

Bayesian analysis of rapid climate change during the last glacial using Greenland $\delta^{18}\text{O}$ data

D. Peavoy¹ and C. Franzke²

¹Complexity Science Centre, University of Warwick, Coventry, UK

²British Antarctic Survey, Cambridge, UK

Received: 12 May 2010 – Published in Clim. Past Discuss.: 15 June 2010

Revised: 25 November 2010 – Accepted: 5 December 2010 – Published: 14 December 2010

Abstract. We present statistical methods to determine climate regimes for the last glacial period using three temperature proxy records from Greenland: measurements of $\delta^{18}\text{O}$ from the Greenland Ice Sheet Project 2 (GISP2), the Greenland Ice Core Project (GRIP) and the North Greenland Ice Core Project (NGRIP) using different timescales. A Markov Chain Monte Carlo method is presented to infer the number of states in a latent variable model along with their associated parameters. By using Bayesian model comparison methods we find that a model with 3 states is sufficient. These states correspond to a gradual cooling during the Greenland Interstadials, more rapid temperature decrease into Greenland Stadial and to the sudden rebound temperature increase at the onset of Greenland Interstadials. We investigate the recurrence properties of the onset of Greenland Interstadials and find no evidence to reject the null hypothesis of randomly timed events.

1 Introduction

Measurements of $\delta^{18}\text{O}$ from Greenland ice cores for the last glacial are a good proxy for regional temperature. These time series indicate large temperature changes occurring on time scales of centuries or less. They are characterised by a rapid warming of up to 16°C , followed by a slow cooling period. They are known as Dansgaard-Oeschger (DO) events and were numbered and dated by Dansgaard et al. (1993) using a time scale derived from a simple ice flow model and constrained by well established events at 11.5 kyr and 110 kyr BP (before present). The DO events correspond to the onset of a

Greenland Interstadial (GIS), the warm phase of the glacial period. An alternation between GIS and the colder Greenland Stadial (GS) will be referred to as a DO cycle.

There is evidence that these dramatic climate shifts at high latitudes in the Northern Hemisphere had global scale effects. For example, Blunier et al. (1998) observed an out of phase signal in Antarctic ice cores, whereas Wang et al. (2001) note general agreement of temperature changes in GISP2 and changes in the hydrological cycle from stalagmite records from the Hulu cave in China. A new common time scale of Greenland/Antarctica temperature back to Marine Isotope Stage 5 will allow more accurate tests of modelling ability to simulate the global impacts of these events (Capron et al., 2010).

Bond et al. (1993) used the abundance of planktic foraminifera in sediment records to show that sea surface temperatures in the North Atlantic (NA) are correlated with DO cycles. They also note that there are longer periods of cooling, which each include several DO cycles. At the termination of these 10–15 kyr cycles there are signatures of large scale ice rafting events that correlate with a GS. Known as Heinrich events, they are associated with collapse of the Laurentide ice sheet, thereby inundating the NA with freshwater, preventing the formation of North Atlantic Deep Water (NADW) and reducing the thermohaline circulation. The subsequent depletion of icebergs in the NA is thought to provoke a rapid reorganisation of the climate system and a temperature increase into a prominent GIS followed by gradual cooling punctuated by a series of DO cycles. This does not explain the mechanism driving the shorter DO cycles, although they are considered to be linked with similar ice sheet calving affecting the salinity of the NA and the formation of NADW.



Correspondence to: D. Peavoy
(d.peavoy@warwick.ac.uk)

The apparent regular occurrence of DO events has led some authors to suggest an external forcing or internal oscillation of the climate system. In particular the Fourier spectrum of the temperature record using the GISP2 time line has a significant peak at 1470 years (Grootes and Stuiver, 1997). However, the $\delta^{18}\text{O}$ record from the Greenland Ice Core Project GRIP, completed the same year as GISP2, whilst also exhibiting DO cycles does not have this spectral peak. Neither does that from the North Greenland Ice Core Project (NGRIP) using the GICC05 time scale.

Alley et al. (2001) suggest a simple model for the dynamics that includes weak periodic forcing and a stochastic resonance mechanism. Stocker and Johnsen (2003) link this to a bipolar thermal “see-saw”, explaining the observed response in Antarctica and Braun et al. (2005) propose a resonance between two faster solar modes of variability to give a resulting weak signal at 1470 years.

It has been argued that Fourier methods are not suitable to detect recurrence patterns for such non-linear discrete events that are likely generated by a thresholding process. Rahmstorf (2003) measures the deviation of the events from a perfect 1470 years signal and finds that their timing is well within the assumed dating error, although there are several 1470 year periods where no event occurs at all. Ditlevsen et al. (2007) assess the significance of the observed periodicity by defining null models and computing the Rayleigh measure and standard deviation of residuals for ensembles of surrogate data. They find no evidence to reject the null model of randomly timed events. The exception is when using the GISP2 ice core record and omitting DO event 9 as numbered by Dansgaard et al. (1993). However, this does not change the result when using the more recent NGRIP ice core record.

The event detection algorithm of Rahmstorf (2003), which uses a criteria of 2‰ $\delta^{18}\text{O}$ increase over 200 years, does not include DO 9, potentially influencing the conclusion. Ditlevsen et al. (2005) defines a DO event as first up crossings of an upper level following a lower level, identifying several more events and concluding that they are randomly paced. Given the sensitivity of results regarding the existence of periodicity to the definition of the events it is desirable to have a method that objectively detects their occurrence.

Kwasniok and Lohmann (2009) show that a stochastic double well model could be fitted using the unscented Kalman filter to capture stadial-interstadial transitions, while Ditlevsen (1999) argues that non-Gaussian “heavy tailed” noise is suitable to capture extreme events within Greenland paleoclimate data. Livina et al. (2010) used GRIP and NGRIP $\delta^{18}\text{O}$ data to study the number of states in the climate for the last 60 kyrs using a polynomial fitting algorithm to windows of the data. They detect the two states corresponding to the stadial and interstadials of the last glacial and find that these merge to a single state around 25 kyrs BP. Although these methods allow for asymmetry between climate states they do not include the characteristic “saw tooth” features of DO temperature cycles.




Fig. 1: ($\delta^{18}\text{O}$) from NGRIP, GRIP and GISP2 expressed in ‰ with respect to Vienna Standard Mean Ocean Water using GICC05, ss09sea and Meese/Sowers time scales respectively.

The data sets we study are shown in Fig. 1. The NGRIP core is dated using GICC05 (Andersen et al., 2006), GRIP chronology from the ss09sea time scale (Johnsen et al., 2001) and GISP2 according to the Meese/Sowers time scale (Meese et al., 1997). We consider the period 82–11.0 kyr BP for GRIP and GISP2, and 60–11.0 because of limited resolution at older dates.

The aim of this paper is to study models that are able to capture the features of the cycles (sudden temperature changes and small fluctuations) while also including the associated uncertainty in the identification of climate states. This uncertainty can then be included in subsequent analysis. We consider there to be hidden states corresponding to different climate regimes and identify these in a probabilistic sense. We generate an ensemble of data sets derived from NGRIP, GRIP or GISP2 using Gaussian Process regression. This incorporates variability into the data and reflects our uncertainty about the dating and measuring process while retaining the characteristic features of the DO cycles. It will also provide a means of assessing how robust our inference method is.

We apply a Bayesian statistical analysis to each ensemble member, which makes no assumptions about the parameters (e.g. mean and variance of the temperature change in a given state). Instead these are computed from the posterior distribution using Bayes’ rule. We discuss Gaussian Process Regression in Sect. 2 and the modelling strategy in Sect. 3. In Sect. 4 we analyse the results from the posterior simulation and find that models with three climate regimes fit the data well. In Sect. 4 we analyse the timing of the model state corresponding to the DO events.

2 Gaussian process regression

We consider discrete time models and so require data at regular intervals. An interpolation method such as cubic splines will smooth the data removing any jumps or discontinuities.




Fig. 2: Gaussian Process (GP) fit to the GISP2 data set. The raw data is shown in red with the mean from the GP in blue and the upper and lower 3σ levels in grey.

Since it is the jumps we are interested in we choose to use Gaussian process regression (GPR), which can be tuned to capture the “roughness” of the data. GPR is a linear least squares estimation algorithm for interpolation where the sample path is considered to be a random process. It also provides uncertainty estimates of the interpolated sample path.

The method relies upon the specification of a covariance function between data points. We use the Matern class with smoothness parameter set equal to 1 (see Rasmussen and Williams, 2006). Then points r apart have covariance $K(r) = \exp(-r/l)$, where l is the length scale. This is then equivalent to a first order autoregressive (AR(1)) covariance, generating continuous but non-differentiable time series. The length scale, process noise and observation noise are set using the maximum likelihood software of Rasmussen and Williams (2006).

The method generates a distribution over processes as shown in Fig. 2 for GISP2. It is clear that the uncertainty increases as we go back in time but decreases around the observed values. The “saw tooth” feature of the DO cycles is preserved.

We also apply the method to all three data sets and obtain an ensemble of 100 members from the Gaussian process distribution for each ice core, each with a time step of 50 years.

3 Latent state models

We assume that there are M unobserved (latent) states governing the temperature. Rather than consider states corresponding to absolute temperature, for example, a high temperature interstadial and a cooler stadial, we model the temperature increments of the series. For example, there may be one state for the sudden temperature increase at the onset of a Greenland Interstadial (GIS) and another for the slow cooling during the GIS. In this way we aim to capture the asymmetry between warming and cooling phases of the DO cycles. Considering only the increments implies that the method will not

be affected by longer time scale variations in the data such as the Milankovitch cycles.

The parameters of the model include the mean temperature change per unit time and the variance within each state. These are not fixed but are inferred directly from the data using Bayes’ rule. We analyse models with up to five hidden states. The occurrence of each state and all of the associated parameters are estimated in a Bayesian framework from the posterior distribution. We also determine the most suitable model for the data by estimating their marginal likelihoods. This involves approximating the integral over the whole parameter space to obtain an estimate of the probability of the model given the data. In this way, models with too many parameters, that over fit the data, are penalised.

We assume that the time series can be modelled as a random walk, dependent upon the state of the climate. This is intended to capture the differing distributions of the increments that occur during different regimes. Note that one could use a more general first order process, however we found that this did not significantly affect the identification of climate states.

Under the random walk model the data X_i at time i evolves according to

$$X_{i+1} = X_i + \mu_{S_i} + \sigma_{S_i} \epsilon_i, \quad (1)$$

where $\epsilon_i \sim \mathcal{N}(0, 1)$ is the standard normal distribution, μ_{S_i} is the forcing and σ_{S_i} is the standard deviation in state S_i . For N data points, the likelihood function of the data given the unobserved state sequence is

$$P(X|\mu, \sigma, S) = \prod_{i=0}^{N-1} \mathcal{N}(X_{i+1} - X_i - \mu_{S_i}, \sigma_{S_i}^2). \quad (2)$$

To capture our uncertainty about the time spent in each state an extra layer is introduced into the hierarchy that conditions S on an unknown parameter λ_j , $j \in \{0 \dots M\}$. This is the probability that S_i is in state j , implying a multinomial distribution for S_i .

In Bayesian modelling one considers the unknown parameters as random variables and aims to compute their probability distribution conditional upon the observed data. For generic parameter θ and data X the posterior distribution is proportional to the product of likelihood function and prior distribution $P(\theta|X) \propto P(X|\theta)P(\theta)$. The prior distribution can incorporate any previous knowledge about θ or it can be designed to be as “uninformative” as possible.

In our problem λ_j , μ_j and σ_j for $j \in \{0 \dots M\}$ are random variables which need a prior distribution. We use conjugate priors, which means that the posterior distributions for a given parameter will be of the same family as the priors. The full posterior will be intractable but using conjugate priors implies that one can sample each parameter from their conditional posterior separately. One then cycles through the conditional posterior distributions to obtain a sample from the full joint distribution. This is a Gibbs sampling algorithm

and will generate dependent samples from the full posterior distribution after discarding some initial convergence period (for an introduction to Bayesian inference using Markov Chain Monte Carlo (MCMC) methods see Gilks et al., 1995).

The conjugate prior for λ_j is given by a Dirichlet distribution with hyperparameter α . This respects the constraint $\sum_j \lambda_j = 1$. For σ_j one defines $\tau_j = 1/\sigma_j^2$ and assigns a Gamma prior $\tau_j \sim \Gamma(a, b)$, whereas μ_j is a priori normally distributed $\mu_j \sim \mathcal{N}(0, \sigma_\mu^2)$. The hyperparameters a, b and σ_μ are fixed. The posterior distributions will then be Dirichlet, Normal and Gamma respectively. Their particular form is given in the appendix.

To estimate the number of states we compute the probability of the model given the data, i.e the marginal likelihood. For model M , the marginal likelihood is

$$P(X|M) = \sum_S \int P(X|\mu, \tau, S) P(S|\lambda) P(\mu, \tau, \lambda|M) d\mu d\tau d\lambda. \quad (3)$$

To estimate this integral/sum one could simply sample from the prior and compute an average of the likelihood function for these values. However, this is likely to be a poor estimator due to the discrepancy between prior and posterior distributions (Kass and Raftery, 1995). Another option would be to use the Harmonic Mean estimator. This is calculated using samples from the posterior so is more accurate although it does not obey a Gaussian central limit theorem and so exhibits instabilities due to the occasional sample with very low likelihood value (Kass and Raftery, 1995).

We choose instead to use the Laplace-Metropolis estimator of Raftery (1996). This is an adaptation of the classic Laplace method, which is known to provide more efficient estimates than those based on posterior simulation. However, it is not always applicable since it requires the availability of the posterior mode and covariance matrix. Instead one can estimate these quantities from the posterior simulation output.

As the hidden states are conditionally independent given θ the above sum can be performed analytically and each integral approximated separately

$$P(X|M) \approx \sum_S (2\pi)^{d/2} |\Psi|^{1/2} P(X|\hat{\mu}, \hat{\tau}, S) P(S|\hat{\lambda}) P(\hat{\mu}, \hat{\tau}, \hat{\lambda}|M), \quad (4)$$

where $\hat{\mu}$, $\hat{\tau}$ and $\hat{\lambda}$ is the posterior mode and Ψ is the covariance matrix computed from the MCMC output.

4 Estimation of climate states

In all cases we use uninformative priors with large variance. For example $\mu \sim \mathcal{N}(0, 100)$. Repeated simulations reveal that the results are not sensitive to the prior. To determine the model most supported by the data we estimate marginal likelihoods using the method described above. The logarithm of the marginal likelihoods is computed from each of the 100 ensemble members of the Gaussian Process regression using

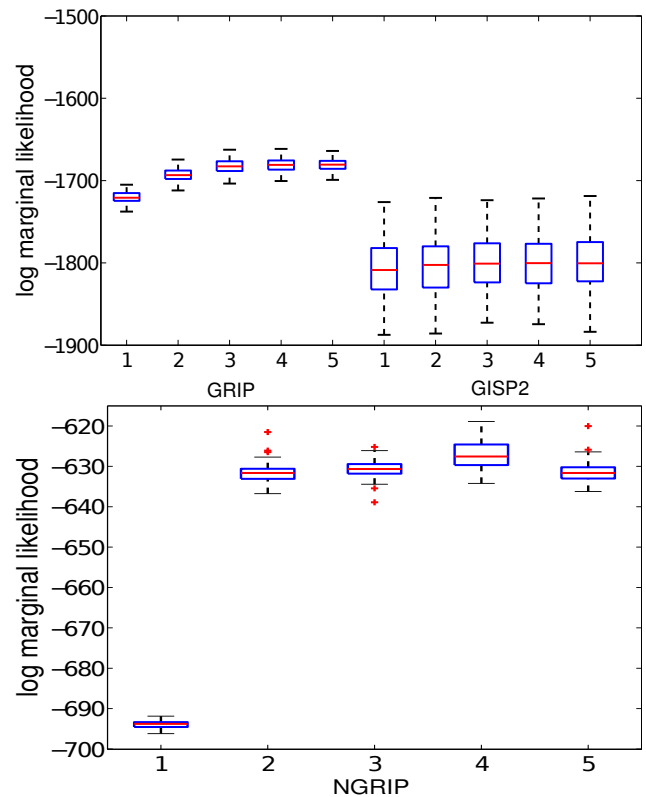


Fig. 3: Marginal likelihoods for the ensemble derived from each data set for models with 1–5 states. The red line indicates the median of the ensemble, the box is the 25 and 75 percentiles, the whiskers are the maximum and minimum of the distribution while the red markers indicate outliers.

10 000 samples from the posterior. The results are shown in boxplot form in Fig. 3.

There is a wider spread in results for GISP2 than GRIP or NGRIP due to the broad distribution of sampled series shown in Fig. 2. The values for NGRIP are much higher due to the smaller data set. If we consider the Bayes Factors BF_{ij} for model i versus j for the medians alone we find that for GRIP and GISP2 the three state model is supported. For GRIP we have $BF_{32} = 4.1 \times 10^4$ regarded as strong by the scale of Kass and Raftery (1995), whereas for GISP2 we have $BF_{32} = 5.2$, regarded as substantial evidence. For NGRIP we have $BF_{43} = 22$, favouring four states. In the further analysis we work with the three state model as no significant extra modelling ability is given by including more than three states for two out of the three data sets.

Using the three state model we find that one state $j = 1$, corresponds very closely to the DO events. The state $j = 0$ occurs during stadial and interstadial periods and corresponds to small fluctuations. There is also some agreement between the temperature decrease state $j = 2$, and the termination of GIS. However, this does not occur for all DO cycles so we focus our attention upon the state corresponding to

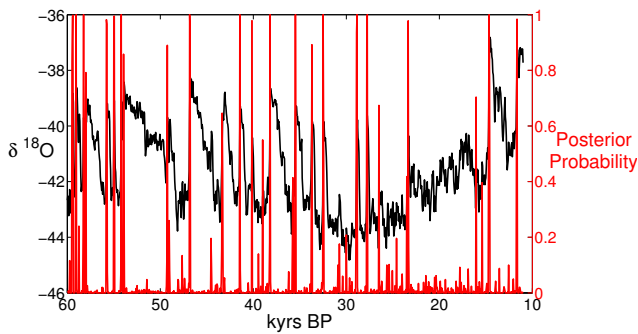



Fig. 4: Posterior probability of DO event as a function of time for NGRIP on the GICC05 timescale.


DO events. Figure 4 shows the posterior probability of a DO event for the NGRIP data, i.e. for each time i the probability that $S_i = 1$. The posterior distribution of hidden states is similar for the other data sets, although more events are detected in the GISP2 data (not shown). The events numbered by Dansgaard et al. (1993), all have high posterior probability, implying that the model is able to distinguish these events. There are, however several other regions of high probability, for example at 26.5 kyr BP. Although not typically regarded as DO events in the literature, these jumps in temperature share similar statistics and may be important in the analysis of recurrence patterns. Many of the high probability regions cluster together and would previously have been classed as a single event. Here, we only regard two events as one if they occur in immediate succession. In this way we make little restriction on the minimum time period between events.

For each data set we retained the hidden state sequence with the largest posterior probability and calculated its distribution of waiting times. These are shown in Fig. 7 along with their best fit exponential distribution. The 95% confidence interval is computed according to the bounds derived in Massart (1990). The theoretical distribution is within this bound for all three data sets implying that we can not reject the model of a Poisson process with no periodicity.


The posterior distributions for the parameter of interest μ are shown in Fig. 6. Note that this parameter has units of $\delta^{18}\text{O}\text{‰}$ and can be considered to be proportional to the local temperature change within a 50 year period for each climate state. There is a clear distinction in this forcing between different climate states. μ_0 , corresponding to the relaxation state, has a mean of $-0.05 - 0.0$ for the three data sets. The termination state has mean -2.2 to -1.7 . There is good agreement for the DO state between the three data sets with a mean of approximately 2.0. This corresponds to a rate of change of 0.04‰ per year. This is larger than the 0.01‰ per year of Rahmstorf (2003), although our model allows noise in the state which would allow a DO event to have a slower temperature change.



(a) NGRIP, 1320



(b) GRIP, 2700



(c) GISP2, 2900

Fig. 5: Empirical cumulative waiting times (red) with the best fit exponential distribution (black) and the 95% confidence interval. Also given are the mean waiting times in years.

5 Analysis of periodicity

As in Ditlevsen et al. (2007) we use the Rayleigh measure to quantify the amount of periodicity for discrete events. This is a measure for the average phase coherence for N events. For period T it is given by $R(T) = \frac{1}{N} \left| \sum_{n=1}^N \left(\cos \frac{2\pi t_n}{T} + i \sin \frac{2\pi t_n}{T} \right) \right|$. For each period T we compute the Rayleigh measure and assess the significance of the result by comparing it to the distribution generated from a

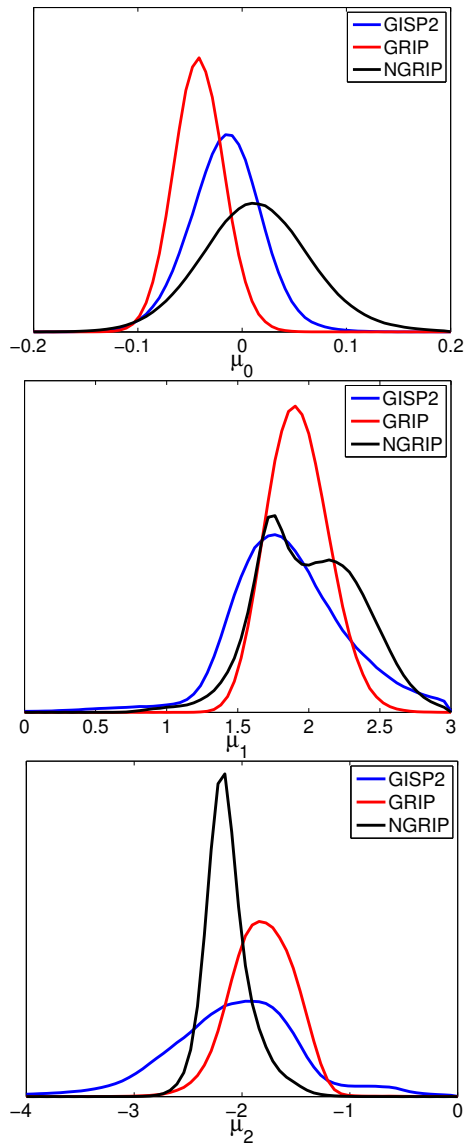


Fig. 6: Posterior distributions of forcing μ for three hidden states. There is significant agreement between NGRIP and GRIP for all three states.

null ensemble of the original data. The null ensemble is generated by sampling the data without replacement so that it is randomly ordered. Then hidden state sequences are generated by sampling from the posterior distribution conditioned upon the randomly ordered data and the other parameters (inferred for the original data). This means that the climate states in the ensemble will be equal in distribution to those in the actual data but their occurrence will be randomised.

Figure 7 shows the Rayleigh measure computed for the state corresponding to the DO events over a range of periods T with the mean and 2σ values for the null ensemble. Using this method no significant periodicities occur for any of the data sets agreeing with the result of Fig. 7 for randomly timed events.




Fig. 7: Rayleigh measures for ice core data (solid lines) shown together with the ensemble means (dashed lines) and upper 2σ values (dotted lines) in each case. Shown are NGRIP (black), GRIP (red) and GISP2 (blue).

6 Conclusions

We propose a statistical modelling strategy able to identify different climate regimes corresponding to differing rates of temperature change. This is a more objective method than setting threshold criteria based on the appearance of the data. The method allows for uncertainty in the detection of events, avoiding the sensitivity of results due to absolute inclusion or exclusion of a single event according to an authors preferred interpretation of the data.

To obtain a regular 50 year sampling, and to check the robustness of the procedure, we apply the method to an ensemble of time series derived from a Gaussian Process regression, separately for three Greenland $\delta^{18}\text{O}$ series: NGRIP, GRIP and GISP2 on different time scales. Using a Bayesian model comparison method we find that models with three states offer a good fit to all three data sets.

The inferred model states have distinct posterior distributions for their rates of temperature change with good agreement for all data sets. We found that one of the states closely corresponds to the onset of Greenland Interstadials and accounts for all of the canonical Dansgaard-Oeschger events of Dansgaard et al. (1993). This state had a mean rate of temperature change of 0.04% per year, which is greater than the 0.01% per year threshold criteria of Rahmstorf (2003). Contrary to Rahmstorf (2003), DO event 9 at 40 kyr BP is included in this state while the Allerød event, occurring at the start of the Younger Dryas, is not. Although the threshold criteria is larger the method allows for greater flexibility in the identification of states due to the incorporation of noise.

Having identified states we analyse their recurrence properties. From the waiting time distribution we find no evidence to reject the null hypothesis of randomly timed events. We also apply the Rayleigh periodicity measure used by Ditlevsen et al. (2007) to an ensemble derived from the

posterior distribution of shuffled data and find no evidence of significant periodicity within the $\delta^{18}\text{O}$ series. Thus, we cannot reject our null hypothesis that DO events are random events. This suggests that DO events are likely triggered by internal dynamics of the climate system. The periodicity seen in GISP2 by previous authors is apparently due to a small number of events that occur at the regular interval of 1450–1500 years. The method we present is robust to these events. We conclude that there is no lasting periodic signal for the rapid warming events of the last glacial period.

We applied the analysis to the GRIP ice core record on the GICC05 time scale giving parameter estimates similar to those in Fig. 6. Contrasting with the results for NGRIP using this time scale there is more support for a single state model, similar to the results for the other time scales shown in Fig. 3a. This implies that some of the differences in the results could be due to regional effects (different locations of drilling sites for GRIP and NGRIP) rather than the different time scale. Indeed the NGRIP data is less smooth, with larger jumps than the other data sets.

Here, we have been primarily concerned with statistical methods to detect jump events within noisy, irregularly sampled data. To make conclusions about variability of high latitude glacial climate one should relate the signal to temperature while controlling for regional affects. Past changes in atmospheric transport from the source of evaporation to the site of measurement can affect the stable isotope ratios. The seasonality of precipitation can also affect the temperature reconstruction. Simulation studies are one way to gain a better understanding of the influences on $\delta^{18}\text{O}$ signals (Krinner et al., 1997). Multi-proxy studies may also be useful to control for factors that could affect $\delta^{18}\text{O}$ besides the temperature of condensation. Masson-Delmotte et al. (2006) reconstruct temperatures allowing for changes in seasonality of precipitation and seawater isotope composition using a different slope between temperature and isotope ratio for cold and warm periods. The temperatures are shown to be consistent with the borehole and gas fractionation temperature reconstruction.

Future work could include applying our analysis directly to temperature reconstructions. Extensions could include an extra component to the likelihood function incorporating the dating uncertainty. This would be a more challenging statistical problem as these errors are not independent. This analysis would benefit from including the estimated maximum annual layer counting errors available for the NGRIP GICC05 time scale.

Appendix A

Conditional posterior distributions

For μ_j , we have

$$\mu_j \sim \mathcal{N} \left(\frac{(\sum_{i=0}^{N-1} I[S_i = j] X_{i+1} - \sum_{i=0}^{N-1} I[S_i = j] X_i)}{\sum_{i=0}^{N-1} I[S_i = j] + \sigma_j^2 / \sigma_\mu^2}, \frac{\sigma_j^2 \sigma_\mu^2}{\sigma_\mu^2 \sum_{i=0}^{N-1} I[S_i = j] + \sigma_j^2} \right), \quad (\text{A1})$$

where $I[S_i = j] = 1$ if $S_i = j$ or 0 otherwise. The posterior for τ_j is the Gamma distribution

$$\tau_j \sim \Gamma \left(a + \sum_{i=0}^{N-1} \frac{I[S_i = j]}{2}, b + \frac{1}{2} \sum_{i=0}^{N-1} I[S_i = j] (X_{i+1} - X_i - \mu_{S_i})^2 \right). \quad (\text{A2})$$

Each unobserved state is sampled from the updated multinomial distribution with probability

$$P(S_i = j) \propto \lambda_j \sqrt{\tau_j} \exp(-\frac{\tau_j}{2} (X_{i+1} - X_i - \mu_j)^2). \quad (\text{A3})$$

whereas λ is sampled as a vector, from the updated Dirichlet distribution

$$\lambda \sim \text{Dir}(\hat{\alpha}), \quad (\text{A4})$$

where $\hat{\alpha}_j = \alpha + \sum_{i=0}^{N-1} I[S_i = j]$.

A1 Laplace-Metropolis estimator

The approximation of the marginal likelihood by the Laplace-Metropolis method.

$$P(X) \approx (2\pi)^{d/2} |\Psi|^{1/2} P(X|\hat{\theta}) P(\hat{\theta}), \quad (\text{A5})$$

where X is the observed data, $\hat{\theta}$ is the posterior mode and Ψ is the empirical covariance matrix. In this equation the missing variables S have been integrated out. In this model this is computationally tractable because the S_i are independent for all i and also X_i is independent of X_j given S_i and S_j . The likelihood term is then

$$\log P(X|\hat{\theta}) = \sum_{i=1}^n \log P(X_i|\hat{\theta}) \quad (\text{A6})$$

It is easy to evaluate $P(X_i|\hat{\theta})$ as it is simply a sum over all possible J hidden states:

$$P(X_i|\hat{\theta}) = \sum_{j=1}^J P(X_i|S_i = j, \hat{\theta}) P(S_i = j|\hat{\theta}). \quad (\text{A7})$$

The full data likelihood is then $P(X_i|S_i = j, \hat{\theta}) = \mathcal{N}(X_{i+1} - X_i - \mu_j, 1/\sqrt{\tau_j})$, with the prior $P(S_i = j|\hat{\theta}) = \lambda_j$.

Acknowledgements. This study is part of the British Antarctic Survey Polar Science for Planet Earth Programme. It was partly funded by The Natural Environment Research Council. D. Peavoy is grateful for support from the Engineering and Physical Sciences Research Council. The authors would also like to thank the Centre for Ice and Climate at the Niels Bohr Institute for providing the data. We are also grateful for comments by E. Wolff, P. Ditlevsen and M. Crucifix and two anonymous referees which helped to improve this manuscript.

Edited by: V. Masson-Delmotte

References

- Alley, R. B., Anandakrishnan, S., and Jung, P.: Stochastic Resonance in the North Atlantic, *Paleoceanography*, 16, 190–198, 2001.
- Andersen, K. K., Svensson, A., Johnsen, S. J., Rasmussen, S. O., Bigler, M., Rthlisberger, R., Ruth, U., Siggaard-Andersen, M. L., Steffensen, J. P., Dahl-Jensen, D., Vinther, B. M., and Clausen, H. B.: The Greenland Ice Core Chronology 2005, 15–42 ka, Part 1: constructing the time scale, *Quaternary Science Reviews*, 25, 3246–3257, critical Quaternary Stratigraphy, 2006.
- Blunier, T., Chappelliaz, J., Schwander, J., Dallenbach, Stauffer, B., Stocker, T. F., Raynaud, D., Jouzel, J., Clausen, H. B., Hammer, C. U., and Johnsen, S.: Asynchrony of Antarctic and Greenland climate change during the last glacial period., *Nature*, 384, 739–743, 1998.
- Bond, G., Broecker, W., Johnsen, S., McManus, J., Labeyrie, L., Jouzel, J., and Bonani, G.: Correlations between climate records from North Atlantic sediments and Greenland ice, *Nature*, 365, 143–147, 1993.
- Braun, H., Christi, M., Rahmstorf, S., Ganopolski, A., Mangini, A., C., K., Roth, K., and Kromer, B.: Possible solar origin of the 1470-year glacial climate cycle demonstrated in a coupled model, *Nature*, 438, 208–211, 2005.
- Capron, E., Landais, A., Chappellaz, J., Schilt, A., Buiron, D., Dahl-Jensen, D., Johnsen, S. J., Jouzel, J., Lemieux-Dudon, B., Loulergue, L., Leuenberger, M., Masson-Delmotte, V., Meyer, H., Oerter, H., and Stenni, B.: Millennial and sub-millennial scale climatic variations recorded in polar ice cores over the last glacial period, *Clim. Past*, 6, 345–365, doi:10.5194/cp-6-345-2010, 2010.
- Dansgaard, W., Johnsen, S. J., Clausen, H. B., Dahl-Jensen, D., Gundestrup, N. S., Hammer, C. U., Hvidberg, C. S., Steffensen, J. P., Sveinbjornsdottir, A. E., Jouzel, J., and Bond, G.: Evidence for general instability of past climate from a 250-kyr ice-core record, *Nature*, 364, 218–220, 1993.
- Ditlevsen, P.: Observation of α -stable noise induced millennial climate changes from an ice-core record, *Geophys. Res. Lett.*, 26, 1441–1444, 1999.
- Ditlevsen, P. D., Kristensen, M. S., and Andersen, K. K.: The Recurrence Time of Dansgaard – Oeschger Events and Limits on the Possible Periodic Component, *J. Climate*, 18, 2594–2603, 2005.
- Ditlevsen, P. D., Andersen, K. K., and Svensson, A.: The DO-climate events are probably noise induced: statistical investigation of the claimed 1470 years cycle, *Clim. Past*, 3, 129–134, doi:10.5194/cp-3-129-2007, 2007.
- Gilks, W. R., Richardson, S., and Spiegelhalter, D., (Eds.): *Markov Chain Monte Carlo In Practice*, Chapman and Hall, 1995.
- Grootes, P. M. and Stuiver, M.: Oxygen 18/16 variability in Greenland snow and ice with 103 to 105-year time resolution, *J. Geophys. Res.*, 102(26), 455–470, 1997.
- Johnsen, S. J., Dahl-Jensen, D., Gundestrup, N., Steffensen, J. P., Clausen, H. B., Miller, H., Masson-Delmotte, V., Sveinbjornsdottir, A. E., and White, J.: Oxygen isotope and palaeotemperature records from six Greenland ice-core stations: Camp Century, Dye-3, GRIP, GISP2, Renland and NorthGRIP, *J. Quaternary Sci.*, 16, 299–307, 2001.
- Kass, R. E. and Raftery, A. E.: Bayes Factors, *Journal of the American Statistical Association*, 90, 773–795, 1995.
- Krinner, G., Genthon, C., and Jouzel, J.: GCM analysis of local influences on ice core δ signals, *Geophys. Res. Lett.*, 24, 2825–2828, 1997.
- Kwasniok, F. and Lohmann, G.: Deriving dynamical models from paleoclimatic records: Application to glacial millennial-scale climate variability, *Phys. Rev. E*, 80, 066104, doi:10.1103/PhysRevE.80.066104, 2009.
- Livina, V. N., Kwasniok, F., and Lenton, T. M.: Potential analysis reveals changing number of climate states during the last 60 kyr, *Clim. Past*, 6, 77–82, doi:10.5194/cp-6-77-2010, 2010.
- Massart, P.: The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality, *Ann. Probab.*, 18, 1269–1283, 1990.
- Masson-Delmotte, V., Dreyfus, G., Braconnot, P., Johnsen, S., Jouzel, J., Kageyama, M., Landais, A., Loutre, M.-F., Nouet, J., Parrenin, F., Raynaud, D., Stenni, B., and Tuenter, E.: Past temperature reconstructions from deep ice cores: relevance for future climate change, *Clim. Past*, 2, 145–165, doi:10.5194/cp-2-145-2006, 2006.
- Meese, D. A., Gow, A. J., Alley, R. B., Zielinski, G. A., Grootes, P. M. Ram, M., Taylor, K. C., Mayewski, P. A., and Bolzan, J. F.: The Greenland Ice Sheet Project 2 depth-age scale: Methods and results, *J. Geophys. Res.*, 102, 26411–26423, 1997.
- Raftery: *Markov Chain Monte Carlo In Practice*, chap. Hypothesis testing and model selection via posterior simulation, 163–188, Chapman and Hall, 1996.
- Rahmstorf, S.: Timing of abrupt climate change: A precise clock, *Geophys. Res. Lett.*, 30, 1510–1513, 2003.
- Rasmussen, C. E. and Williams, C. K. I.: *Gaussian Processes for Machine Learning*, MIT Press, 2006.
- Stocker, T. F. and Johnsen, S. J.: A minimum thermodynamic model for the bipolar seesaw, *Paleoceanography*, 18(11), 1–9, 2003.
- Wang, Y. J., Cheng, H., Edwards, R. L., An, Z. S., Wu, J. Y., Shen, C.-C., and Dorale, J. A.: A High-Resolution Absolute-Dated Late Pleistocene Monsoon Record from Hulu Cave, China, *Science*, 294, 2345–2348, doi:10.1126/science.1064618, 2001.