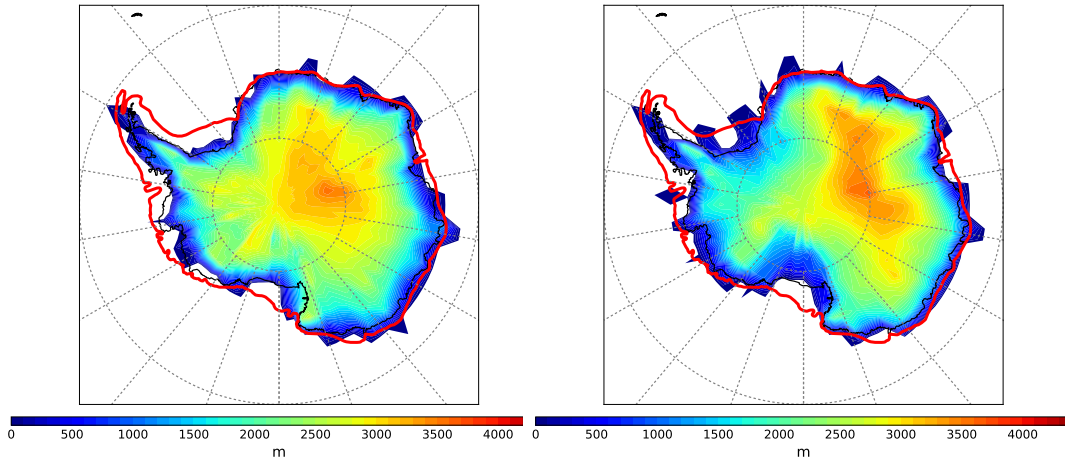
(am)

(an)



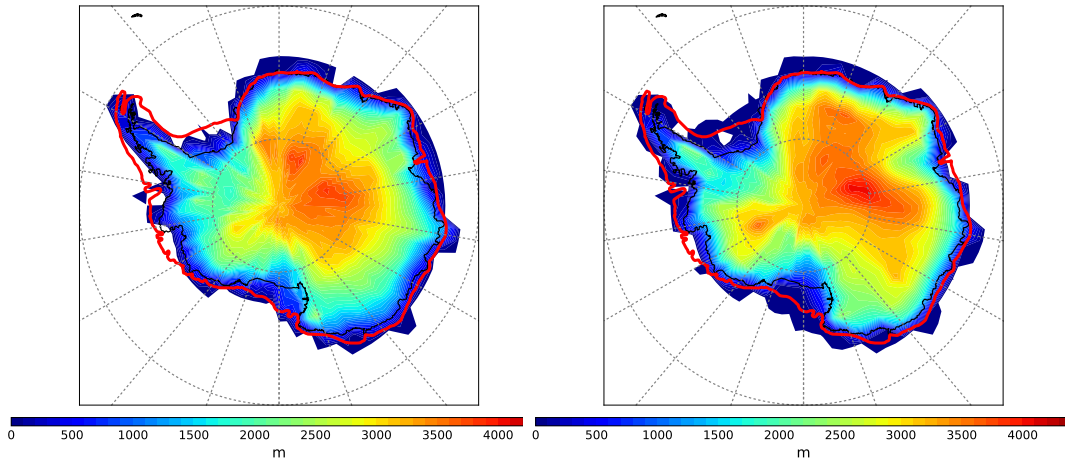
(ao)

(ap)



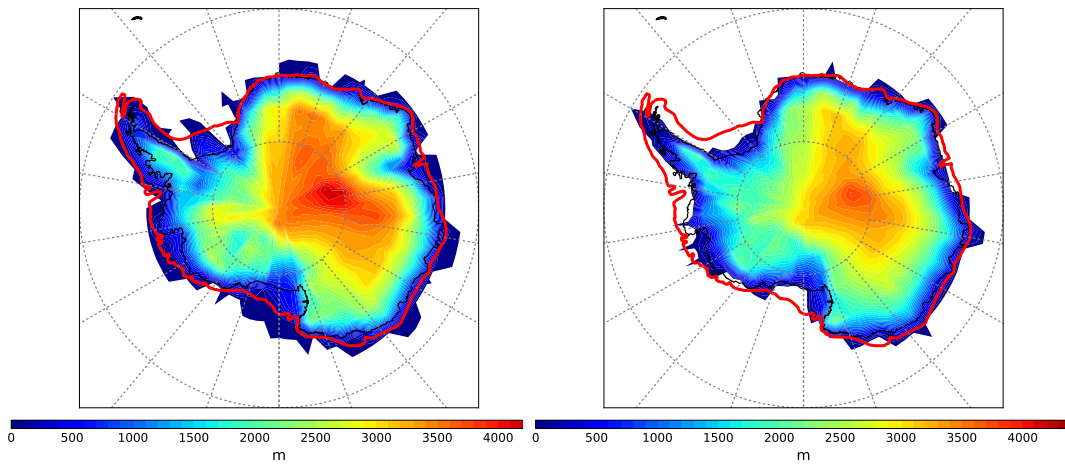
(aq)

(ar)



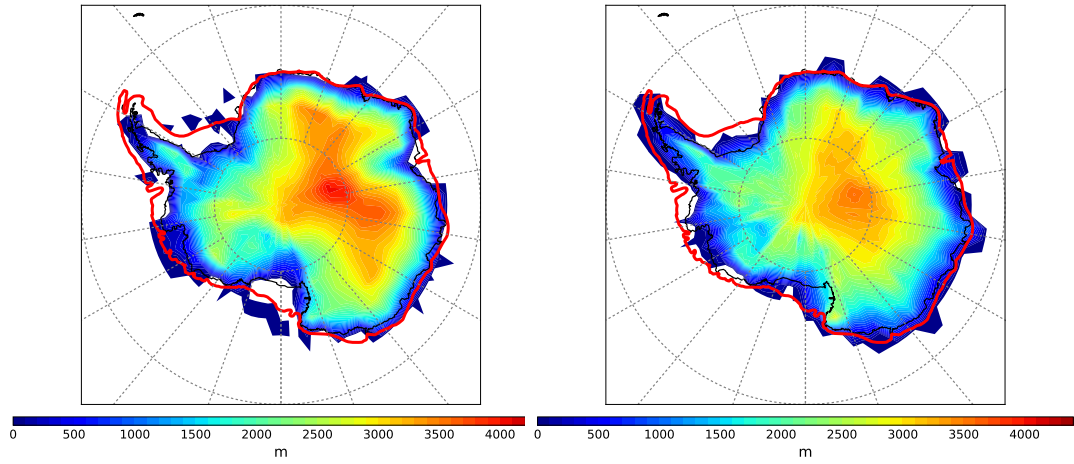
(as)

(at)



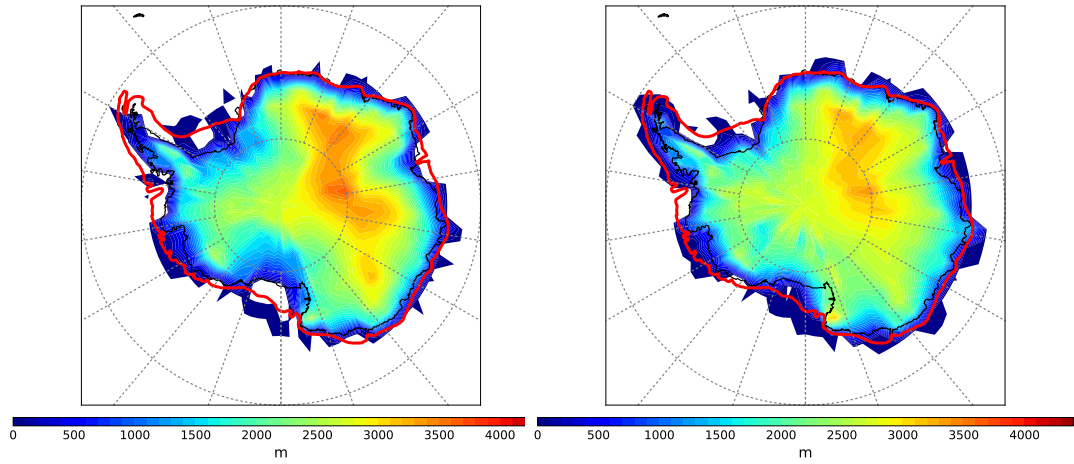
(au)

(av)



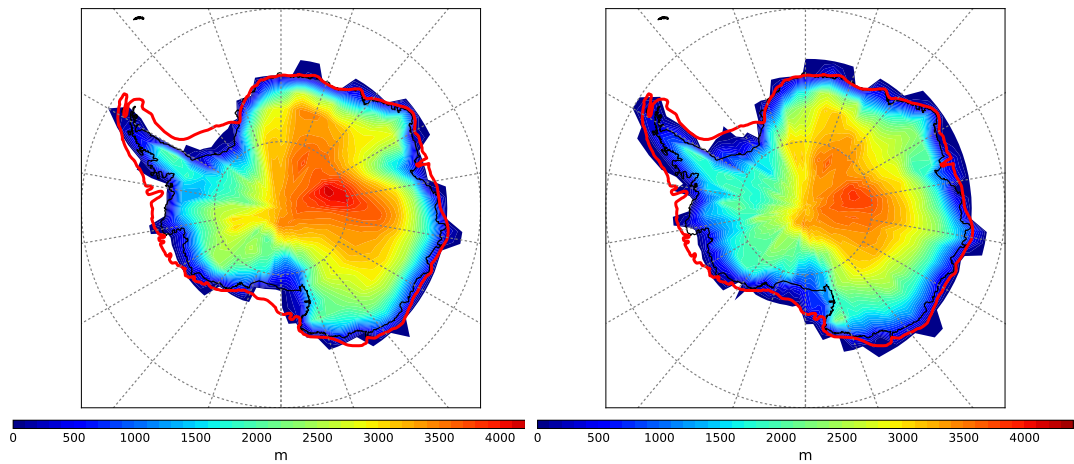
(aw)

(ax)



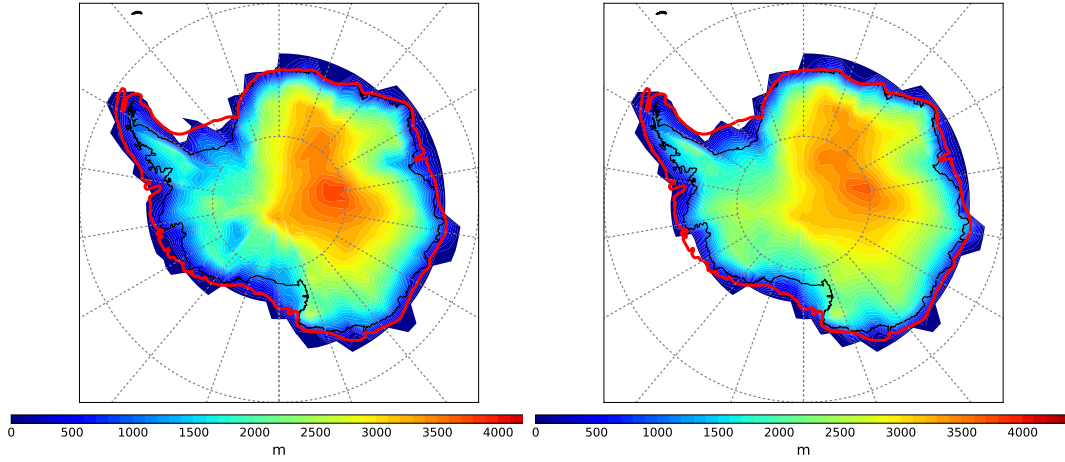
(ay)

(az)



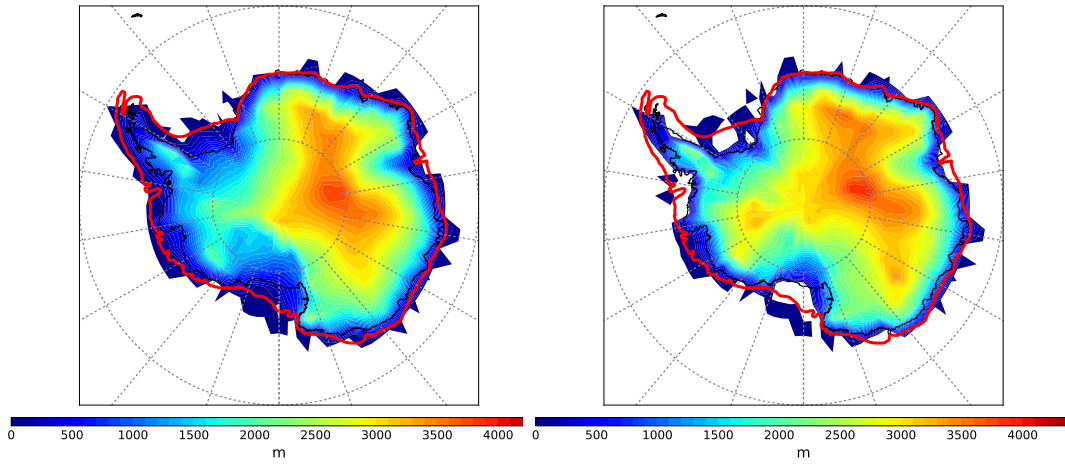
(ba)

(bb)



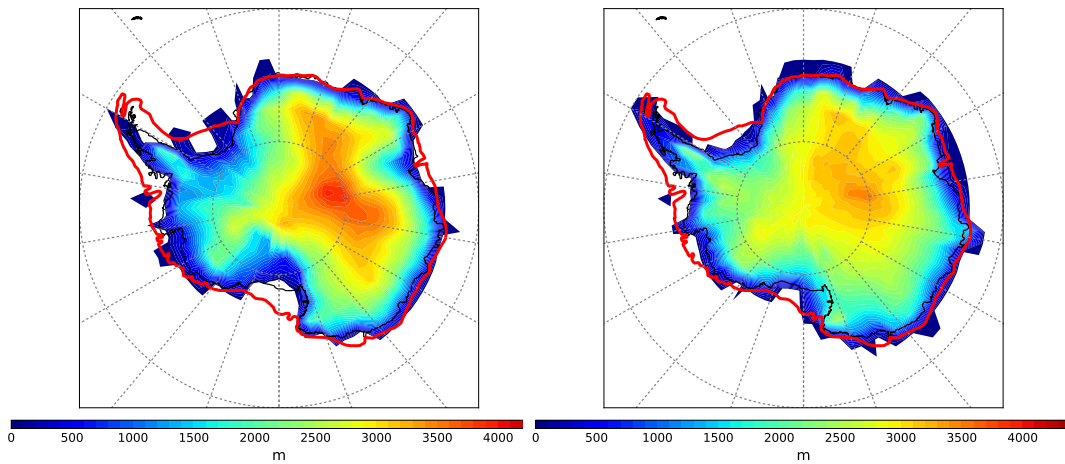
(bc)

(bd)



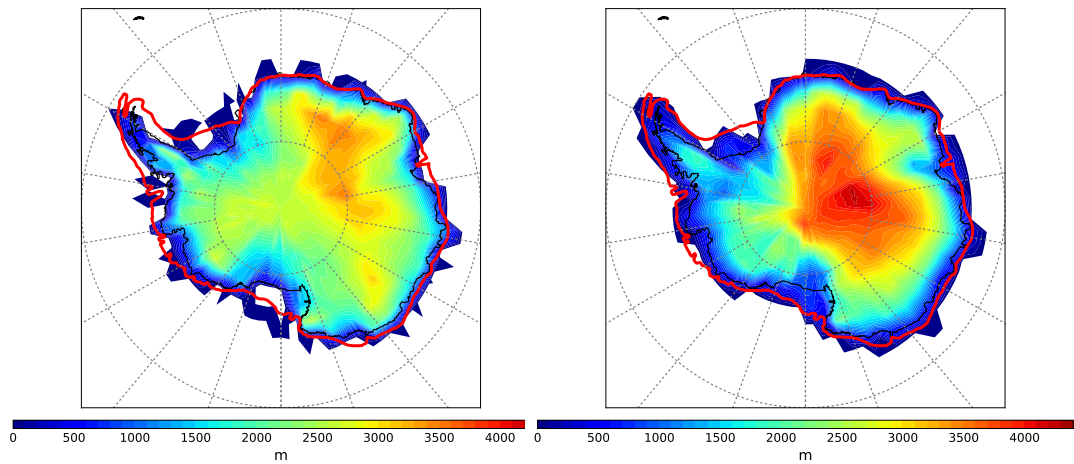
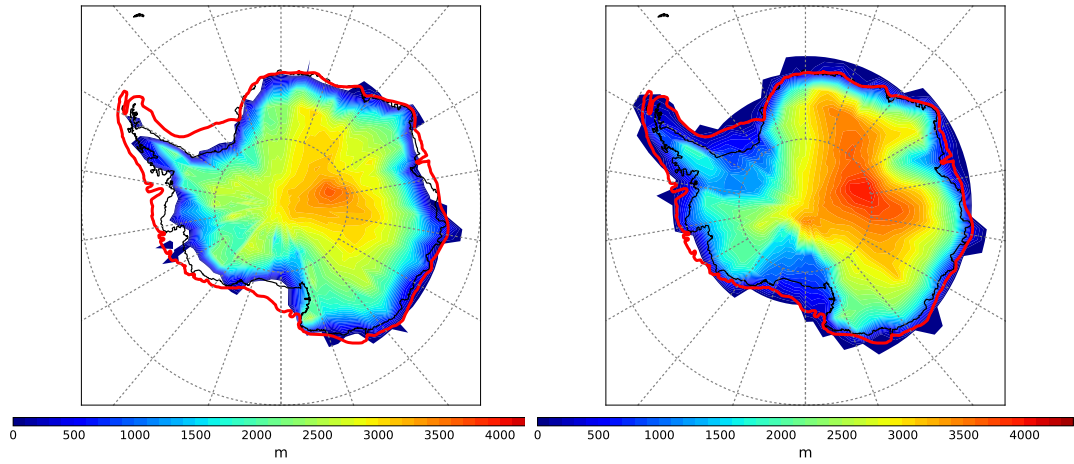
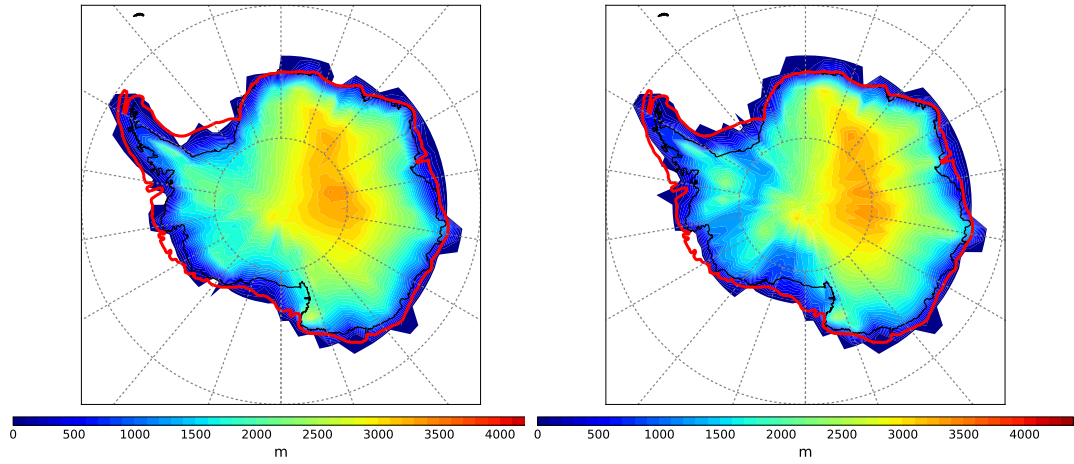
(be)

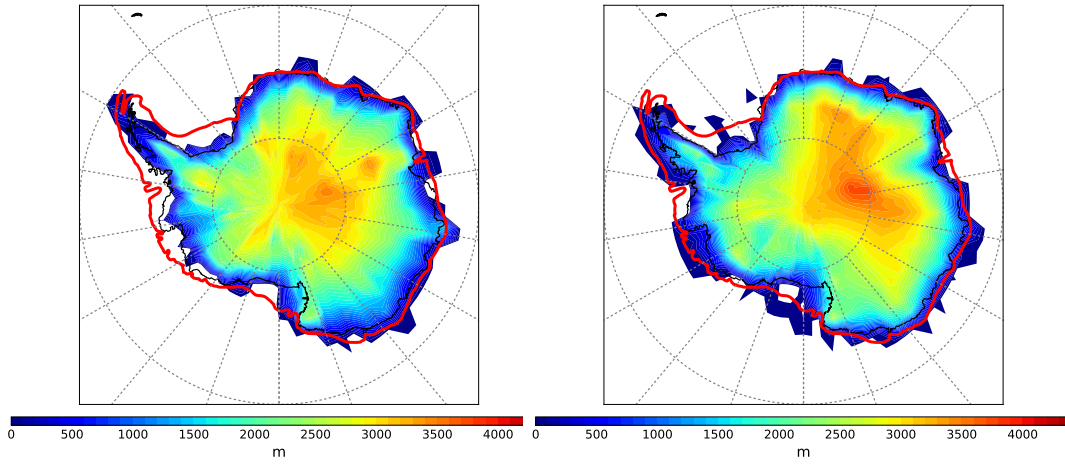
(bf)



(bg)

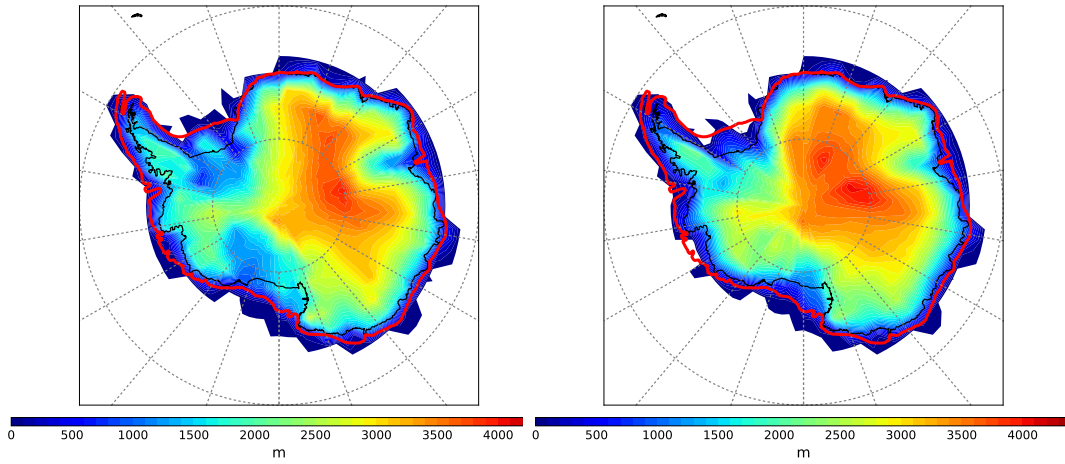
(bh)





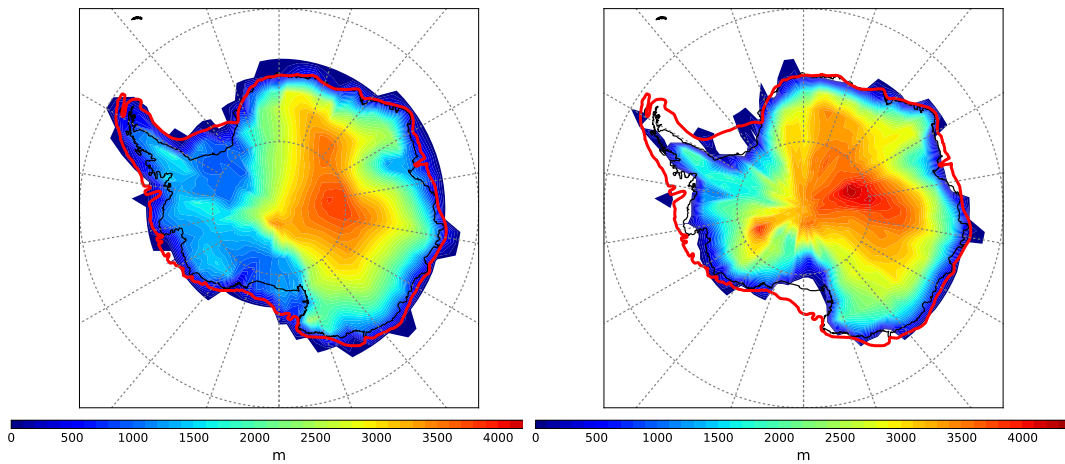
(bo)

(bp)



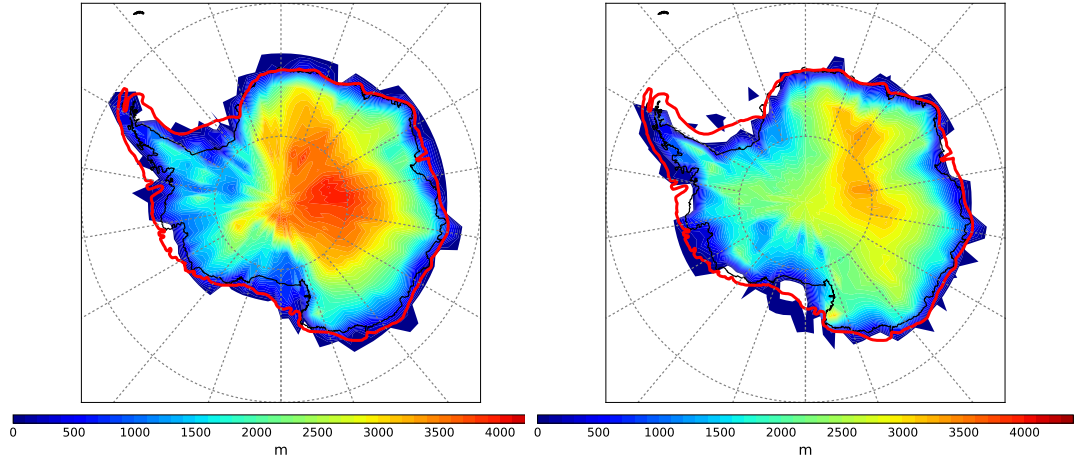
(bq)

(br)



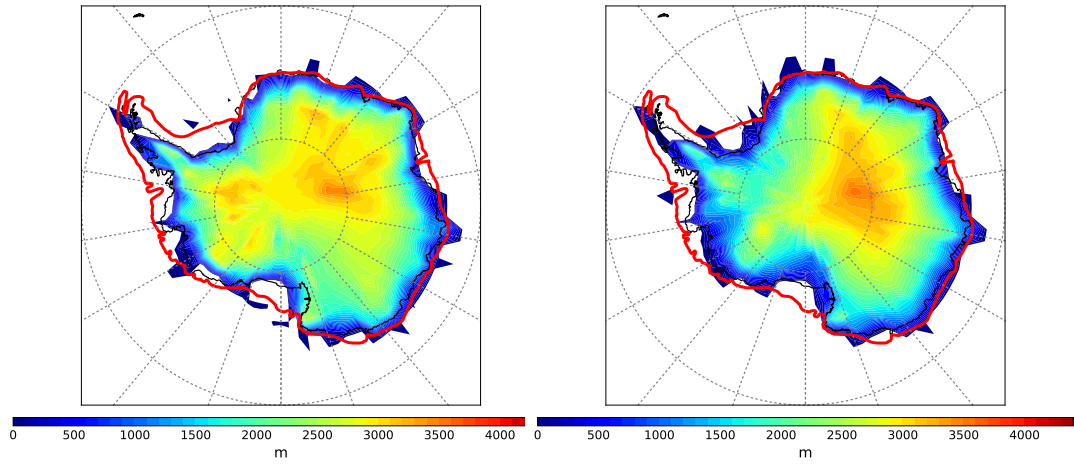
(bs)

(bt)



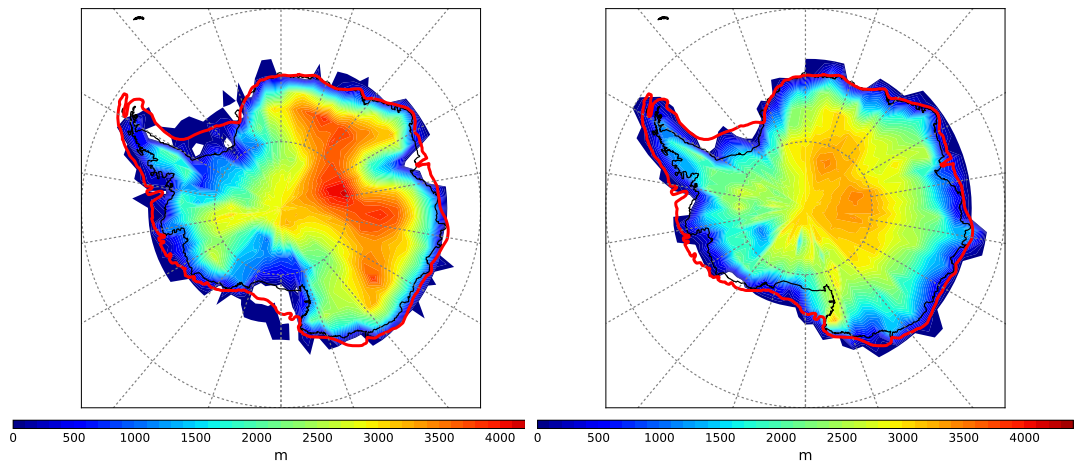
(bu)

(bv)



(bw)

(bx)



(by)

(bz)

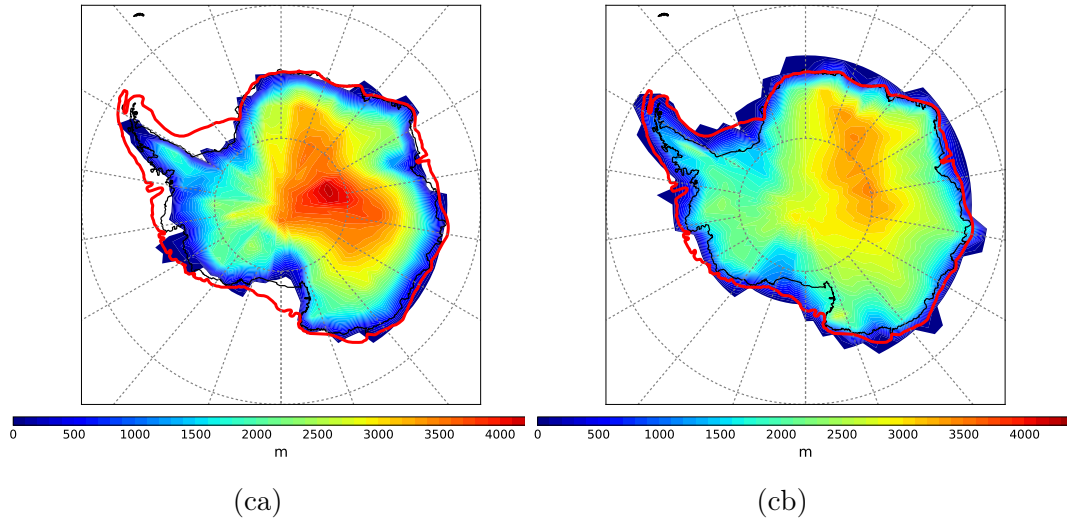


Figure 1

B Post-processing scripts for HadCM3 output.

B.1 data_process_PIrun.py

```
#!/usr/bin/env python2.7
# -----
# Author: Max Holloway, Date: Oct 2016
# Use: Date processing script using iris to extract variables from
# individual UM .pp files (using STASH codes) and merge into single time
# series cube and save as monthly time series in 'expID'/monthly directory
# script use; python /home/users/mholloway/python/data_process.py experiment_name
# -----
import sys
import os
exec(open('/Users/fionaturner/Documents/HADCM3-output/Post-processing-scripts/
startup.py').read())
exec(open('/Users/fionaturner/Documents/HADCM3-output/Post-processing-scripts/
my_defs.py').read())
exec(open('/Users/fionaturner/Documents/HADCM3-output/Post-processing-scripts
/my_cmaps.py').read())
exec(open('/Users/fionaturner/Documents/HADCM3-output/Post-processing-scripts
/my_plots.py').read())
data_root='/Volumes/shared/wilkinson_uq1/User/smq15fet'
# -----
```

```
# get variable stash codes
# pd files:
# water_sublimation_flux_in_timestep m01s03i231
# Evaporation flux from open sea m01s03i232
# air_pressure_at_sea_level m01s16i222
# air_temperature / (K) m01s03i236
# precipitation_flux / (kg m-2 s-1) m01s05i216
# relative_humidity / (%) m01s03i245
# sea_ice_area_fraction / (1) m01s00i031
# sea_ice_thickness / (m) m01s00i032
# snowfall_amount / (kg m-2) m01s00i023
# specific_humidity / (1) m01s03i237
# surface_air_pressure / (Pa) m01s00i001
# surface_temperature / (K) m01s00i024
# surface_upward_latent_heat_flux / (W m-2) m01s03i234
# surface_upward_sensible_heat_flux / (W m-2) m01s03i217
# toa_incoming_shortwave_flux / (W m-2) m01s01i207
# toa_outgoing_longwave_flux / (W m-2) m01s02i205
# toa_outgoing_shortwave_flux / (W m-2) m01s01i208
# wind speed m01s03i230
# y_wind m01s03i226
# x_wind m01s03i225
# 3D fields - pc files
# air_temperature / (K) m01s16i203
# geopotential_height / (m) m01s16i202
# relative_humidity / (%) m01s16i204
```

208

```
# specific_humidity m01s15i226
# x_wind / (m s-1) m01s15i201
# y_wind / (m s-1) m01s15i202
#
# pd_name_list=['water_sublimation_flux','evap_flux,air_pressure_at_sea_level',
'air_temperature', \
#       'precipitation_flux','relative_humidity','sea_ice_area_fraction',
'sea_ice_thickness','snowfall_amount', \
#       'specific_humidity','surface_air_pressure','surface_temperature',
'surface_upward_latent_heat_flux', \
#       'surface_upward_sensible_heat_flux','toa_incoming_shortwave_flux','
toa_outgoing_longwave_flux', \
#       'toa_outgoing_shortwave_flux','wind_speed','y_wind','x_wind']
#
# -----

# pass experiment to be processed
expID=str(sys.argv[1]) # 'xluba'
print( "Processing experiment", expID)

pd_stash_list=['m01s03i231','m01s03i232','m01s16i222','m01s03i236','m01s05i216',
m01s03i245', \
        'm01s00i031','m01s00i032','m01s00i023','m01s03i237','m01s00i001',
        'm01s00i024', \
        'm01s01i207','m01s02i205','m01s01i208','m01s03i230','m01s03i226',
        'm01s03i225']
```

```

for VAR in range(len(pd_stash_list)):
    #filename = iris.sample_data_path(''+data_root+'/'+expID+'/pcpd/'
    #expID+'a@pd*')
    filename = ''+data_root+'/'+expID+'/'+expID+'a.pd*'
    #print(filename)
    stash_constraint = iris.AttributeConstraint
    (STASH=str(pd_stash_list[VAR]))
    #if os.path.getsize(filename) > 0:
    cube = iris.load_cube(filename, stash_constraint)
    cube_name = cube[-1].standard_name
    if cube[-1].standard_name == None:
        cube_name = cube[-1].long_name
    print( "saving", str(cube_name))
    print( str(cube_name), ": start date =", str(cube.coord('time')[0]))
    print( str(cube_name), ": end date =", str(cube.coord('time')[-1]))
    iris.save(cube, ''+data_root+'/'+expID+'/monthly/'+expID+'.'+
    +str(cube_name)+'
    .monthly.nc')#, netcdf_format="NETCDF3_CLASSIC")

# -----
# Isotopes in precipitation are in STASH code 338 -
# levels of the field correspond as
# follows:
# 0=large scale rain 16o
# 1=large scale rain 18o

```


210

```
# 2=large scale rain deuterium
# 3=large scale snow 16o
# 4=large scale snow 18o
# 5=large scale snow deuterium
# 6=convective rain 16o
# 7=convective rain 18o
# 8=convective rain deuterium
# 9=convective snow 16o
# 10=convective snow 18o
# 11=convective snow deuterium
# -----

filename = ''+data_root+'/'+'expID+'/'+'expID+'a.pc*'
stash_constraint = iris.AttributeConstraint(STASH='m01s00i338')
precip_isotopes = iris.load(filename, stash_constraint)
iris.save(precip_isotopes, ''+data_root+'/'+'expID+'/'+'monthly/'+'expID+'.'
precip_isotopes.monthly.nc')

isotopes = ['d160_lsrain', 'd180_lsrain', 'd2H_lsrain', 'd160_lssnow', 'd180_lssnow',
'd2H_lssnow', \
    'd160_cnrain', 'd180_cnrain', 'd2H_cnrain', 'd160_cnsnow', 'd180_cnsnow',
    'd2H_cnsnow'] # Define strings to hold different isotope species
    (ordered as above)

for typ in range(len(isotopes)):
    locals()[isotopes[typ]] = precip_isotopes[typ]
```

```

H2016 = d160_lsrain + d160_cnrain + d160_lssnow + d160_cnsnow
H2018 = d180_lsrain + d180_cnrain + d180_lssnow + d180_cnsnow
d180=(((H2018/H2016)/Rsmow018)-1)*1000) # OXYGEN-18
d180.data[d180.data > 100] = N.nan
d180.data[d180.data < -100] = N.nan
d180.data=N.ma.masked_invalid(d180.data)
d180.var_name,d180.long_name = u'd180',u'd180' # change cube name
print("d180")
iris.save(d180, ''+data_root+'/' +expID+' /monthly/' +expID+'.d180.monthly.nc')

H2 = d2H_lsrain + d2H_cnrain + d2H_lssnow + d2H_cnsnow
d2H=(((H2/H2016)/RvsmowD)-1)*1000) # DEUTERIUM
d2H.data[d2H.data > 100] = N.nan
d2H.data[d2H.data < -1000] = N.nan
d2H.data=N.ma.masked_invalid(d2H.data)
d2H.var_name,d2H.long_name = u'deuterium',u'deuterium' # change cube name
print("deuterium")
iris.save(d2H, ''+data_root+'/' +expID+' /monthly/' +expID+'.deuterium.monthly.nc')

```

B.2 data_process.py

```
#!/usr/bin/env python2.7
```

```
# -----
```

```
# Author: Max Holloway, Date: Oct 2016
# Use: Date processing script using iris to extract variables from
# individual UM .pp files (using STASH codes) and merge into single time
# series cube and save as monthly time series in 'expID'/monthly directory
# script use; python /home/users/mholloway/python/data_process.py experiment_name
# -----
import sys
import os
exec(open('/Users/fionaturner/Documents/HADCM3-output/startup.py').read())
exec(open('/Users/fionaturner/Documents/HADCM3-output/my_defs.py').read())
exec(open('/Users/fionaturner/Documents/HADCM3-output/my_cmaps.py').read())
exec(open('/Users/fionaturner/Documents/HADCM3-output/my_plots.py').read())
data_root='/Volumes/shared/wilkinson_uq1/User/smq15fet'
# -----
# get variable stash codes
# pd files:
# water_sublimation_flux_in_timestep m01s03i231
# Evaporation flux from open sea m01s03i232
# air_pressure_at_sea_level m01s16i222
# air_temperature / (K) m01s03i236
# precipitation_flux / (kg m-2 s-1) m01s05i216
# relative_humidity / (%) m01s03i245
# sea_ice_area_fraction / (1) m01s00i031
# sea_ice_thickness / (m) m01s00i032
# snowfall_amount / (kg m-2) m01s00i023
# specific_humidity / (1) m01s03i237
```

```
# surface_air_pressure / (Pa) m01s00i001
# surface_temperature / (K) m01s00i024
# surface_upward_latent_heat_flux / (W m-2) m01s03i234
# surface_upward_sensible_heat_flux / (W m-2) m01s03i217
# toa_incoming_shortwave_flux / (W m-2) m01s01i207
# toa_outgoing_longwave_flux / (W m-2) m01s02i205
# toa_outgoing_shortwave_flux / (W m-2) m01s01i208
# wind speed m01s03i230
# y_wind m01s03i226
# x_wind m01s03i225
# 3D fields - pc files
# air_temperature / (K) m01s16i203
# geopotential_height / (m) m01s16i202
# relative_humidity / (%) m01s16i204
# specific_humidity m01s15i226
# x_wind / (m s-1) m01s15i201
# y_wind / (m s-1) m01s15i202
#
# pd_name_list=['water_sublimation_flux','evap_flux,air_pressure_at_sea_level',
'air_temperature', \
#       'precipitation_flux','relative_humidity','sea_ice_area_fraction',
'sea_ice_thickness','snowfall_amount', \
#       'specific_humidity','surface_air_pressure','surface_temperature',
'surface_upward_latent_heat_flux', \
#       'surface_upward_sensible_heat_flux','toa_incoming_shortwave_flux',
'toa_outgoing_longwave_flux', \
```

214

```
#      'toa_outgoing_shortwave_flux', 'wind_speed', 'y_wind', 'x_wind']
#
# -----

# pass experiment to be processed
expID=str(sys.argv[1]) # 'xluba'
print( "Processing experiment", expID)

pd_stash_list=['m01s03i231', 'm01s03i232', 'm01s16i222', 'm01s03i236', 'm01s05i216',
'm01s03i245', \
              'm01s00i031', 'm01s00i032', 'm01s00i023', 'm01s03i237', 'm01s00i001',
              'm01s00i024', \
              'm01s01i207', 'm01s02i205', 'm01s01i208', 'm01s03i230', 'm01s03i226',
              'm01s03i225']

for VAR in range(len(pd_stash_list)):
    #filename = iris.sample_data_path(''+data_root+'/' +expID+' /pcpd/' +expID+
    'a@pd*')
    filename = ''+data_root+'/' +expID+' /pcpd/' +expID+'a@pd*'
    #print(filename)
    stash_constraint = iris.AttributeConstraint(STASH=
    str(pd_stash_list[VAR]))
    #if os.path.getsize(filename) > 0:
    cube = iris.load_cube(filename, stash_constraint)
    cube_name = cube[-1].standard_name
    if cube[-1].standard_name == None:
```

```

        cube_name = cube[-1].long_name
    print( "saving", str(cube_name))
    print( str(cube_name), ": start date =", str(cube.coord('time')[0]))
    print( str(cube_name), ": end date =", str(cube.coord('time')[-1]))
    iris.save(cube, ''+data_root+'/' +expID+' /monthly/' +expID+' .
        '+str(cube_name)+' .monthly.nc')#, netcdf_format="NETCDF3_CLASSIC")

# -----
# Isotopes in precipitation are in STASH code 338 - levels of the field
# correspond as follows:
# 0=large scale rain 16o
# 1=large scale rain 18o
# 2=large scale rain deuterium
# 3=large scale snow 16o
# 4=large scale snow 18o
# 5=large scale snow deuterium
# 6=convective rain 16o
# 7=convective rain 18o
# 8=convective rain deuterium
# 9=convective snow 16o
# 10=convective snow 18o
# 11=convective snow deuterium
# -----

filename = ''+data_root+'/' +expID+' /pcpd/' +expID+'a@pc*'
stash_constraint = iris.AttributeConstraint(STASH='m01s00i338')
```

```

precip_isotopes = iris.load(filename, stash_constraint)
iris.save(precip_isotopes, ''+data_root+'/' +expID+' /monthly/' +expID+
'.precip_isotopes.monthly.nc')

isotopes = ['d160_lsrain', 'd180_lsrain', 'd2H_lsrain', 'd160_lssnow', 'd180_lssnow',
'd2H_lssnow', \
    'd160_cnrain', 'd180_cnrain', 'd2H_cnrain', 'd160_cnsnow', 'd180_cnsnow',
    'd2H_cnsnow']
    # Define strings to hold different isotope species (ordered as above)

for typ in range(len(isotopes)):
    locals()[isotopes[typ]] = precip_isotopes[typ]

H2016 = d160_lsrain + d160_cnrain + d160_lssnow + d160_cnsnow
H2018 = d180_lsrain + d180_cnrain + d180_lssnow + d180_cnsnow
d180=(((H2018/H2016)/RsmowO18)-1)*1000) # OXYGEN-18
d180.data[d180.data > 100] = N.nan
d180.data[d180.data < -100] = N.nan
d180.data=N.ma.masked_invalid(d180.data)
d180.var_name,d180.long_name = u'd180',u'd180' # change cube name
print("d180")
iris.save(d180, ''+data_root+'/' +expID+' /monthly/' +expID+' .d180.monthly.nc')

H2 = d2H_lsrain + d2H_cnrain + d2H_lssnow + d2H_cnsnow
d2H=(((H2/H2016)/RvsmowD)-1)*1000) # DEUTERIUM
d2H.data[d2H.data > 100] = N.nan

```

```

d2H.data[d2H.data < -1000] = N.nan
d2H.data=N.ma.masked_invalid(d2H.data)
d2H.var_name,d2H.long_name = u'deuterium',u'deuterium' # change cube name
print("deuterium")
iris.save(d2H, ''+data_root+'/'+'expID+'/'monthly/'+'expID+'.deuterium.monthly.nc')

```

B.3 my_cmmaps.py

```

# CUSTOM COLOURMAPS
#execfile('custom_cmmaps.py')
import matplotlib.colors as mcol
#####
## GLOBALS
##
COLOR1 = '#2C722F' #green
COLOR2 = '#8C201E' #red
COLOR3 = '#224E73' #purple
COLOR4 = '#DEA000' #yellow
COLOR5 = '#078C7A' #cyan
COLOR6 = '#EE1904' #lighter red
COLOR7 = '#32AE1C' # lighter green

COLDARKGREEN = '#007D1C'
COLDARKRED = '#A61000'

```


218

COLDARKBLUE = '#015666'

COLDARKYELLOW = '#A65100'

COLLIGHTGREEN = '#38E05D'

COLLIGHTRED = '#FF5240'

COLLIGHTBLUE = '#37B6CE'

COLLIGHTYELLOW = '#FF9D40'

COLDARKGREEN = '#007D1C'

COLDARKRED = '#9F0025'

COLDARKPURPLE = '#42036F'

COLDARKBLUE = '#06266F'

COLDARKYELLOW = '#A66F00'

COLGREEN = '#1AD644'

COLRED = '#D61A1A'

COLPURPLE = '#D61AAC'

COLBLUE = '#1A4ED6'

COLYELLOW = '#E4EB2F'

COLLIGHTGREEN = '#38E05D'

COLLIGHTRED = '#B12D4C'

COLLIGHTBLUE = '#4671D5'

COLLIGHTYELLOW = '#FFBF40'

COLLIGHTPURPLE = '#963FD5'

```
COLLIST1 = mcol.LinearSegmentedColormap.from_list('hagcolors', [COLDARKGREEN,
COLDARKRED, COLDARKPURPLE, COLDARKBLUE, COLDARKYELLOW])

COLLIST2 = mcol.LinearSegmentedColormap.from_list('hagcolors', [COLDARKPURPLE,
COLDARKBLUE, 'White', 'White', COLDARKYELLOW, COLDARKRED])

COLLIST3 = mcol.LinearSegmentedColormap.from_list('hagcolors', [COLDARKPURPLE,
COLDARKBLUE, 'White', COLDARKYELLOW, COLDARKRED])

COLMAP1 = mpl.colors.LinearSegmentedColormap.from_list(
    'COLMAP1',
    [COLLIGHTGREEN, COLDARKGREEN , COLDARKGREEN ,
    'White', COLLIGHTRED, COLLIGHTRED,
    COLLIGHTPURPLE ], N=200)

COLMAP2 = mpl.colors.LinearSegmentedColormap.from_list(
    'COLMAP1',
    [COLLIGHTGREEN, COLDARKGREEN, 'White', COLDARKRED, COLLIGHTRED], N=200)

#name='hess','milagro'
#'GC','GC_DISCRETE','wigner_cmap'

# -----

COLDARKBLUE = '#015666'
COLLIGHTBLUE = '#37B6CE'
```

220

COLDARKBLUE = '#06266F'

COLBLUE = '#1A4ED6'

COLLIGHTBLUE = '#4671D5'

COLDARKGREEN = '#007D1C'

COLLIGHTGREEN = '#38E05D'

COLDARKGREEN = '#007D1C'

COLGREEN = '#1AD644'

COLLIGHTGREEN = '#38E05D'

COLDARKYELLOW = '#A65100'

COLLIGHTYELLOW = '#FF9D40'

COLDARKYELLOW = '#A66F00'

COLYELLOW = '#E4EB2F'

COLLIGHTYELLOW = '#FFBF40'

COLDARKRED = '#A61000'

COLLIGHTRED = '#FF5240'

COLDARKRED = '#9F0025'

COLRED = '#D61A1A'

COLLIGHTRED = '#B12D4C'

COLDARKPURPLE = '#42036F'

COLPURPLE = '#D61AAC'

COLLIGHTPURPLE = '#963FD5'

```
COLLIST1 = mcol.LinearSegmentedColormap.from_list('hagcolors', [COLDARKBLUE,
COLBLUE, COLLIGHTBLUE, 'White', 'White', #COLDARKGREEN, COLLIGHTGREEN,
#COLDARKGREEN, COLGREEN, COLLIGHTGREEN, COLLIGHTYELLOW,
COLYELLOW, COLDARKYELLOW, COLLIGHTRED, COLRED, COLDARKRED,
COLLIGHTPURPLE#, COLDARKPURPLE
])
```

```
COLLIST2 = mcol.LinearSegmentedColormap.from_list('hagcolors', [COLDARKBLUE,
COLBLUE, COLLIGHTBLUE, 'White', 'White', #COLDARKGREEN, COLLIGHTGREEN,
#COLDARKGREEN, COLGREEN, COLLIGHTGREEN, COLLIGHTYELLOW,
COLYELLOW, COLDARKYELLOW, COLLIGHTPURPLE, COLLIGHTRED, COLRED,
COLDARKRED, #, COLDARKPURPLE
])
```

```
COLLIST3 = mcol.LinearSegmentedColormap.from_list(
'colors', [
'MidnightBlue', 'Navy', 'CornflowerBlue', 'DarkSlateBlue', 'SlateBlue',
'MediumSlateBlue', 'MediumBlue', 'RoyalBlue', 'Blue', 'DodgerBlue',
'DeepSkyBlue', 'SkyBlue', 'LightSkyBlue', 'SteelBlue', 'LightSteelBlue',
'LightBlue', 'PowderBlue', 'DarkTurquoise', 'MediumTurquoise',
'Turquoise', 'Cyan', 'LightCyan',
'White', 'White', 'White', 'White', 'White', 'White', 'White', 'White',
'LightYellow', 'Yellow', 'Gold', 'Goldenrod', 'DarkGoldenrod',
'SandyBrown', 'Tan', 'Chocolate', 'Firebrick', 'Brown',
'DarkSalmon', 'Salmon', 'LightSalmon', 'Orange', 'DarkOrange', 'Coral',
'LightCoral', 'Tomato', 'OrangeRed', 'Red',
```

```

'Violet'#, 'Plum', 'Orchid', 'MediumOrchid', 'DarkOrchid', 'DarkViolet',
'BlueViolet', 'Purple', 'MediumPurple'
])

```

```

COLLIST4 = mcol.LinearSegmentedColormap.from_list(
    'colors', [
        'MidnightBlue', 'Navy', '#CornflowerBlue', '#DarkSlateBlue', 'SlateBlue',
        'MediumSlateBlue', '#MediumBlue', '#RoyalBlue',
        'Blue', '#DodgerBlue', #
        'DeepSkyBlue', '#SkyBlue', '#LightSkyBlue', '#SteelBlue', 'LightSteelBlue',
        '#LightBlue',
        '#PowderBlue', '#DarkTurquoise', 'MediumTurquoise', '#Turquoise',
        'Cyan', '#LightCyan',
        'White', 'White', '#LightYellow',
        'Yellow', 'Gold', 'Orange', '#SandyBrown',
        'Goldenrod', 'DarkGoldenrod', '#Tan',
        'Chocolate', '#Brown', '#DarkSalmon', 'Salmon', 'LightSalmon', '#DarkOrange',
        '#Coral', 'LightCoral', '#Tomato', '#OrangeRed',
        'OrangeRed', '#Tomato',
        'Red', 'Firebrick', '#Violet'#, 'Plum', 'Orchid', 'MediumOrchid', 'DarkOrchid',
        'DarkViolet'#, 'BlueViolet', 'Purple', 'MediumPurple'
    ])

```

```

COLLIST5 = mcol.LinearSegmentedColormap.from_list(
    'colors', [
        'MidnightBlue', '#Navy', '#CornflowerBlue', '#DarkSlateBlue', 'SlateBlue',

```

```

'MediumSlateBlue',
'MediumBlue',#'RoyalBlue',
'Blue', 'DodgerBlue',#'DeepSkyBlue',
'SkyBlue', 'LightSkyBlue',#'SteelBlue', 'LightSteelBlue', 'LightBlue',
'PowderBlue',#
'DarkTurquoise', 'MediumTurquoise', 'Turquoise',
'Cyan',#'LightCyan',
'White', 'White', 'White', 'white',
'LightYellow', 'Yellow', 'Gold', 'Goldenrod',#'DarkGoldenrod',
'Orange', 'DarkOrange', 'SandyBrown',#'Tan',
'Chocolate', 'Firebrick', 'Brown',#'DarkSalmon', 'Salmon', 'LightSalmon',
#'Coral',
'LightCoral',
'Tomato', 'OrangeRed', 'Red',
'Violet',#'Plum', 'Orchid', 'MediumOrchid', 'DarkOrchid',
'DarkViolet',#'BlueViolet', 'Purple', 'MediumPurple'
])

```

```
# colormap for absolute T
```

```
red = N.array([0, 0, 0, 10, 0, 0, 0, 60, 72, 127, 155, 221, 229, 239, 228, 205,
161, 116, 77]) / 256.
```

```
green = N.array([0, 0, 66, 144, 170, 191, 206, 209, 224, 255, 255, 242, 235,
190, 128, 72, 33, 29, 30]) / 256.
```

```
blue = N.array([128, 255, 255, 255, 255, 255, 209, 204, 208, 212, 250, 243,
99, 63, 39, 27, 22, 29, 27]) / 256.
```

```

my_colormap0 = N.array([red, green, blue]).T
# colormap for T anomalies
red = N.array([0, 0, 0, 10, 0, 60, 72, 221, 255,255, 250, 255, 229, 239, 228,
205, 161, 116, 77]) / 256.
green = N.array([0, 0, 66, 144, 191, 209, 255, 242, 255,255, 250, 255, 215,
190, 128, 72, 33, 29, 30]) / 256.
blue = N.array([128, 255, 255, 255, 255, 204, 250, 243, 255, 255, 210, 0, 99,
63, 39, 27, 22, 29, 27]) / 256.
my_colormap = N.array([red, green, blue]).T

```

B.4 my_defs.py

```

RvsmowD=155.76e-6; #\'b10.05 x 10-6 [P. Fritz and J.Ch. Fontes, Handbook of
Environmental Isotope Geochemistry, Vol. 1, pp. 1-19 (1980)].

```

```

#http://deuterium.nist.gov/standards.html

```

```

RsmowO18=0.0020052; #ratio of O18/O16 from http://epswww.unm.edu/facstaff/
zsharp/bio2.htm

```

```

def oxygen18_iris(d160_lsrain,d160_cnrain,d160_lssnow,d160_cnsnow,d180_lsrain,
d180_cnrain,d180_lssnow,d180_cnsnow):
    '''Calculates d180 from 160/180 ls/cn rain/snow'''
    H2016 = d160_lsrain + d160_cnrain + d160_lssnow + d160_cnsnow
    H2018 = d180_lsrain + d180_cnrain + d180_lssnow + d180_cnsnow
    d180=(((H2018/H2016)/RsmowO18)-1)*1000)

```

```

d180.data=N.ma.masked_invalid(d180.data)
return d180[:, -1, :, :]

def deu_iris(d16_lsrain,d16_cnrain,d16_lssnow,d16_cnsnow,d02_lsrain,d02_cnrain,
d02_lssnow,d02_cnsnow):
    '''Calculates deuterium from 16O/02H ls/cn rain/snow'''
    H2016 = d16_lsrain + d16_cnrain + d16_lssnow + d16_cnsnow
    H2 = d02_lsrain + d02_cnrain + d02_lssnow + d02_cnsnow
    d2H=(((H2/H2016)/RvsmowD)-1)*1000)
    d2H.data=N.ma.masked_invalid(d2H.data)
    return d2H[:, -1, :, :]

def month2annpw(precip, isotope):
    years = N.zeros([(len(precip[:])//12),len(precip[0,:]),len(precip[0,0,:])])
    k=0 #counter variable to index array
    for i in range (0,(len(precip[:])//12)):
        isoyr=isotope.data[k:k+12, :, :]
        preyr=precip[k:k+12, :, :]
        A=isoyr.reshape(12, len(isoyr[0,:])*len(isoyr[0,0,:]))
        B=preyr.reshape(12, len(preyr[0,:])*len(preyr[0,0,:]))
        tyrsum=N.zeros(len(A[0,:]))
        for xy in range(len(A[0,:])):
            tyrsum[xy] = (N.sum(B[:,xy]*A[:,xy]))/sum(B[:,xy],0)#(N.sum(preyr))

    tyrsum=N.ma.masked_invalid(tyrsum)
    tyrgrid=tyrsum.reshape(len(isoyr[0,:]),len(isoyr[0,0,:]))

```



```

#reshape back to (x,y)
years[i,:,:]=tyrgrid #add the summed year to the output array
k=k+12 #add '12 months' #month2ann=years

years=isotope[-(len(years)):]+years-isotope[-(len(years)):]
# convert back to iris cube
return years[:]

def precip_weight(precip, isotope):
    tyrgrid = N.zeros([len(precip[0,:]),len(precip[0,0,:])])
    A=isotope.data.reshape(len(isotope.data[:]),
len(isotope.data[0,:])*len(isotope.data[0,0,:]))
    B=precip.reshape(len(precip[:]), len(precip[0,:])*len(precip[0,0,:]))
    tyrsum=N.zeros(len(A[0,:]))
    for xy in range(len(A[0,:])):
        tyrsum[xy] = (N.sum(B[:,xy]*A[:,xy]))/sum(B[:,xy],0)#(N.sum(precip))

    tyrsum=N.ma.masked_invalid(tyrsum)
    tyrgrid=tyrsum.reshape(len(isotope.data[0,:]),len(isotope.data[0,0,:]))
    #reshape back to (x,y)
    tyrgrid=isotope[-1]+tyrgrid-isotope[-1] # convert back to iris cube
    tyrgrid.data=N.ma.masked_invalid(tyrgrid.data)
    return tyrgrid[:]

def load_d180(expID, expName,tlen):
    """load precip weighted d180 - averaged over last 30 years"""

```

```

print( str(expName), ": load precip weighted d180 and d180ann -
averaged over last", str(tlen), "years")

# HadCM3 precipitation -----
cube = iris.load(''+data_root+'/' + expID + '/monthly/' + expID + '.
precip_mm_srf.monthly.nc')
hp = cube[0]
hp = hp[-tlen:,-1,:,:]
hp.data=N.ma.masked_invalid(hp.data)
neutron = ['16','02','18']
raintype = ['lsrain','lssnow','cnrain','cnsnow']
for i in range(len(neutron)):
    neu = neutron[i]
    for j in range(len(raintype)):
        typ = raintype[j]
        cube = iris.load(''+data_root+'/' + expID + '/monthly/' + expID +
'.isotope'+neu+
'_'+typ+'_mm_srf.monthly.nc')
        locals()["d"+str(neutron[i])+"_"+str(raintype[j])] = cube[0]
# OXYGEN-18 -----
d180 = oxygen18_iris(d16_lsrain,d16_cnrain,d16_lssnow,d16_cnsnow,d18_lsrain,
d18_cnrain,d18_lssnow,d18_cnsnow)
d180 = d180[-tlen:,:,:]
# calculate precip weighted d180
d180Pi = precip_weight(hp.data, d180)
d180Pi.data[d180Pi.data > 100] = N.nan
d180Pi.data[d180Pi.data < -100] = N.nan

```

```

d180Pi.data=N.ma.masked_invalid(d180Pi.data)
return d180Pi, month2annpw(hp.data, d180)

def earth_cube(cube):
    '''sorts out coordinates to refer to earth
    (needed for area weighting stuff)''';
    cs = iris.coord_systems.GeogCS(6371229) #define geographic coord system with
    radius of Earth
    for c in ['latitude', 'longitude']: #modify lat and lon coords to use cs
        cube.coord(c).guess_bounds()
        cube.coord(c).coord_system=cs

def movingaverage(interval, window_size):
    window = N.ones(int(window_size))/float(window_size)
    return N.convolve(interval, window, 'same')

def assign_model_txy(var,lat_data,lon_data,lat_model,lon_model):
    model_coresites = N.zeros([len(var[:]),len(lat_data)])
    for core_site in range(len(lat_data)):
        lati = lat_data[core_site]
        loni = lon_data[core_site]
        index1,grid_lat = min(enumerate(lat_model), key=lambda x:
        abs(x[1]-(lati)))
        index2,grid_lon = min(enumerate(lon_model), key=lambda x:
        abs(x[1]-(loni)))
        model_coresites[:,core_site] = var[:,index1,index2]

```

```
    return model_coresites

def define_ann(var):
    '''create month time coord'''
    iris.coord_categorisation.add_year(var, 'time', name='year')
    year = var.coord('year')
    ann_mean = var.aggregated_by(['year'], iris.analysis.MEAN)
    print( repr(ann_mean))
    return year, ann_mean

def define_annt(var):
    '''create month time coord'''
    iris.coord_categorisation.add_year(var, 't', name='year')
    year = var.coord('year')
    ann_mean = var.aggregated_by(['year'], iris.analysis.MEAN)
    print( repr(ann_mean))
    return year, ann_mean

def define_month(var):
    '''create month time coord'''
    iris.coord_categorisation.add_month(var, 'time', name='month')
    month = var.coord('month')
    month_mean = var.aggregated_by(['month'], iris.analysis.MEAN)
    print( repr(month_mean))
    return month, month_mean
```

B.5 my_plots.py

```

def basemap_ais_plot(var, lat, lon, levels, colormap):
    lon0=180
    vmin,vmax=levels.min(),levels.max()
    fig = plt.figure(figsize=(5,5))#cm2inch(9.5), cm2inch(23))
    ax = fig.add_subplot(111)
    plt.tight_layout(pad=1)
    fig.subplots_adjust(bottom=0.15)
    var,lonsout = mpl_toolkits.basemap.shiftgrid(lon0, var.data, lon.points,
    start=False, cyclic=360.0)
    var,lonsout = mpl_toolkits.basemap.addcyclic(var, lonsout)
    map = Basemap(projection='spaeqd',boundinglat=-55,lon_0=180,resolution='l')
    map.drawcoastlines(color = 'grey')
    map.drawparallels(N.arange(-80.,81.,20.), color='grey')
    map.drawmeridians(N.arange(-180.,181.,20.), color='grey')
    map.drawmapboundary(fill_color='white')
    x,y = map(*N.meshgrid(lonsout,lat.points))
    contour_result = map.contourf(x,y,var, levels, cmap=colormap,
    extend='both')
    ax.annotate('A', xy=(0, 1), xycoords='axes fraction',weight='bold',
    fontsize=10,xytext=(5, -5), textcoords='offset points',ha='left', va='top')
    plt1_ax = plt.gca()
    left, bottom, width, height = plt1_ax.get_position().bounds
    first_plot_left = plt1_ax.get_position().bounds[0]

```

```

width = left - first_plot_left + width
colorbar_axes = fig.add_axes([first_plot_left, bottom - 0.06, width, 0.03])
cbar = plt.colorbar(contour_result, colorbar_axes, orientation= 'horizontal')
cbar.set_label(u'$\u2030$')
cbar.ax.tick_params(labelsize=8)

def basemap_gis_plot(var,lat,lon,levels,colormap):
    lon0=180
    vmin,vmax=levels.min(),levels.max()
    fig = plt.figure(figsize=(5,5))#cm2inch(9.5), cm2inch(23))
    ax = fig.add_subplot(111)
    plt.tight_layout(pad=1)
    fig.subplots_adjust(bottom=0.15)
    var,lonsout = mpl_toolkits.basemap.shiftgrid(lon0, var.data, lon.points,
start=False, cyclic=360.0)
    var,lonsout = mpl_toolkits.basemap.addcyclic(var, lonsout)
    m = Basemap(width=4000000,height=4000000,resolution='l',projection='eqdc',
lat_1=50.,lat_2=80,lat_0=70,lon_0=-20.)
    m.drawcoastlines(color = 'grey')
    m.drawparallels(N.arange(-80.,81.,20.))
    m.drawmeridians(N.arange(-180.,181.,20.))
    x,y = m(*N.meshgrid(lonsout,lat.points))
    contour_result = m.contourf(x,y,var, levels=levels, cmap=colormap,
extend='both')
    ax.annotate('A', xy=(0, 1), xycoords='axes fraction',weight='bold',
fontsize=10,xytext=(5, -5), textcoords='offset points',ha='left', va='top')

```

```

plt1_ax = plt.gca()
left, bottom, width, height = plt1_ax.get_position().bounds
first_plot_left = plt1_ax.get_position().bounds[0]
width = left - first_plot_left + width
colorbar_axes = fig.add_axes([first_plot_left, bottom - 0.06, width, 0.03])
cbar = plt.colorbar(contour_result, colorbar_axes, orientation= 'horizontal')
cbar.set_label(u'$\u2030$')
cbar.ax.tick_params(labelsize=8)

def basemap_gis_plot2(var,lat,lon,levels,colormap):
    lon0=180
    lat_0=73.; lon_0=-40.
    vmin,vmax=levels.min(),levels.max()
    fig = plt.figure(figsize=(5,5))#cm2inch(9.5), cm2inch(23))
    ax = fig.add_subplot(111)
    plt.tight_layout(pad=1)
    fig.subplots_adjust(bottom=0.15)
    var,lonsout = mpl_toolkits.basemap.shiftgrid(lon0, var.data, lon.points,
start=False, cyclic=360.0)
    var,lonsout = mpl_toolkits.basemap.addcyclic(var, lonsout)
    m = Basemap(projection='ortho',lat_0=lat_0,lon_0=lon_0,resolution=None)
    width = m.urcrnrx - m.llcrnrx
    height = m.urcrnry - m.llcrnry
    coef = 0.25
    width = width*coef
    height = height*coef

```

```

m = Basemap(projection='ortho',lon_0=lon_0,lat_0=lat_0,resolution='l',
llcrnrx=-0.5*width,llcrnry=-0.5*height,urcrnrx=0.5*width,urcrnry=0.5*height)
m.drawcoastlines(color = 'grey')
m.drawcountries()
m.drawparallels(np.arange(-90.,120.,30.))
m.drawmeridians(np.arange(0.,360.,60.))
m.drawmapboundary()
x,y = m(*N.meshgrid(lonsout,lat.points))
contour_result = m.contourf(x,y,var,levels,cmap=colormap, extend='both')
ax.annotate('A', xy=(0, 1), xycoords='axes fraction',weight='bold',
fontsize=10,xytext=(5, -5), textcoords='offset points',ha='left', va='top')
plt1_ax = plt.gca()
left, bottom, width, height = plt1_ax.get_position().bounds
first_plot_left = plt1_ax.get_position().bounds[0]
width = left - first_plot_left + width
colorbar_axes = fig.add_axes([first_plot_left, bottom - 0.06, width, 0.03])
cbar = plt.colorbar(contour_result, colorbar_axes, orientation= 'horizontal')
cbar.set_label(u'$\u2030$')
cbar.ax.tick_params(labelsize=8)

def plt_smoothed_timeseries(var, varsmo, color, cutoff):
    almost_black = '#262626'
    fig = plt.figure(figsize=(8,5)#cm2inch(10), cm2inch(15))
    #(cm2inch(11.5), cm2inch(9.5))#(8, 8))
    ax = fig.add_subplot(1,1,1)
    plt.tight_layout()

```



```

xii=N.arange(len(var))
plt.plot(xii[cutoff:-cutoff],varsmo[cutoff:-cutoff],c=color,linewidth=1.5)
plt.scatter(xii,var,marker='o',edgecolor=color,c=color,alpha=0.1)
ax.annotate('A', xy=(0, 1), xycoords='axes fraction',weight='bold',
fontsize=10,xytext=(5, -5), textcoords='offset points',ha='left', va='top')
plt.gca().set_xlim(xmax=xii.max(),xmin=xii.min())

```

B.6 process_annual.py

```

#!/usr/bin/env python2.7
# -----
# Author: Max Holloway, Date (last modified): Mar 2017
# Use: Date processing script using iris to load (processed) monthly time
# series files (created using 'data_process.py' and stored in
# 'expID'/monthly directory), annual mean data and save new cube in
# 'expID'/annual directory.
# script use; python /home/users/mholloway/python/process_annual.py
experiment_name
# -----
import sys
exec(open('/Users/fionaturner/Documents/HADCM3-output/Post-processing-scripts
/startup.py').read())
exec(open('/Users/fionaturner/Documents/HADCM3-output/Post-processing-scripts
/my_defs.py').read())

```

```
exec(open('/Users/fionaturner/Documents/HADCM3-output/Post-processing-scripts
/my_cmaps.py').read())
exec(open('/Users/fionaturner/Documents/HADCM3-output/Post-processing-scripts
/my_plots.py').read())

data_root='/Volumes/shared/wilkinson_uq1/User/smq15fet'
expID=str(sys.argv[1]) # pass experiment to be processed
print("Processing experiment", expID)
# -----
variables=['sea_ice_area_fraction', 'air_pressure_at_sea_level',
'air_temperature', 'precipitation_flux', 'surface_air_pressure', \
'surface_temperature', 'wind_speed', 'y_wind', 'x_wind']

climatology_months=50*12 # calculate annual mean climatologies over final
50 years

for VAR in range(len(variables)): # Calculate and save annual mean variables
    """calculate annual mean values"""
    cube = iris.load(''+data_root+'/' +expID+' /monthly/' +expID+' .
'+variables[VAR]+' .monthly.nc')
    cube_name = cube[-1].standard_name
    print("calculate annual mean", str(cube_name), "for final",
    climatology_months/12, "years")
    yr, cube_ann = define_ann(cube[0][:-climatology_months:])
    #locals()[str(cube_name)] = cube[0]
    locals()[str(cube_name)+'_ann'] = cube_ann
```

```

iris.save(cube_ann, ''+data_root+'/' +expID+' /annual/' +expID+'.' +cube_name+'
.annual.nc')

# -----
#climatology_months=50*12
ISO=['d180', 'deuterium']
pf = iris.load(''+data_root+'/' +expID+' /monthly/' +expID+' .precipitation_flux.
monthly.nc')
# Load Precipitation
for II in range(len(ISO)): # Calculate and save precip-weighted and
annual mean isotopes
    cube = iris.load(''+data_root+'/' +expID+' /monthly/' +expID+'.' +ISO[II]+' .
monthly.nc')
    if len(cube[0].data) < len(pf[0].data): # check that precip and isotopes are
the same length
        print("do final time coords match?")
        print(pf[0][ -len(cube[0].data):].coord('time')[ -1] == cube[0].
coord('time')[ -1])
        # check that last time coords match
        precip=pf[0][ -len(cube[0].data):].copy() # align time coords and
remove extra dt's
    else:
        precip=pf[0].copy()
    print("calculate precipitation-weighted and annual mean", str(ISO[II]),
"for final", climatology_months/12, "years")
    isotope_cube = cube[0].copy()

```

```

yr, cube_ann = define_ann(isotope_cube.copy()) # define annual time coords
to pass
to annual mean isotope cubes
print("calculate climatology over final", climatology_months/12, "years")
isoPi = precip_weight(precip[-climatology_months:].data, cube[0]
[-climatology_months:])
isoPi.units = None # remove precip units from isotope cube
isoPi.var_name,isoPi.long_name = str(ISO[II]),str(ISO[II])
# change cube name
locals()[ISO[II]+"_Pi"] = isoPi # allocate variable for dxs calculation
iso_ann = month2annpw(precip[-climatology_months:].data, cube[0]
[-climatology_months:])
isoANN = cube_ann[-len(iso_ann.data):].copy()
# hack time units from months into years!
isoANN.data = iso_ann.data # hack time units from months into years!
isoANN.units = None # remove precip units from isotope cube
isoANN.var_name,isoANN.long_name = str(ISO[II]),str(ISO[II])
# change cube name
locals()[ISO[II]+"_ann"] = isoANN # allocate variable for dxs calculation
iris.save(isoPi, ''+data_root+'/' +expID+' /annual/' +expID+'.'+ISO[II]+'
_pw.climate.nc')
iris.save(isoANN, ''+data_root+'/' +expID+' /annual/' +expID+'.'+ISO[II]+'
_pw.annual.nc')

print( "calculate precipitation-weighted and annual mean d-excess")
dxs_Pw = deuterium_Pi[:] - 8*d180_Pi[:] # (Dansgaard,1964)

```

```

dxs_Pw.var_name,dxs_Pw.long_name = u'deuterium-excess',u'deuterium-excess'
# change cube name
iris.save(dxs_Pw, ''+data_root+'/'+'expID+'/'annual/'+'expID+'.dxs_pw.climate.nc')
dxs_ann = deuterium_ann[:] - 8*d180_ann[:]
dxs_ann.var_name,dxs_ann.long_name = u'deuterium-excess',u'deuterium-excess'
# change cube name
iris.save(dxs_ann, ''+data_root+'/'+'expID+'/'annual/'+'expID+'.dxs_pw.annual.nc')

```

B.7 process_monthly_means.py

```

#!/usr/bin/env python2.7
# -----
# Author: Max Holloway, Date (last modified): Mar 2017
# Use: Date processing script using iris to load (processed) monthly time
# series files (created using 'data_process.py' and stored in
# 'expID'/monthly directory), monthly mean data (to generate cube of size
# t=12) and save new cube in 'expID'/monthly directory.
# script use; python /home/users/mholloway/python/process_monthly_means.py
experiment_name
# -----
import sys
exec(open('/Users/fionaturner/Documents/HADCM3-output/Post-processing-scripts
/startup.py').read())
exec(open('/Users/fionaturner/Documents/HADCM3-output/Post-processing-scripts

```

```
/my_defs.py').read()
exec(open('/Users/fionaturner/Documents/HADCM3-output/Post-processing-scripts
/my_cmaps.py').read())
exec(open('/Users/fionaturner/Documents/HADCM3-output/Post-processing-scripts
/my_plots.py').read())

data_root='/Volumes/shared/wilkinson_uq1/User/smq15fet'
expID=str(sys.argv[1]) # pass experiment to be processed
print( "Processing experiment", expID)
# -----
climatology=50*12 # final 50 years used for monthly mean climatologies
variables=['sea_ice_area_fraction', 'air_pressure_at_sea_level',
'air_temperature', 'precipitation_flux', 'surface_air_pressure', \
'surface_temperature', 'wind_speed', 'y_wind', 'x_wind', 'd180', 'deuterium']

for VAR in range(len(variables)): # Calculate and save monthly averages
    """calculate monthly mean values"""
    cube = iris.load(''+data_root+''+expID+'/monthly/'+expID+'.
'+variables[VAR]+' .monthly.nc')
    cube_name = variables[VAR]
    print( "calculate monthly mean", str(cube_name), "over final",
    climatology/12, "years")
    month, cube_month = define_month(cube[0][:-climatology:])
    iris.save(cube_month, ''+data_root+''+expID+'/monthly/'+expID+'.
'+str(cube_name)+' .month_mean.nc')
```

```

## when doing it for d180 and deuterium, must first run
# iris.FUTURE.netcdf_promote = True
# iris.FUTURE.netcdf_no_unlimited = True
## then remember to change file name in directory before running again

#iris.FUTURE.netcdf_promote = True
#iris.FUTURE.netcdf_no_unlimited = True

#cube = iris.load(''+data_root+'/' +expID+' /monthly/' +expID+'.' +variables[9]+'
.monthly.nc')
#cube_name = cube[-1].standard_name
#print( "calculate monthly mean", str(cube_name), "over final", climatology/12,
"years")
#month, cube_month = define_month(cube[0] [-climatology:])
#iris.save(cube_month, ''+data_root+'/' +expID+' /monthly/' +expID+'.'
+str(cube_name)+' .month_mean.nc')

```

B.8 startup.py

```

#!/usr/bin/env python2.7
#execfile('Documents/HADCM3-output/startup.py')
import os
import sys

```

```
import time

import matplotlib as mpl
import mpl_toolkits
#mpl.use('Qt4Agg')
#mpl.use('gtkagg')
import matplotlib.pyplot as plt
import matplotlib.mlab as mlab
import numpy as N
import scipy.io.netcdf as S
from scipy.io.netcdf import NetCDFFile
import scipy as sp
from scipy import interpolate
from scipy import stats
import mpl_toolkits.basemap as bm
from mpl_toolkits.basemap import Basemap
import matplotlib.font_manager as fm
from pylab import *
from matplotlib.colors import LogNorm
import netCDF4
from netCDF4 import Dataset
import iris
import iris.plot as iplt
import iris.quickplot as qplt
import iris.analysis.cartography
import iris.coord_categorisation
```


242

```
#import iris.cube.Cube.interpolate()
```

```
import cartopy.crs as ccrs
```

```
import cartopy.feature as cfeature
```

```
data_root='/Volumes/shared/wilkinson_uq1/User/smq15fet'
```

Bibliography

- Abrahamsen, E. (2012). *Oceanographic conditions beneath Fimbul Ice Shelf, Antarctica*. PhD thesis, University of Southampton.
- Abram, N. J., Mulvaney, R., Wolff, E. W., Triest, J., Kipfstuhl, S., Trusel, L. D., Vimeux, F., Fleet, L., and Arrowsmith, C. (2013). Acceleration of snow melt in an antarctic peninsula ice core during the twentieth century. *Nature Geosci*, 6(5):404–411.
- Albert, I., Donnet, S., Guihenneuc-Jouyaux, C., Low-Choy, S., Mengersen, K., and Rousseau, J. (2012). Combining expert opinions in prior elicitation. *Bayesian Anal.*, 7(3):503–532.
- Argus, D. F., Peltier, W. R., Drummond, R., and 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.
- Basu, K., Saha, A., and Chatterjee, S. (2017). *Large-Scale Quadratically Constrained Quadratic Program via Low-Discrepancy Sequences*. Curran Associates, Inc.
- Bentley, M. J., Cofaigh, C. Á., 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., Crosta, X., Dickens, W. A., Domack, E., Dowdeswell, J. A., Dunbar, R., Ehrmann, W., Evans, J., Favier, V., Fink, D., Fogwill, C. J., Glasser, N. F., Gohl, K., Golledge, N. R., Goodwin, I., Gore, D. B., Greenwood, S. L., Hall, B. L., Hall, K., Hedding, D. W., Hein, A. S., Hocking, E. P., Jakobsson, M., Johnson, J. S., Jomelli, V., Jones, R. S., Klages, J. P., Kristoffersen, Y., Kuhn, G., Leventer, A., Licht, K., Lilly, K., Lindow, J., Livingstone, S. J., Masson, G., McGlone, M. S., McKay, R. M., Melles, M., Miura,

- H., Mulvaney, R., Nel, W., Nitsche, F. O., O'Brien, P. E., Post, A. L., Roberts, S. J., Saunders, K. M., Selkirk, P. M., Simms, A. R., Spiegel, C., Stollendorf, T. D., Sugden, D. E., van der Putten, N., van Ommen, T., Verfaillie, D., Vyverman, W., Wagner, B., White, D. A., Witus, A. E., and 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. Reconstruction of Antarctic Ice Sheet Deglaciation (RAISED).
- Blunier, T. and Brook, E. J. (2001). Timing of millennial-scale climate change in antarctica and greenland during the last glacial period. *Science*, 291(5501):109–112.
- Bracegirdle, T. J. and Stephenson, D. B. (2012). Higher precision estimates of regional polar warming by ensemble regression of climate model projections. *Climate Dynamics*, 39(12):2805–2821.
- Briggs, R. D., Pollard, D., and Tarasov, L. (2014). A data-constrained large ensemble analysis of antarctic evolution since the eemian. *Quaternary Science Reviews*, 103:91 – 115.
- Brook, E. J., White, J. W., Schilla, A. S., Bender, M. L., Barnett, B., Severinghaus, J. P., Taylor, K. C., Alley, R. B., and 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):1333 – 1343.
- Chang, W., Applegate, P. J., Haran, M., and Keller, K. (2014a). Probabilistic calibration of a greenland ice sheet model using spatially resolved synthetic observations: toward projections of ice mass loss with uncertainties. *Geoscientific Model Development*, 7(5):1933–1943.
- Chang, W., Haran, M., Applegate, P., and Pollard, D. (2016a). Calibrating an ice sheet model using high-dimensional binary spatial data. *Journal of the American Statistical Association*, 111(513):57–72.
- Chang, W., Haran, M., Applegate, P., and Pollard, D. (2016b). Improving ice sheet model calibration using paleoclimate and modern data. *ArXiv e-prints*.
- Chang, W., Haran, M., Olson, R., and Keller, K. (2014b). Fast dimension-reduced climate model calibration and the effect of data aggregation. *Ann. Appl. Stat.*, 8(2):649–673.
- Clark, P. U., Dyke, A. S., Shakun, J. D., Carlson, A. E., Clark, J., Wohlfarth, B.,

- Mitrovica, J. X., Hostetler, S. W., and McCabe, A. M. (2009). The last glacial maximum. *Science*, 325(5941):710–714.
- Clem, K. R., Fogt, R. L., Turner, J., Lintner, B. R., Marshall, G. J., Miller, J. R., and Renwick, J. A. (2020). Record warming at the south pole during the past three decades. *Nature Climate Change*, 10(8):762–770.
- Cuffey, K. M., Clow, G. D., Steig, E. J., Buizert, C., Fudge, T., Koutnik, M., Waddington, E. D., Alley, R. B., and Severinghaus, J. P. (2016). Deglacial temperature history of west antarctica. *Proceedings of the National Academy of Sciences*, 113(50):14249–14254.
- D’Andrilli, J., Foreman, C. M., Sigl, M., Priscu, J. C., and McConnell, J. R. (2017). A 21 000-year record of fluorescent organic matter markers in the wais divide ice core. *Climate of the Past*, 13(5):533–544.
- de Boer, B., Haywood, A. M., Dolan, A. M., Hunter, S. J., and 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.
- Dias, L. C., Morton, A., and Quigley, J. (2018). Elicitation. *Springer International Publishing*. MR3700912. doi: <https://doi.org/10.1007/978-3-319-65052-4>, 1(2):3.
- Domingo, D., Malmierca-Vallet, I., Sime, L., Voss, J., and 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. e2019JF005237 10.1029/2019JF005237.
- EFSA (2014). Guidance on expert knowledge elicitation in food and feed safety risk assessment. *EFSA Journal*, 12(6):3734.
- Einstein, A. (1916). The foundation of the general theory of relativity. *Annalen Phys.*, 49(7):769–822.
- 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., Frenzel, A., Fritzsche, D., Fundel, F., Gabrielli, P., Gaspari, V., Gersonde, R., Graf, W., Grigoriev, D., Hamann, I., Hansson, M., Hoffmann, G., Hutterli, M. A., Huybrechts, P., Isaksson, E., Johnsen, S., Jouzel, J., Kaczmarek, M., Karlin, T., Kaufmann, P., Kipfstuhl, S., Kohno, M., Lambert, F., Lambrecht, A., Landais, A., Lawer, G., Leuenberger, M., Littot, G., Loulergue, L., L’Aijthi, D., Maggi, V., Marino, F., Masson-Delmotte,

- V., Meyer, H., Miller, H., Mulvaney, R., Narcisi, B., Oerlemans, J., Oerter, H., Parrenin, F., Petit, J. R., Raisbeck, G., Raynaud, D., RÅúthlisberger, R., Ruth, U., Rybak, O., Severi, M., Schmitt, J., Schwander, J., Siegenthaler, U., Siggaard-Andersen, M. L., Spahni, R., Steffensen, J. P., Stenni, B., Stocker, T. F., Tison, J. L., Traversi, R., Udisti, R., Valero-Delgado, F., van den Broeke, M. R., van de Wal, R. S. W., Wagenbach, D., Wegner, A., Weiler, K., Wilhelms, F., Winther, J. G., and Wolff, E. (2006). One-to-one coupling of glacial climate variability in greenland and antarctica. *Nature*, 444(7116):195 – 198.
- Fogt, R. L., Jones, J. M., Goergens, C. A., Jones, M. E., Witte, G. A., and Lee, M. Y. (2016). Antarctic station-based seasonal pressure reconstructions since 1905: 2. variability and trends during the twentieth century. *Journal of Geophysical Research: Atmospheres*, 121(6):2836–2856. 2015JD024565.
- Garthwaite, P. H., Kadane, J. B., and O’Hagan, A. (2005). Statistical methods for eliciting probability distributions. *Journal of the American Statistical Association*, 100(470):680–701.
- Geyer, C. J. (1992). Practical markov chain monte carlo. *Statistical Science*, 7(4):473–483.
- Golledge, N. R., Fogwill, C. J., Mackintosh, A. N., and 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.
- 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., and Hall, B. L. (2013). Glaciology and geological signature of the last glacial maximum antarctic ice sheet. *Quaternary Science Reviews*, 78(Supplement C):225 – 247.
- Grootes, P., Steig, E., Stuiver, M., Waddington, E., and Morse, D. (1999). Gisp2_ - taylor dome oxygen isotope ratios. *Quat Res*, 283:1712–1714.
- Haywood, A., Dowsett, H., and Dolan, A. (2016). Integrating geological archives and climate models for the mid-pliocene warm period. *Nature Communications*, 7:10646 EP–.
- Hoffman, M. D. and Gelman, A. (2011). The no-u-turn sampler: Adaptively setting path lengths in hamiltonian monte carlo. *ArXiv e-prints*.
- Holloway, M. D., Sime, L. C., Singarayer, J. S., Tindall, J. C., and Valdes, P. J. (2016). Reconstructing paleosalinity from $\delta^{18}\text{o}$: coupled model simulations of the

- last glacial maximum, last interglacial and late holocene. *Quaternary Science Reviews*, 131:350–364.
- IMBIE (2018). Mass balance of the antarctic ice sheet from 1992 to 2017. *Nature*, 558:219–222.
- Johnson, S. R., Tomlinson, G. A., Hawker, G. A., Granton, J. T., Grosbein, H. A., and Feldman, B. M. (2010). A valid and reliable belief elicitation method for bayesian priors. *Journal of Clinical Epidemiology*, 63(4):370 – 383.
- Jolliffe, I. T. (1986). *Principal Components in Regression Analysis*, pages 129–155. Springer New York, New York, NY.
- Jones, J. M., Gille, S. T., Goosse, H., Abram, N. J., Canziani, P. O., Charman, D. J., Clem, K. R., Crosta, X., de Lavergne, C., Eisenman, I., England, M. H., Fogt, R. L., Frankcombe, L. M., Marshall, G. J., Masson-Delmotte, V., Morrison, A. K., Orsi, A. J., Raphael, M. N., Renwick, J. A., Schneider, D. P., Simpkins, G. R., Steig, E. J., Stenni, B., Swingedouw, D., and Vance, T. R. (2016). Assessing recent trends in high-latitude southern hemisphere surface climate. *Nature Climate Change*, 6(10):917–926.
- Joseph, V. R. and Hung, Y. (2008). Orthogonal-maximin latin hypercube designs. *Statistica Sinica*, 18(1):171–186.
- Jouzel, J., Delaygue, G., Landais, A., Masson-Delmotte, V., Risi, C., and Vimeux, F. (2013). Water isotopes as tools to document oceanic sources of precipitation. *Water Resources Research*, 49(11):7469–7486.
- 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., and Watanabe, O. (2007). Northern hemisphere forcing of climatic cycles in antarctica over the past 360,000 years. *Nature*, 448(7156):912–916.
- Keller, E. D., Baisden, W. T., Bertler, N. A., Emanuelsson, B. D., Canessa, S., and Phillips, A. (2018). Calculating uncertainty for the rice ice core continuous flow analysis water isotope record. *Atmospheric Measurement Techniques*, 11(8):4725–4736.
- Kennedy, M. C. and O’Hagan, A. (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(3):425.

- Kopp, R. E., Simons, F. J., Mitrovica, J. X., Maloof, A. C., and Oppenheimer, M. (2009). Probabilistic assessment of sea level during the last interglacial stage. *Nature*, 462(7275):863–867.
- Landais, A., Barkan, E., and Luz, B. (2008). Record of $\delta^{18}\text{O}$ and 17O -excess in ice from Vostok Antarctica during the last 150,000 years. *Geophysical Research Letters*, 35(2).
- Loeppky, J. L., Sacks, J., and Welch, W. J. (2009). Choosing the sample size of a computer experiment: A practical guide. *Technometrics*, 51(4):366–376.
- Malmierca-Vallet, I., Sime, L. C., Tindall, J. C., Capron, E., Valdes, P. J., Vinther, B. M., and Holloway, M. D. (2018). Simulating the last interglacial Greenland stable water isotope peak: The role of arctic sea ice changes. *Quaternary Science Reviews*, 198:1–14.
- McInerney, G. J., Chen, M., Freeman, R., Gavaghan, D., Meyer, M., Rowland, F., Spiegelhalter, D. J., Stefaner, M., Tessarolo, G., and Hortal, J. (2014). Information visualisation for science and policy: engaging users and avoiding bias. *Trends in Ecology & Evolution*, 29(3):148 – 157.
- Met-Office (2016). Met office climate prediction model: Hadcm3. <http://www.metoffice.gov.uk/research/modelling-systems/unified-model/climate-models/hadcm3>. Accessed: 27-06-2017.
- Notz, D. (2015). How well must climate models agree with observations? *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 373(2052).
- O’Hagan, A., Buck, C. E., Daneshkhah, A., Eiser, J. R., Garthwaite, P. H., Jenkinson, D. J., Oakley, J. E., and Rakow, T. (2006). *Uncertain Judgements: Eliciting Experts’ Probabilities*. Wiley.
- Parnell, A. C. (2005). *The statistical analysis of former sea level*. PhD thesis, University of Sheffield.
- Peltier, W. (2004). Global glacial isostasy and the surface of the ice-age earth: the ice-5g (vm2) model and GRACE. *Annu. Rev. Earth Planet Sci.*, 32(1):111–149.
- Peltier, W. R. (1994). Ice age paleotopography. *Science*, 265(5169):195–201.
- Peltier, W. R., Argus, D. F., and Drummond, R. (2015). Space geodesy constrains ice age terminal deglaciation: The global ice-6g_c (vm5a) model. *Journal of Geophysical Research: Solid Earth*, 120(1):450–487.

- 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., PÄLpin, L., Ritz, C., Saltzman, E., and Stievenard, M. (1999). Climate and atmospheric history of the past 420,000 years from the vostok ice core, antarctica. *Nature*, 399(6735):429–436.
- Pollard, D. and DeConto, R. M. (2009). Modelling west antarctic ice sheet growth and collapse through the past five million years. *Nature*, 458:329–333.
- 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. Copyright - Database copyright ProQuest LLC; ProQuest does not claim copyright in the individual underlying works; Last updated - 2019-10-17.
- Pörtner, H.-O., Roberts, D., Masson-Delmotte, V., Zhai, P., Tignor, M., Poloczanska, E., Mintenbeck, K., Alegría, A., Nicolai, M., Okem, A., Petzold, J., Rama, B., and Weyer, N. (2019). *IPCC, 2019: Summary for Policymakers. In: IPCC Special Report on the Ocean and Cryosphere in a Changing Climate*. In press.
- Pronzato, L. and Müller, W. G. (2012). Design of computer experiments: space filling and beyond. *Statistics and Computing*, 22(3):681–701.
- Salzmann, U., Haywood, A. M., and Lunt, D. J. (2009). The past is a guide to the future? comparing middle pliocene vegetation with predicted biome distributions for the twenty-first century. *Philosophical Transactions: Mathematical, Physical and Engineering Sciences*, 367(1886):189–204.
- Sandhu, R., Pettit, C., Khalil, M., Poirel, D., and Sarkar, A. (2017). Bayesian model selection using automatic relevance determination for nonlinear dynamical systems. *Computer Methods in Applied Mechanics and Engineering*, 320:237 – 260.
- Schneider, D. P., Steig, E. J., van Ommen, T. D., Dixon, D. A., Mayewski, P. A., Jones, J. M., and Bitz, C. M. (2006). Antarctic temperatures over the past two centuries from ice cores. *Geophysical Research Letters*, 33(16):n/a–n/a. L16707.
- Screen, J. A., Bracegirdle, T. J., and Simmonds, I. (2018). Polar climate change as manifest in atmospheric circulation. *Current Climate Change Reports*, 4(4):383–395.
- Sime, L. C., Wolff, E. W., Oliver, K. I. C., and Tindall, J. C. (2009). Evidence for warmer interglacials in east antarctic ice cores. *Nature*, 462(7271):342–345.

- Simms, A. R., Lisiecki, L., Gebbie, G., Whitehouse, P. L., and Clark, J. F. (2019). Balancing the last glacial maximum (lgm) sea-level budget. *Quaternary Science Reviews*, 205:143–153.
- Steig, E. J., Ding, Q., White, J. W. C., K&Aijttel, 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., Burgener, L., Williams, J., and 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.
- Steig, E. J., Morse, D. L., Waddington, E. D., Stuiver, M., Grootes, P. M., Mayewski, P. A., Twickler, M. S., and 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.
- 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., Landais, A., Lemieux-Dudon, B., Maggi, V., Masson-Delmotte, V., Mazzola, C., Minster, B., Montagnat, M., Mulvaney, R., Narcisi, B., Oerter, H., Parrenin, F., Petit, J. R., Ritz, C., Scarchilli, C., Schilt, A., Sch&Aijpbach, S., Schwander, J., Selmo, E., Severi, M., Stocker, T. F., and Udisti, R. (2011). Expression of the bipolar see-saw in antarctic climate records during the last deglaciation. *Nature Geoscience*, 4(1):46–49.
- Stocker, T., Qin, D., Plattner, G.-K., Tignor, M., Allen, S., Boschung, J., Nauels, A., Xia, Y., Bex, V., and Midgley, P. (2013). *IPCC, 2013: Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press.
- Thomas, E. R., Bracegirdle, T. J., Turner, J., and Wolff, E. W. (2013). A 308 year record of climate variability in west antarctica. *Geophysical Research Letters*, 40(20):5492–5496. 2013GL057782.
- Thomas, E. R., Dennis, P. F., Bracegirdle, T. J., and Franzke, C. (2009). Ice core evidence for significant 100-year regional warming on the antarctic peninsula. *Geophysical Research Letters*, 36(20):n/a–n/a. L20704.
- Tindall, J. C., Valdes, P. J., and Sime, L. C. (2009). Stable water isotopes in hadcm3:

- Isotopic signature of the southern oscillation and the tropical amount effect. *Journal of Geophysical Research: Atmospheres*, 114(D4):n/a–n/a. D04111.
- Turner, J., Lu, H., White, I., King, J. C., Phillips, T., Hosking, J. S., Bracegirdle, T. J., Marshall, G. J., Mulvaney, R., and Deb, P. (2016). Absence of 21st century warming on antarctic peninsula consistent with natural variability. *Nature*, 535(7612):411–415.
- 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., Fegyveresi, J. M., Ferris, D., Fitzpatrick, J. J., Johnson, J., Hargreaves, G., Lee, J. E., Maselli, O. J., Mason, W., McGwire, K. C., Mitchell, L. E., Mortensen, N., Neff, P., Orsi, A. J., Popp, T. J., Schauer, A. J., Severinghaus, J. P., Sigl, M., Spencer, M. K., Vaughn, B. H., Voigt, D. E., Waddington, E. D., Wang, X., and Wong, G. J. (2013). Onset of deglacial warming in west antarctica driven by local orbital forcing. *Nature*, 500(7463):440–444.
- Werner, M., Jouzel, J., Masson-Delmotte, V., and Lohmann, G. (2018). Reconciling glacial antarctic water stable isotopes with ice sheet topography and the isotopic paleothermometer. *Nature communications*, 9(1):1–10.
- Whitehouse, P. L., Bentley, M. J., Milne, G. A., King, M. A., and 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.