

Inference of past climate from borehole temperature data using Bayesian Reversible Jump Markov chain Monte Carlo

Peter O. Hopcroft,¹ Kerry Gallagher² and Chris C. Pain¹

¹Department Earth Science & Engineering, Imperial College, London SW7 2AZ, UK. E-mail: peter.hopcroft00@imperial.ac.uk

²Géosciences Rennes, Université de Rennes 1, Rennes, 35042, France

Accepted 2007 August 27. Received 2007 August 2; in original form 2007 May 21

SUMMARY

Estimates of past climate derived from borehole temperatures are assuming a greater importance in context of the millennial temperature variation debate. However, recovery of these signals is usually performed with regularization which can potentially lead to underestimation of past variation when noise is present. In this work Bayesian inference is applied to this problem with no explicit regularization. To achieve this Reversible Jump Markov chain Monte Carlo is employed, and this allows models of varying complexity (i.e. variable dimensions) to be sampled so that it is possible to infer the level of ground surface temperature (GST) history resolution appropriate to the data. Using synthetic examples, we show that the inference of the GST signal back to more than 500 yr is robust given boreholes of 500 m depth and moderate noise levels and discuss the associated uncertainties. We compare the prior information we have used with the inferred posterior distribution to show which parts of the GST reconstructions are independent of this prior information. We demonstrate the application of the method to real data using five boreholes from southern England. These are modelled both individually and jointly, and appear to indicate a spatial trend of warming over 500 yr across the south of the country.

Key words: ground surface temperature history, inversion of temperature logs, Markov chain Monte Carlo.

1 INTRODUCTION

Records of past surface temperatures are valuable data for placing recent warming in a long-term context and for constraining climate model sensitivities so that they can be used in prediction of future climate scenarios. However, historical instrumental records are limited over the majority of the globe to the past 150 yr (Jones & Mann 2004). For this reason, a variety of methods have been developed to indirectly measure the past variations using proxy indicators. One approach used to reconstruct pre-instrumental surface temperature variations which has begun to take a more significant role (e.g. Gonzalez-Rouco *et al.* 2003, 2006; Hegerl *et al.* 2007) involves examining temperature–depth profiles measured in terrestrial boreholes. These profiles exhibit perturbations at depth which, in the absence of other thermal influences, can be directly related to past variations in ground surface temperatures (GST) by the law of heat conduction (see Pollack & Huang 2000, and the references therein).

In order to infer the transient surface temperature record, given a temperature depth profile, we need to solve an ill-posed inverse problem. In this, the GST history and equilibrium thermal regime must be estimated from measured temperature profiles and thermophysical properties. Many different approaches to this problem have been used, (Shen & Beck 1991; Kakuta 1992; Mareschal &

Beltrami 1992; Shen *et al.* 1992; Huang *et al.* 1996; Dahl-Jensen *et al.* 1998; Woodbury & Ferguson 2006). Many of these methods involve a constraint on the smoothness or resolution of the reconstructed GST history. This is designed to deal with non-uniqueness of the inverse problem. Hartmann & Rath (2005) show that it is difficult to choose the magnitude of this model smoothing that should be applied and they demonstrate with synthetic examples using noisy data that very smooth reconstructions are obtained which have difficulty capturing temperature variability before 200 yr before present. In this work, we therefore, explore an alternative approach to inversion and uncertainty analysis using Bayesian inference.

Bayesian inference is cast in a probabilistic framework (e.g. Mosegaard & Tarantola 1995; Malinverno 2000; Tarantola 2005) which uses Bayes' theorem to map a prior probability density function (pdf), via the data likelihood, into a posterior pdf. Thus the posterior pdf describes the distributions of all model parameters conditioned on the data and prior information. The model uncertainty or non-uniqueness is then characterized by the width and shape of this posterior pdf. The prior distribution takes the place of the regularization by constraining model values to lie close to values considered realistic.

For many practical problems, sampling from the posterior pdf requires Markov chain Monte Carlo methods. In this work we use Reversible Jump Markov chain Monte Carlo (RJ-MCMC) Green

(1995) which additionally incorporates the number of model parameters as one of the model variables. Thus for this problem the GST model complexity is also inferred from the data. This leads to a reduced dependence on the user-defined parametrization. Recent applications of RJ-MCMC and Bayesian inference to other geophysical problems include Malinverno (2000, 2002), Andersen *et al.* (2003), Malinverno & Briggs (2004), Malinverno & Leaney (2005), Stephenson *et al.* (2006) and Sambridge *et al.* (2006).

In this paper, we first describe the forward problem, that is, solving for the temperature field in 1-D. We then give a brief introduction to Bayesian inference in general and the use of the RJ-MCMC sampling algorithm. The set up of this particular problem in a Bayesian framework is detailed followed by examples using synthetic and real data. We conclude with a discussion of the results obtained.

2 FORWARD MODEL AND BAYESIAN FORMULATION

2.1 Forward model

In order to find a solution to an inverse problem a forward model formulation is required. This provides the forward path from the system's control parameters to the output of that system. In this work, we solve the 1-D heat conduction equation, subject to two boundary conditions. The heat conduction equation is

$$\rho C \frac{\partial T(z, t)}{\partial t} = \frac{\partial}{\partial z} \left[\kappa(z) \frac{\partial T(z, t)}{\partial z} \right], \quad (1)$$

where z is depth which is positive into the ground, t is time, T is the temperature and κ , ρ and C are the thermal conductivity, density and specific heat capacity, respectively. The boundary conditions are an upper surface Robin boundary condition which is used to express the heat transfer from the air to the Earth through the surface heat transfer coefficient and a lower surface Neumann boundary condition for background geothermal heat flow:

$$q(0, t) = \beta [T_{\text{air}}(t) - T(0, t)] \quad (2)$$

$$q(z_{\text{max}}, t) = -\kappa(z_{\text{max}}) \frac{\partial T}{\partial z}, \quad (3)$$

where β is the heat transfer coefficient, T_{air} is the air temperature, T are temperatures in the ground and $q(z_{\text{max}}, t)$ is the geothermal heat flux, which is assumed to be constant with time.

In order to solve the forward problem efficiently we have chosen to use a finite element approximation to eq. (1) (e.g. Lewis *et al.* 1996). This allows us to take account of both radiogenic heat sources within the lithologies and also the depth dependence of the thermophysical properties. We use both a steady state and transient simulations: the former is used to calculate the equilibrium temperature depth profile and this is then perturbed with the latter. The finite element model has nodal spacing of 5 m at which the temperatures and all physical properties are specified. The model is run for 150–300 time steps of 2–5 yr. For 5 yr time steps and a typical GST history (2000 yr with a 1 K step change at 1000 yr), the maximum discretization error is 6×10^{-4} K, well below typical data noise values. This error is not very sensitive to the spatial discretization but scales with the length of the time steps used.

2.2 Bayesian inference

In a Bayesian formulation, all model parameters are described by probability distributions which effectively encompass the uncer-

tainty associated with a particular parameter. The aim of Bayesian inference is to quantify the posterior probability distribution which characterizes all of the model parameters given prior information and the data available (see, for example, Bernardo & Smith 1994; Jaynes 2003). Encompassing this relationship is Bayes' theorem:

$$p(\mathbf{m} | \mathbf{d}, \wp) = \frac{p(\mathbf{d} | \mathbf{m}, \wp) p(\mathbf{m} | \wp)}{p(\mathbf{d} | \wp)}, \quad (4)$$

where $p()$ stands for probability and 'a | b' implies conditional dependence of a on b, that is, a given b. \mathbf{m} is the model considered, \mathbf{d} is the data and \wp is the prior information. Thus $p(\mathbf{d} | \mathbf{m}, \wp)$ is the probability of observing the data given the model and the prior information \wp , and $p(\mathbf{m} | \wp)$ is the probability distribution of the model parameters given the prior information (which includes the model formulation).

The prior probability distribution used in inverse problems introduces information into the model so that the inverse solution is constrained in some sense. In many methods this is achieved by using a constraint on the model smoothness. However, because Bayesian inference is naturally parsimonious (Denison *et al.* 2002) there is no need to explicitly deal with the familiar trade-off between data fit and model resolution. Thus, when involved in choosing models of different complexity, the Bayesian approach will prefer the simpler model (in the Occam's razor sense: Jefferys & Berger 1992; Malinverno 2002; MacKay 2003; Sambridge *et al.* 2006). Assuming that we have a complex and simple model that both fit the observed data, then, for non-diffuse priors, the support from the data (or evidence, see below) for a simpler model will be higher in the region of the observed data as it is spread over a smaller region in data space (MacKay 2003).

The denominator of eq. (4), often referred to as the marginal likelihood of the data or evidence (Sambridge *et al.* 2006; Sivia & Skilling 2006) is generally an intractable integral and so the direct calculation of the posterior is not possible. However, this can be overcome by the use of Monte Carlo sampling using Markov chains (for a general introduction see Gilks *et al.* 1996) for which the posterior is only needed up to a constant of proportionality.

RJ-MCMC (Green 1995) is a generalization of the more well-known Metropolis–Hastings algorithm. It allows sampling of models over a range of dimensions in the same sampling scheme. Like the more well known fixed dimension MCMC, RJ-MCMC constitutes a two stage process of proposing a model, based on the current one, and then accepting or rejecting this proposal. The proposal is made by drawing from a probability distribution $q(\mathbf{m}' | \mathbf{m})$ such that a new proposed model \mathbf{m}' is conditional only on the current model \mathbf{m} . In the most general case of a proposal, in which we might allow a model to change dimensions, or simply changing the model parameters in fixed dimensions, we are always transforming one model to another. This transformation needs to be accounted for in the acceptance criterion, in accordance with the requirement of detailed balance (see Green 1995, 2003). This is taken into account by the exact form of the acceptance probability, α . When the number of dimensions does not change during a proposal, α can be written in the same form as the well-known Metropolis–Hastings sampler. However, the more general form of acceptance probability is:

$$\begin{aligned} \alpha &= \min \left[1, \frac{p(\mathbf{m}' | \wp)}{p(\mathbf{m} | \wp)} \cdot \frac{p(\mathbf{d} | \mathbf{m}', \wp)}{p(\mathbf{d} | \mathbf{m}, \wp)} \cdot \frac{q(\mathbf{m} | \mathbf{m}')}{q(\mathbf{m}' | \mathbf{m})} \cdot |\mathbf{J}| \right] \\ &= \min[1, (\text{prior ratio}) \cdot (\text{likelihood ratio}) \\ &\quad \times (\text{proposal ratio}) \cdot |\text{Jacobian}|], \end{aligned} \quad (5)$$

where \mathbf{J} , the Jacobian, accounts for the transformation from one model to another and is given by:

$$\mathbf{J} = \frac{\partial(\mathbf{m}', \mathbf{u}')}{\partial(\mathbf{m}, \mathbf{u})}, \quad (6)$$

where \mathbf{u} and \mathbf{u}' are used to transform the current model \mathbf{m} to the proposed model, \mathbf{m}' which may or may not be of a different dimension. In most practical fixed dimension cases we have encountered, the Jacobian is 1, and can be discarded. For variable dimension models the form of the Jacobian depends on the form of the transformation between dimensions. The total dimension of the variables and parameters before and after a dimension change must balance (Green 1995), such that

$$d(\mathbf{m}) + d(\mathbf{u}) = d(\mathbf{m}') + d(\mathbf{u}'). \quad (7)$$

For example, a model of dimensionality k when subject to a birth move, leads a model of dimensionality $(k + 1)$. The extra dimension coming from \mathbf{u} , which is size 1 and \mathbf{u}' is then size 0. We return to the details of this later for our specific problem.

If the new model is accepted, the current model \mathbf{m} is replaced by \mathbf{m}' which becomes the current model for the next iteration. If it is not accepted, the current model \mathbf{m} is retained for another iteration and a new modification is proposed. This process is then iterated many times so that, after a period of initial exploration of the model space (referred to as the burn-in), we collect a series of samples of the model parameters. It can be shown that for any proposal distribution, the stationary probability distribution sampled in this way will be a good estimate of the true probability distribution (Gilks *et al.* 1996).

2.3 Model setup

The GST history is set up as a series of linear segments with the nodes of these segments being the k GST model parameters, that is, $\mathbf{m} = (T_i, t_i; i = 1, k)$, where T is temperature and t is time. We allow between k_{\min} ($=2$) and k_{\max} ($=20$) of these nodes with a maximum separation of L ($=$ the duration of the reconstruction in years) and a minimum of the time step used in the finite element approximation ($=2-5$ yr). The RJ-MCMC algorithm samples these k model parameters as well as the equilibrium heat flow from below the borehole and the pre-reconstruction mean temperature (the steady state GST applied which is usually referred to as the pre-observational mean or POM). In Fig. 1 a schematic of the model set up is shown.

The parameter k represents, for this particular parametrization (linear segments), a measure of the complexity of the model. By allowing k to vary during the MCMC sampling we can address the question of model resolution more satisfactorily than if k was fixed. This is because the data explicitly influence the complexity of the model through the acceptance criterion (which includes both the data likelihood and the model prior). However, in the context of a continuous underlying process, k is not a physical quantity and clearly cannot be directly measured. Consequently, its relevance to model complexity only really applies to comparing the different models sampled by RJ-MCMC, which in our case can be regarded hierarchically (i.e. more parameters implies more complexity). Overall, this modelling scheme corresponds to what statisticians call an open modelling perspective (Bernardo & Smith 1994; Denison *et al.* 2002) whereby none of the models used is considered the ‘true’ model. This is true of any GST inversion model as past temperature changes of a variety of timescales are removed from the data by thermal diffusion. Thus real past temperature changes that occurred on scales from hourly to the decadal and even centennial are lost and cannot be represented in the model.

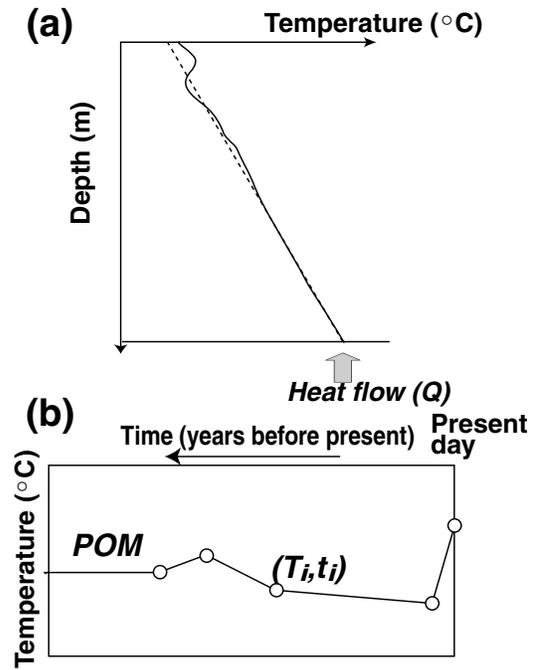


Figure 1. (a) Example data plot with heat flow Q . (b) Model setup: circles denote the interpolation points for the GST history. These can move in time and new points can be added or deleted to vary the complexity of the GST history. The additional model parameters are POM, the pre-observational mean GST and the heat flow, which is the lower boundary condition on the thermal model.

In simultaneous inversion (Beltrami & Mareschal 1992; Beltrami *et al.* 1992; Pollack *et al.* 1996; Beltrami *et al.* 1997) a number of boreholes are inverted together, that is we look for a common GST history but allow individual heat flow and POM values to be inferred. The GST history applied to each borehole is a deviation from the individual POM which allows for differences in long term means due to latitude but preserves the overall climatic trend (if similar across boreholes). All of the GST values are corrected for elevation using a lapse rate of 6.5 K km^{-1} .

2.4 Prior, proposals and likelihood

2.4.1 Prior

In this work the prior information takes three forms. For the GST temperature values we specify a most likely value and an associated probability distribution designed to take account of the full range of possible past variations. The second places a constraint on the time points of the GST history, such that they are probabilistically spaced out over the time domain. The equilibrium heat flow, POM values and the number of GST points are sampled from uniform priors, with specified upper and lower limits.

The prior on the GST values is set to a multivariate Gaussian distribution centred on the present mean GST value (i.e. the value at zero years):

$$p(\mathbf{T}_{\text{GST}} | \wp) = \frac{1}{[(2\pi)^k \det \mathbf{C}_{pr}]^{1/2}} \times \exp \left[-\frac{1}{2} (\mathbf{T}_{\text{GST}} - \mathbf{T}_{\text{prior}})^T \mathbf{C}_{pr}^{-1} (\mathbf{T}_{\text{GST}} - \mathbf{T}_{\text{prior}}) \right]. \quad (8)$$

We assumed no correlation between the GST parameters, so \mathbf{C}_{pr} is a diagonal matrix with entries equal to 1.0 K^{-2} (but see Serban &

Jacobsen 2001, for a discussion of correlated prior information in this context). \mathbf{T}_{GST} is the current GST model, $\mathbf{T}_{\text{prior}}$ is the prior mean value for the GST model and k is the number of GST interpolation nodes.

The time positions of the interpolation points are drawn using order statistics from a uniform distribution over the time interval examined (Green 1995, 2001). Thus the prior distribution on the time locations is given by:

$$p(\mathbf{t} | \wp) = \frac{k! I[0 < t_1 < t_2 < \dots < t_k < L]}{L^k}, \quad (9)$$

where \mathbf{t} are the time points of the GST history, L is the length of the time domain ($= t_{\text{max}} - t_{\text{min}}$) and I is the order statistic uniform distribution. This form of prior helps to ensure that very closely spaced interpolation points, which have negligible effect on the likelihood and are inconsistent with the data (because of the spreading with time of the thermal signal in the ground) do not emerge in the model (Green 1995).

Setting a prior distribution on a non-physical quantity such as k , we must consider questions such as what information do we have, and what does it mean, to differentiate between models with different numbers of model parameters. In this context we are looking at the expected resolution given that the original signal has been strongly filtered by thermal diffusion. Typical GST estimates derived from borehole temperature are of very low resolution, for example Huang *et al.* (1996) recommend century-long trends. Hence for reconstructions considered here, we might expect k to be around 5 if we had equally spaced nodes. By setting a uniform prior over the range [2, 20] we are including the resolution level recommended by Huang *et al.* (1996) in the range of models considered. However, as the Bayesian approach is naturally parsimonious, we avoid introducing unwarranted complexity (i.e. more model parameters) if the data do not require it. For example, if the true solution was a constant temperature over time, we would infer 1 GST parameter (assuming the present day value is fixed). Although we could use many more parameters and still produce a flat solution (i.e. constant temperature) over time with exactly the same data fit, the posterior probability is lower than for the simpler model. This is because we effectively penalize the complex model through the increased number of terms in the prior.

The priors on heat flows and POM values are uniform over the intervals [0, 100] m Wm⁻² and [0, 15] °C, respectively.

2.4.2 Proposal functions

In the generalized variable dimension case there are five types of changes to the model, each with an associated proposal function, to allow us to sample the model space. One of these change types is selected at random at each iteration of the RJ-MCMC algorithm:

- (i) Perturb one temperature value, T_i .
- (ii) Perturb one time value, t_i .
- (iii) Create a new GST point (birth).
- (iv) Delete one GST point (death).
- (v) Perturb the heat flow and POM values for one borehole.

The probability of selecting one of these change types is set to 1/5 for all iterations, except for when the number of GST model points reaches k_{min} or k_{max} . In the case that $k = k_{\text{min}}$, the death change and change v. (perturbing the heat flow and POM) have probability set to 0 and in the case that $k = k_{\text{max}}$, the birth step and change v. have probability set to zero. Option v. is not selected to improve efficiency

and the birth and death steps are limited in this way to keep k in the range [k_{min} , k_{max}].

Although the choice of proposal distribution used in the MCMC is essentially arbitrary, poor choices lead to very slow movement around the model space, such that convergence to the stationary distribution and exploration of the posterior distribution can take a very long time. It is therefore, desirable to choose proposal distributions carefully such that the model space search is as efficient as possible.

2.4.3 Likelihood

The likelihood term is based on a least-squares measure of the data misfit:

$$p(\mathbf{d}_{\text{obs}} | \mathbf{m}, \wp, k) = \frac{1}{[(2\pi)^n \det \mathbf{C}_d]^{1/2}} \times \exp \left[-\frac{1}{2} (\mathbf{d}_{\text{sim}} - \mathbf{d}_{\text{obs}})^T \mathbf{C}_d^{-1} (\mathbf{d}_{\text{sim}} - \mathbf{d}_{\text{obs}}) \right], \quad (10)$$

where there are n data points, \mathbf{d}_{obs} are observed subsurface temperatures and \mathbf{C}_d is the data covariance matrix (here assumed to be diagonal). The errors on the data values are, therefore, assumed to be Gaussian and uncorrelated. The values used for \mathbf{C}_d are discussed with the descriptions of synthetic and real data applications. \mathbf{d}_{sim} are the simulated underground temperatures which are calculated using the T_{GST} vector (size k) according to the finite element method of Section 2.1.

3 SYNTHETIC DATA EXAMPLES

In order to evaluate the effectiveness of Bayesian RJ-MCMC method, we first test it on synthetic data designed to be representative, in terms of quality, of real data. We can then assess what uncertainties may be associated with the various model parameters. For this purpose we used one multi-proxy surface air temperature (SAT) reconstruction (Moberg *et al.* 2005) and one instrumental SAT series (Jones & Moberg 2003) to produce a 2002-yr SAT history. This was then used as a GST forcing for calculating the subsurface temperature values in one synthetic profile of depth 500 m with values at 5 m intervals. The thermal diffusivity used was 4×10^{-7} m² s⁻¹ and the heat flow and POM were 60 m Wm⁻² and 9.6 °C, respectively. The temperature–depth values were then degraded with normally distributed 0.1 K noise. The algorithm was tested on both the noise free and 0.1 K noise cases. The algorithm was run for 500 000 iterations with the last 400 000 used to generate the posterior distributions of GST, heat flow and POM.

Figs 2 and 3 show the posterior distributions for the noise free and 0.1 K noise cases. The ‘P’ scale is the posterior probability for the GST value at a given time. The posterior distributions of heat flow and POM values can be reasonably well fitted with a Gaussian curve. For this synthetic case the posterior means and standard deviations for the heat flow and POM values were (60.8 ± 0.8) m Wm⁻² and (9.4 ± 0.3) °C, respectively. These values are slightly biased from the true values, possibly by the action of the prior information which does not correspond well with the first 1000 yr of the GST in the true model.

One way to quantify the posterior distribution is to evaluate posterior expectation values of the model (Gilks *et al.* 1996). The expected value is given by the equation:

$$E[f(m)] = \frac{\int f(m)\pi(m) dm}{\int \pi(m) dm}, \quad (11)$$

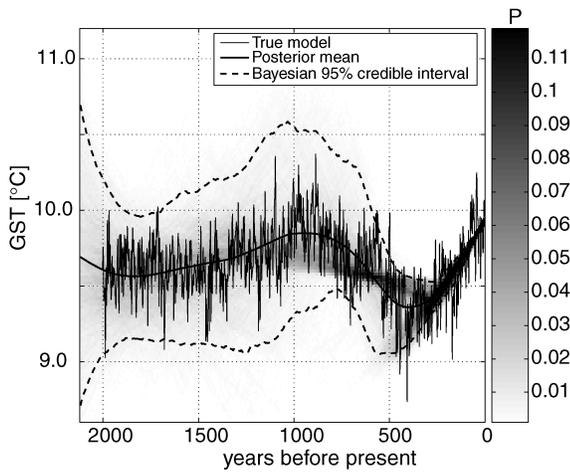


Figure 2. RJ-MCMC inversion: synthetic data with no noise added. The synthetic data are calculated using surface air temperature reconstruction (Moberg *et al.* 2005) and instrumental data (Jones & Moberg 2003) shown as the true model. The posterior mean is the expected value integrated across all models, the posterior distribution is shown in the grey-scale and the 95 per cent credible intervals are shown as the dashed lines.

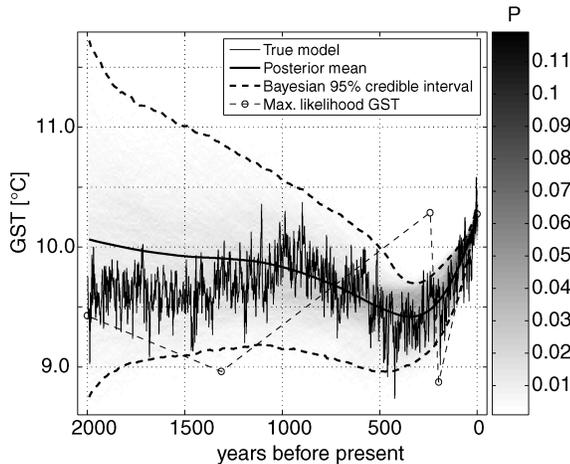


Figure 3. RJ-MCMC inversion: synthetic data with Gaussian random noise added (mean 0.0 K and standard deviation 0.1 K). The max likelihood GST history is shown for comparison with the posterior mean. The mean shows much better agreement with the true model as it combines information from all the models sampled whereas the maximum likelihood model is just the single model which provides a best fit to the data. The posterior distribution is shown in the grey-scale and the 95 per cent credible intervals are shown to illustrate the increasing uncertainty with time.

where $\pi(m)$ is the posterior distribution and $f(m)$ is some function of interest. It is then possible to use the posterior samples generated by a Markov chain Monte Carlo algorithm to calculate this expectation value (Gilks *et al.* 1996), using the following equation:

$$E[f(m)] = \frac{1}{n} \sum_1^n f(m). \quad (12)$$

In the case that $f(m)$ is the model value, the expectation value is the posterior mean which is calculated by integrating across all models weighted by their posterior probability values (as models are sampled with frequency proportional to their posterior probability). This has the advantage that it gives a smoother result than any single sampled model (all of which would consist of linear segments).

However, it combines information from all of the sampled models and, therefore, takes account of the uncertainties. Also, in the context of model complexity, this approach of model averaging allows us to assimilate the results from many models, conditioned either on a constant number of model parameters, or over all models allowing for the variable number of model parameters.

To illustrate this we show both the posterior mean model and the maximum likelihood model for case in Fig. 3. The maximum likelihood model does not fit within the 95 per cent credible intervals and is a good example of overfitting the data. The mean however, corresponds well with the true model and is much smoother than the maximum likelihood model.

The Bayesian 95 per cent credible interval is calculated by removing 2.5 per cent of the smallest and largest GST values at each time point to give a range of GST values which have a 95 per cent probability of enclosing the true model. This range can be affected by the choice of prior information but for sensible prior choices gives a good idea of the changing resolution.

In the noise free case the GST has been well recovered across most of the time domain. However, in the more realistic noisy example, this is not the case and the reconstruction has slightly underestimated the warmth at 1000 yr and the cool period around 400 yr. The 95 per cent credible interval becomes wider and more asymmetric than in the noise free example.

In Fig. 4, the posterior for the noisy data simulation is plotted at four times in the past and compared with the prior distribution and the true values (over a range of ± 5 yr). For the recent past the algorithm samples with high probability close to the true model values, as we expect in this form of problem in which the resolution is expected to deteriorate as we go back in time. We can see that further back in time the posterior distribution becomes progressively more influenced by the prior. For example, the posterior at 2000 yr it is effectively the same as the prior, indicating we have learnt nothing more about the GST at this time from the data. This type of observation can be used with real data examples to ascertain which parts of the GST reconstruction are actually supported by the data.

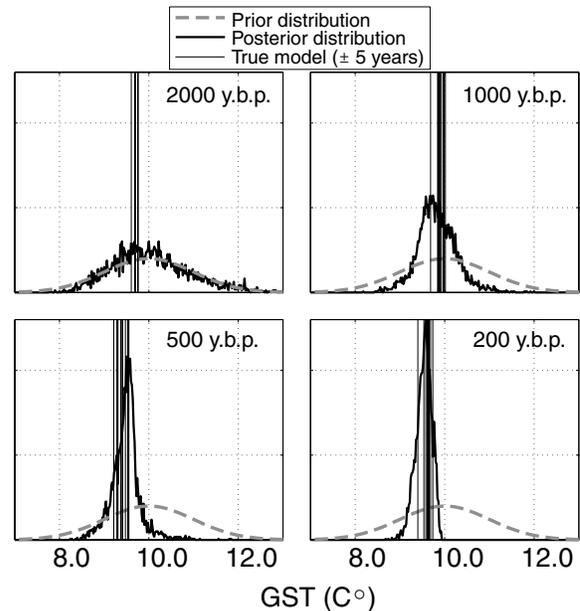


Figure 4. RJ-MCMC inversion: posterior probability distributions at 4 times before present (2000, 1000, 500 and 200 yr).

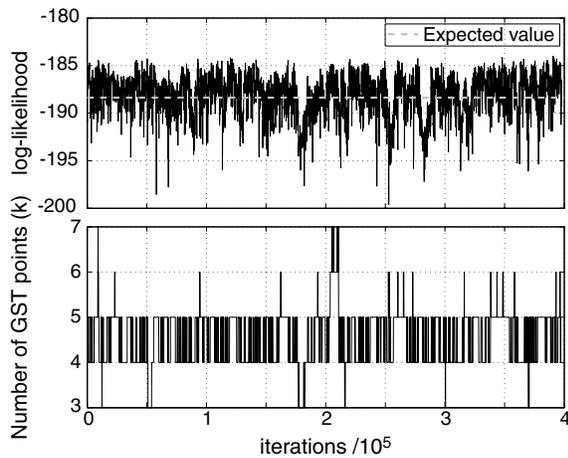


Figure 5. The log-likelihood and the number of time-temperature points in the GST history. The algorithm has explored models from size 3–7 with 4 showing most support from the data.

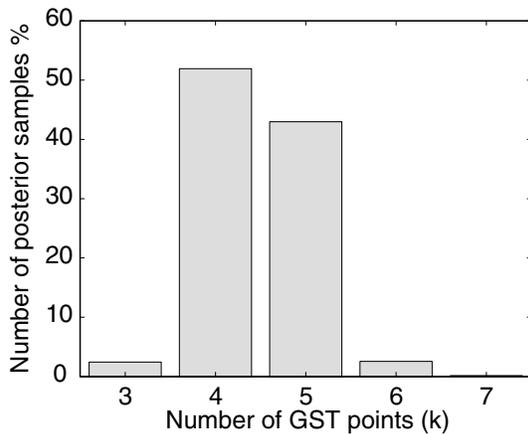


Figure 6. Posterior probability density function of the number of GST points for the noisy synthetic case as sampled by the Reversible Jump MCMC algorithm. The data clearly favour a GST history with only 4 or 5 points.

In Fig. 5, the log likelihood and number of GST points are shown for the noisy case. The likelihood is stable above the expectation value (-188.4) and shows that the algorithm is stable around this value. The data support between 3 and 7 points in the GST history, consistent with the anticipated resolution implied by the recommendation of Huang *et al.* (1996). 52 per cent of sampled models consisted of 4 GST points (see Fig. 6).

Our results show better resolution than the identical synthetic example (with 0.1K normally distributed noise) of Woodbury & Ferguson (2006) where no pre-20th century variation is captured. Another similar synthetic example with 0.1 K noise is described

by Hartmann & Rath (2005) where no variation prior to 1800 is resolved and the average pre-1800 GST signal is somewhat biased compared to the true model. Here we recover some of the variation of the past millennium for a 500 m synthetic profile. Additionally we are able to make inference about resolution changes over time and we can assess which parts of the reconstruction relate to the data or to inverse methodology (in our case prior information, in other methods model smoothing). We are, therefore, confident to use this method to analyse real data examples for signals of past climate over the past 500–1000 yr.

4 REAL DATA: EXAMPLES FROM THE UK

The real data used in this paper are selected from the International Heat Flow Commission borehole climatology database Huang & Pollack (1998) (all data we discuss below are available via <http://www.ncdc.noaa.gov/paleo/borehole/intro.html>). Summary information for the boreholes examined is given in Table 1. These data have been judged to be of good enough quality such that they can be used for palaeoclimate studies and specifically there is no evidence of vertical thermal convection. The five selected borehole data comprise temperature–depth profiles and thermal conductivity measurements with a minimum and maximum depths of 262 and 323 m, respectively. We have used a volumetric heat capacity of $2 \times 10^6 \text{ J m}^{-3} \text{ K}^{-1}$ for all the boreholes. The noise term (standard deviation) used in the likelihood term was found by examining how well the data could be fit using an initial exploratory RJ-MCMC run. This led to a value of 0.05 K chosen for all boreholes.

We inverted the profiles both individually and jointly using simultaneous inversion (Beltrami & Mareschal 1992; Beltrami *et al.* 1992; Pollack *et al.* 1996; Beltrami *et al.* 1997). In the latter case a joint likelihood is used in which we assume that the errors on the temperature data measured in each borehole are statistically independent from those in the other boreholes. This leads to the following likelihood formulation:

$$p(\mathbf{d}_{i=1..n_b} | \mathbf{m}, \wp, k) = p(\mathbf{d}_1 | \mathbf{m}, \wp, k) \times \dots \times p(\mathbf{d}_{n_b} | \mathbf{m}, \wp, k) \\ = \prod_{i=1}^{n_b} p(\mathbf{d}_i | \mathbf{m}, \wp, k), \quad (13)$$

where as before \mathbf{m} refers to the model (GST and heat flow parameters), \wp refers to prior information, k is the model dimensionality, \mathbf{d}_i is the data from borehole number i and there are n_b boreholes.

Because we are dealing with a total likelihood probability rather than functional misfits, this formulation means boreholes for which the climate signal is more uncertain will have a smaller influence on the overall result as their posterior probability values will always be lower than those with a more robust signal. The priors used in the real cases are the same as for the synthetic, with a mean set to the present day temperature and a standard deviation of 1.0 K.

Table 1. Borehole summary information for the five real data examples.

Borehole	Max. depth (m)	No. values	Lat.	Long.	Log date	Heat flow (m W m^{-2})	POM ($^{\circ}\text{C}$)
Wythlycombe (A)	262	80	50.65	-3.37	1983	39.0 ± 2.0	12.5 ± 0.34
Venn (B)	299	92	50.71	-3.32	1983	47.6 ± 1.1	11.4 ± 0.19
Seabarn (C)	320	16	50.62	-2.53	1978	60.0 ± 0.9	9.2 ± 0.19
Chalgrove (D)	323	100	51.66	-1.05	1983	44.8 ± 0.6	11.6 ± 0.15
Worcester (E)	298	92	52.22	-2.2	1983	44.6 ± 1.3	9.5 ± 0.23

The values quoted for the heat flow and POM are the posterior pdf means and standard deviations as sampled by the RJ-MCMC algorithm. These posterior distributions are well approximated by a Gaussian distribution.

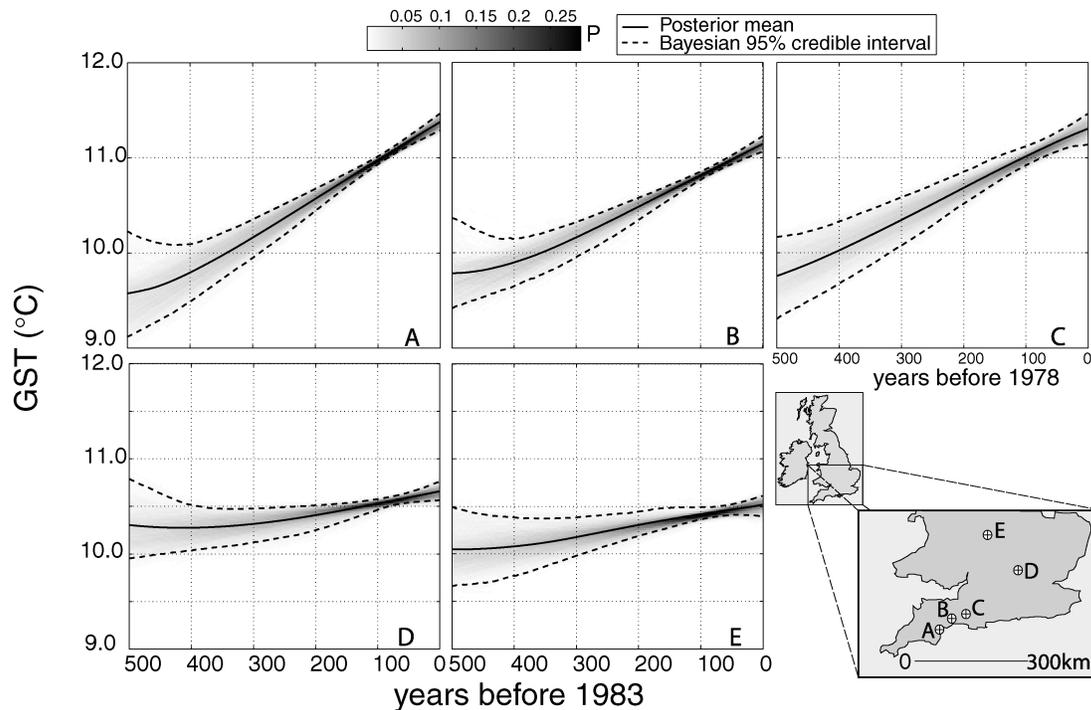


Figure 7. Posterior probability distributions for the individually inverted boreholes. The locations of these are shown in the inset map, and the details are given in Table 1.

The heat flow and POM values for the five boreholes are given in Table 1. Since the posterior distributions for heat flow and POM are close to Gaussian, the mean and standard deviations as sampled by the RJ-MCMC algorithm are quoted. In Fig. 7, the individual posterior distributions of GST for the five boreholes are shown. Again, the ‘P’ scale indicates posterior probability value showing the resolution of the GST signal with time. It is possible to rank the boreholes by comparing the posterior width of the GST signal, for example, by looking at similar plots as are shown in Fig. 4. In this case borehole D has the smallest uncertainties and C has the largest. Borehole D also shows the best agreement with the Central England temperature record (Manley 1974; Parker *et al.* 1992) which dates back to 1659 (not shown).

Intriguingly, a possible spatial trend emerges for these five boreholes. The southern boreholes (A, B and C) show 1.5 K warming over 500 yr and the central set (D and E) showing an average of 0.5 K. This apparent spatial trend clearly needs confirmation using more data, but a similar signal appears in Beltrami & Bourlon (2004) for which more boreholes were analysed.

For the joint inversion, the GST applied to each borehole is defined as deviations from the POM (which is sampled by the RJ-MCMC algorithm also). This allows for differences in the mean conditions due to latitude whilst allowing a common long-term climatic signal to be inferred. The joint posterior probability distribution for the five boreholes is shown in Fig. 8. We see that the signal leads to a warming of approximately 1.0 K over the 350 yr to 1983. It appears that some uncertainty has been introduced by combining the data in this way.

5 DISCUSSION AND CONCLUSIONS

In this work, we have demonstrated the use of Bayesian inference for inferring past climate from borehole temperatures. In order to

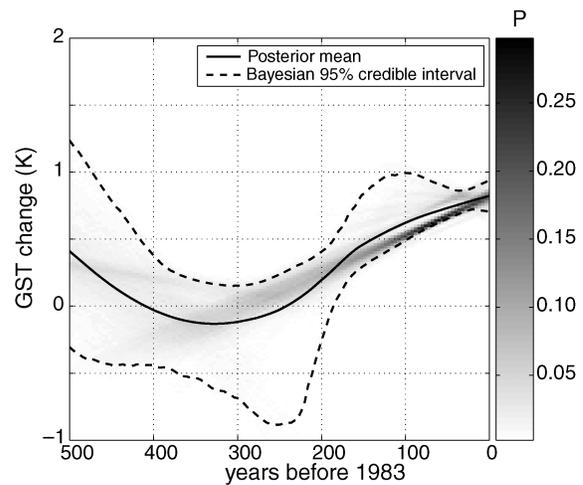


Figure 8. Posterior probability distribution for the simultaneous inversion of the five boreholes shown as deviations from the steady state.

achieve this we have utilized Reversible Jump Markov chain Monte Carlo sampling which allows efficient exploration of model space by sampling models of varying dimension. The advantages of this approach mean that the inverse problem can be treated in a fully non-linear manner and that explicit regularization of the model parameters is not required, thus avoiding the need to find an optimal regularization value (Hartmann & Rath 2005). Instead we come to define a posterior probability distribution of the GST conditioned on the measured data and the prior information. The expected value, determined from the posterior, provides an average over all models and, therefore, better captures the variability in the range of possible solutions than a single (e.g. best) model. This posterior distribution also incorporates the uncertainty due to the information loss from

diffusion of the thermal signal and noise on the data and allows identification of any non-uniqueness or poor resolution of GST or thermal equilibrium parameters. By comparing the posterior distribution with the prior it is possible to infer which parts of the distribution are supported by the data. Using the transdimensional form of MCMC, we can also address the complexity of the GST reconstruction in terms of the support from the data for different models.

In any Bayesian formulation, the outcome is to some extent influenced by the choice of prior information. However, it is impossible to define totally non-informative priors and so sensible choices relating real uncertainties must be made (e.g. Scales & Tenorio 2001). The priors used in this work are constant with time so that *a priori* no GST change is assumed. However, the prior distribution (i.e. its width) is designed to account for likely past variations. This is a good test of the methodology in the first case. However, the prior distribution is a flexible tool and could be used to incorporate information from instrumental data or other palaeoclimate reconstructions as well as the associated uncertainties.

Examples from real data provisionally indicate a spatial trend of GST history over the UK such that boreholes in the south display more warming over the last half millennium than those in central England. This trend would need clarification using more data from the region or explicit spatial modelling.

ACKNOWLEDGMENTS

We would like to thank the two reviewers, one of them Klaus Mosegaard, whose comments helped to greatly improve the content of this paper. Also, POH was supported through a studentship grant from the Environmental Mathematics and Statistics programme of NERC and EPSRC to KG.

REFERENCES

- Andersen, K.E., Brooks, S.P. & Hansen, M.B., 2003. Bayesian inversion of geoelectrical resistivity data, *J. Roy. Stat. Soc., Ser. B, Stat. Methodol.*, **65**, 619–642.
- Beltrami, H. & Bournon, E., 2004. Ground warming patterns in the Northern Hemisphere during the last five centuries, *Earth planet. Sci. Lett.*, **227**(3–4), 169–177.
- Beltrami, H. & Mareschal, J.C., 1992. Ground temperature histories for central and eastern Canada from geothermal measurements—Little Ice-Age signature, *Geophys. Res. Lett.*, **19**(7), 689–692.
- Beltrami, H., Jessop, A.M. & Mareschal, J.C., 1992. Ground temperature histories in eastern and central Canada from geothermal measurements—evidence of climatic-change, *Global Planet. Change*, **98**(2–4), 167–184.
- Beltrami, H., Cheng, L. & Mareschal, J.C., 1997. Simultaneous inversion of borehole temperature data for determination of ground surface temperature history, *Geophys. J. Int.*, **129**(2), 311–318.
- Bernardo, A. & Smith, J., 1994. *Bayesian Theory*, John Wiley, Chichester.
- Dahl-Jensen, D., Mosegaard, K., Gundestrup, N., Clow, G.D., Johnsen, S.J., Hansen, A.W. & Balling, N., 1998. Past temperatures directly from the Greenland Ice Sheet, *Science*, **282**(5387), 268–271.
- Denison, D.G.T., Holmes, C.C., Mallick, B.K. & Smith, A.F.M., 2002. *Bayesian Methods for Nonlinear Classification and Regression*, John Wiley and Sons, Chichester.
- Gilks, W.R., Richardson, S. & Spiegelhalter, D. (eds), 1996. *Markov Chain Monte Carlo in Practice*, Chapman and Hall, London.
- Gonzalez-Rouco, F., von Storch, H. & Zorita, E., 2003. Deep soil temperature as proxy for surface air-temperature in a coupled model simulation of the last thousand years, *Geophys. Res. Lett.*, **30**(21), 2116, doi:10.1029/2003GL018264.
- Gonzalez-Rouco, J.F., Beltrami, H., Zorita, E. & von Storch, H., 2006. Simulation and inversion of borehole temperature profiles in surrogate climates: spatial distribution and surface coupling, *Geophys. Res. Lett.*, **33**(1), L01703, doi:10.1029/2005GL024693.
- Green, P.J., 1995. Reversible jump Markov chain Monte Carlo computation and Bayesian model determination, *Biometrika*, **82**(4), 711–732.
- Green, P.J., 2001. A Primer on Markov chain Monte Carlo, in *Complex Stochastic Systems*, pp. 1–62, eds Barndorff-Nielsen, O., Cox, D.R. & Kluppelberg, C., Chapman and Hall, London.
- Green, P.J., 2003. Trans-dimensional Markov chain Monte Carlo, in Green, P., Hjort, N. & Richardson, S., *Highly Structured Stochastic Systems*, pp. 179–198, Oxford Statistical Science Series, Oxford University Press, Oxford.
- Hartmann, A. & Rath, V., 2005. Uncertainties and shortcomings of ground surface temperature histories derived from inversion of temperature logs, *J. Geophys. Eng.*, **2**(4), 299–311.
- Hegerl, G.C., Crowley, T.J., Allen, M., Hyde, W.T., Pollack, H.N., Smerdon, J. & Zorita, E., 2007. Detection of human influence on a new, validated 1500-year temperature reconstruction, *J. Clim.*, **20**(4), 650–666.
- Huang, S. & Pollack, H., 1998. Global borehole temperature database for climate reconstruction. IGBP PAGES/World Data Center-A for Paleoclimatology data contribution series 1998–044. NOAA/NGDC Paleoclimatology Program, Boulder CO, USA.
- Huang, S.P., Shen, P.Y. & Pollack, H.N., 1996. Deriving century-long trends of surface temperature change from borehole temperatures, *Geophys. Res. Lett.*, **23**(3), 257–260.
- Jaynes, E., 2003. *Probability Theory: The Logic of Science*, Cambridge University Press, Cambridge, UK.
- Jefferys, W.H. & Berger, J.O., 1992. Ockham's razor and Bayesian-analysis, *Am. Scientist*, **80**, 64–72.
- Jones, P.D. & Mann, M.E., 2004. Climate over past millennia, *Rev. Geophys.*, **42**(2), RG2002, doi:10.1029/2003RG000243.
- Jones, P.D. & Moberg, A., 2003. Hemispheric and large-scale surface air temperature variations: an extensive revision and an update to 2001, *J. Clim.*, **16**(2), 206–223.
- Kakuta, S., 1992. Surface-temperature history during the last 1000 years near Prudhoe Bay, Alaska: applying control theory to the inversion of borehole temperature profiles, *Global Planet. Change*, **6**(2–4), 225–244.
- Lewis, R.W., Morgan, K., Thomas, H.R. & Seetharamu, K., 1996. *The Finite Element Method in Heat Transfer Analysis*, Wiley, Chichester.
- MacKay, D., 2003. *Information Theory, Inference, and Learning Algorithms*, Cambridge University Press, Cambridge.
- Malinverno, A., 2000. A Bayesian criterion for simplicity in inverse problem parametrization, *Geophys. J. Int.*, **140**(2), 267–285.
- Malinverno, A., 2002. Parsimonious Bayesian Markov chain Monte Carlo inversion in a nonlinear geophysical problem, *Geophys. J. Int.*, **151**(3), 675–688.
- Malinverno, A. & Briggs, V.A., 2004. Expanded uncertainty quantification in inverse problems: hierarchical Bayes and empirical Bayes, *Geophysics*, **69**(4), 1005–1016.
- Malinverno, A. & Leaney, W.S., 2005. Monte-Carlo Bayesian look-ahead inversion of walkaway vertical seismic profiles, *Geophys. Prospect.*, **53**(5), 689–703.
- Manley, G., 1974. Central England temperatures—monthly means 1659 to 1973, *Q. J. R. Meteorol. Soc.*, **100**(425), 389–405.
- Mareschal, J. & Beltrami, H., 1992. Evidence for recent warming from perturbed geothermal gradients: examples from eastern Canada, *Clim. Dyn.*, **6**, 135–143.
- Moberg, A., Sonechkin, D.M., Holmgren, K., Datsenko, N.M. & Karlen, W., 2005. Highly variable Northern Hemisphere temperatures reconstructed from low- and high-resolution proxy data, *Nature*, **433**(7026), 613–617.
- Mosegaard, K. & Tarantola, A., 1995. Monte-carlo sampling of solutions to inverse problems, *J. Geophys. Res.—Solid Earth*, **100**(B7), 12 431–12 447.
- Parker, D.E., Legg, T.P. & Folland, C.K., 1992. A new daily central England temperature series, 1772–1991, *Int. J. Climatol.*, **12**(4), 317–342.
- Pollack, H.N. & Huang, S.P., 2000. Climate reconstruction from subsurface temperatures, *Ann. Rev. Earth Planet. Sci.*, **28**, 339–365.

- Pollack, H.N., Shen, P.Y. & Huang, S.P., 1996. Inference of ground surface temperature history from subsurface temperature data: interpreting ensembles of borehole logs, *Pure appl. Geophys.*, **147**(3), 537–550.
- Sambridge, M., Gallagher, K., Jackson, A. & Rickwood, P., 2006. Trans-dimensional inverse problems, model comparison and the evidence, *Geophys. J. Int.*, **167**(2), 528–542.
- Scales, J.A. & Tenorio, L., 2001. Prior information and uncertainty in inverse problems, *Geophysics*, **66**(2), 389–397.
- Serban, D.Z. & Jacobsen, B.H., 2001. The use of broad-band prior covariance for inverse palaeoclimate estimation, *Geophys. J. Int.*, **147**(1), 29–40.
- Shen, P.Y. & Beck, A.E., 1991. Least-squares inversion of borehole temperature-measurements in functional space, *J. Geophys. Res.-Solid Earth*, **96**(B12), 19 965–19 979.
- Shen, P.Y., Wang, K., Beltrami, H. & Mareschal, J.C., 1992. A comparative study of inverse methods for estimating climatic history from borehole temperature data, *Global Planet. Change*, **6**(2–4), 113–127.
- Sivia, J. & Skilling, D., 2006. *Data Analysis. A Bayesian Tutorial*, 2nd edn, Oxford University Press, Oxford, UK.
- Stephenson, J., Gallagher, K. & Holmes, C.C., 2006. Low temperature thermochronology and strategies for multiple samples 2: partition modelling for 2D/3D distributions with discontinuities, *Earth planet. Sci. Lett.*, **241**(3–4), 557–570.
- Tarantola, A., 2005. *Inverse Problem Theory and Methods for Model Parameter Estimation*, SIAM, Philadelphia.
- Woodbury, A.D. & Ferguson, G., 2006. Ground surface paleotemperature reconstruction using information measures and empirical Bayes, *Geophys. Res. Lett.*, **33**(6), L06702, doi:10.1029/2005GL025243.

APPENDIX A: PROPOSALS

The GST time and GST proposals are updated individually, and are scaled so that proposals further back in time are larger. This is in order to improve the overall efficiency of the algorithm as the uncertainties are expected to expand with time before the present. If temperature value labelled i is chosen for update and k is the current number of GST model points, the new value T'_i is given by:

$$T'_i = T_i + u_2 \times \sigma_T^m \times \exp\left(\frac{i}{k}\right), \quad (\text{A1})$$

where T and T' are the current and proposed GST values, respectively and u_2 is a random Gaussian draw of mean 0 and standard deviation 1.0, and σ_T^m is a scaling factor for moving the temperature point (set to 0.1 for all cases here). Similarly the time values are updated using:

$$t'_i = t_i + [t_- - t_+] \times u_1 \times \sigma_t^m \times \exp\left(\frac{i}{k}\right), \quad (\text{A2})$$

where t and t' are the current and proposed GST time points, respectively, u_1 is a random Gaussian draw of mean 0 and standard deviation 1.0, and σ_t^m is a scaling factor for moving the time point (set to 0.05 for all cases here). t_- and t_+ are the time values for the two adjacent GST model points. This proposal scales the time proposal to the width of the interval in which it is located and also exponentially with time in the past. These updates were found to give better mixing of the Markov chains, such that less iterations were required to quantify the posterior probability distributions. This move type is fully reversible as time points are not allowed to cross each other.

The birth and death steps need to be designed so that one can exactly reverse the action of the other. See Fig. A1 for an illustration of this birth/death procedure. For a birth move, a time in the history is chosen at random and a GST interpolation point is added there. The current GST temperature at this time is then perturbed using a

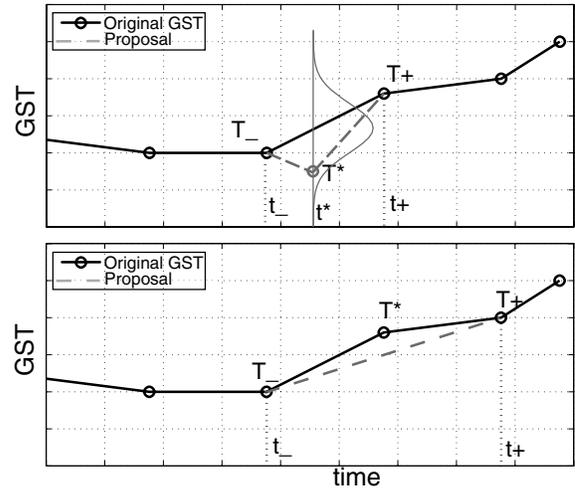


Figure A1. Upper panel: an example birth move, a random draw across the whole time domain determines the time value for the new point and a random draw from a Gaussian distribution determines the GST value at the new point. The distribution width has been exaggerated in order to improve clarity. Lower panel: an example death move.

Gaussian distribution. The equations relating new and old points:

$$t_* = t_- + \sigma_t^b u_1 (t_+ - t_-) \quad (\text{A3})$$

$$T_* = \frac{(T_+ - T_-)}{(t_+ - t_-)} (t_* - t_-) + u_2 \sigma_T^b + T_-, \quad (\text{A4})$$

where t_* and T_* are the new time and temperature values, (t_-, T_-) and (t_+, T_+) are the coordinates of the two points either side, u_1 is a random draw from a uniform distribution over the interval $[0, 1]$ and u_2 is random draw from a Gaussian distribution of mean 0 and standard deviation 1.0. σ_t^b is set to 1.0 and σ_T^b is the standard deviation of the proposal for the new temperature value. In this work $\sigma_T = 10^{-6}$ has been used as it was found that this value led to the algorithm correctly identifying the dimensionality of the GST curves for synthetic cases. When σ_T was set much larger or smaller the number of dimensions tended towards k_{\min} or k_{\max} , respectively.

In the death step, an interpolation point is chosen at random and deleted, such that the GST curve is linearly interpolated from the two adjacent interpolation point temperatures. In both cases the GST history is returned to its current state if the proposal is not accepted (as defined by the acceptance criterion, eq. 5).

The heat flow and POM temperatures are updated simultaneously using a bivariate Gaussian proposal distribution. This is used because the heat flow and equilibrium temperature values for a particular borehole are strongly negatively correlated. The required 2×2 correlation matrix for each borehole was found by running an exploratory MCMC simulation before running the GST sampler.

APPENDIX B: CALCULATING THE ACCEPTANCE PROBABILITY

B1 Fixed dimension moves of time or temperature values

In the case that one T or t value is changed and we have a symmetrical proposal function, the proposal ratio is equal to 1. The acceptance probability is then only dependent on the likelihood ratio and the

prior ratio on the temperature values:

$$\alpha = \min \left[1, \frac{p(\mathbf{T}' | \wp)}{p(\mathbf{T} | \wp)} \cdot \frac{p(\mathbf{d} | \mathbf{m}', \wp)}{p(\mathbf{d} | \mathbf{m}, \wp)} \right] \\ = \min[1, (\text{prior ratio}) \cdot (\text{likelihood ratio})]. \quad (\text{B1})$$

In the case that the time step positions are perturbed the prior ratio from eq. (9) is introduced and so the acceptance term is given by:

$$\alpha = \min \left[1, \frac{(t_+ - t'_*)(t'_* - t_-)}{(t_+ - t_*)(t_* - t_-)} \cdot \frac{p(\mathbf{d} | \mathbf{m}', \wp)}{p(\mathbf{d} | \mathbf{m}, \wp)} \right] \\ = \min[1, (\text{prior ratio}) \cdot (\text{likelihood ratio})], \quad (\text{B2})$$

where t_- and t_+ are the GST model point times which are adjacent to the model point being perturbed (t_*).

B2 Birth/death acceptance term

For the birth (and death) acceptance term, the prior and proposal ratio and also the Jacobian are required. The prior ratio takes account of the number of time steps and the spacing using the order statistics drawn uniformly over the time interval in question, L . The ratio is again calculated from eq. (9):

$$\frac{p(\mathbf{m}' | \wp)}{p(\mathbf{m} | \wp)} = \frac{p(\mathbf{T}' | \wp)}{p(\mathbf{T} | \wp)} \frac{p(k+1)}{p(k)} \frac{k+1}{L} \frac{(t_+ - t_*)(t_* - t_-)}{(t_+ - t_-)}. \quad (\text{B3})$$

The proposal ratio reflects the probability of choosing a point for birth or death move, and is given by

$$\frac{q(\mathbf{m} | \mathbf{m}')}{q(\mathbf{m}' | \mathbf{m})} = \frac{L d_{k+1}}{(k+1) b_k}, \quad (\text{B4})$$

where k is the current model dimension, d_{k+1} is the death move probability and b_k is the birth probability. The ratio $L/(k+1)$ accounts for the probability of choosing a point for a birth or death step. For a birth, the location of the time point is chosen with probability $=1/L$, that is, uniform over the duration of the GST history. For a death an existing point is selected with a probability $=1/(k+1)$ (since we are reversing a birth step from $(k+1)$ to k). The ratio of the two values, b_k and d_{k+1} is $\frac{1}{5}$ as there are 5 proposal types to choose from. When k has reached the value k_{\min} or k_{\max} the proposal ratio for a birth or death change is different because there are less proposals types to choose from. We do not allow selection of option v. (heat flow and POM) when the minimum or maximum dimension limit has been reached. For $k = k_{\min}$ only two options remain: birth or perturb a temperature, T_i . For $k = k_{\max}$ only three options remain: death, perturb a temperature, T_i or perturb a time point t_i . This leads to

the proposal ratios for birth (at $k = k_{\min}$) and death (at $k = k_{\max}$) as follows:

$$\frac{q(\mathbf{m} | \mathbf{m}')}{q(\mathbf{m}' | \mathbf{m})_b} = \frac{L}{(k_{\min} + 1)} \frac{d_{(k_{\min}+1)}}{b_{k_{\min}}} = \frac{2L}{5(k_{\min} + 1)} \\ \frac{q(\mathbf{m} | \mathbf{m}')}{q(\mathbf{m}' | \mathbf{m})_d} = \frac{(k_{\max} - 1)}{L} \frac{b_{k_{\max}-1}}{d_{k_{\max}}} = \frac{3(k_{\max} - 1)}{5L}. \quad (\text{B5})$$

B3 Derivation of birth/death Jacobian term

The equations relating the new temperature and time points to the current points are given again for clarity (eqs A3 and A4):

$$t_* = t_- + \sigma_t u_1 (t_+ - t_-) \quad (\text{B6})$$

$$T_* = \frac{(T_+ - T_-)}{(t_+ - t_-)} (t_* - t_-) + u_2 \sigma_T + T_-, \quad (\text{B7})$$

where t_* and T_* are the new time and temperature values, (t_-, T_-) and (t_+, T_+) are the coordinates of the two points either side. Now substituting t_* into eq. (B7) gives:

$$T_* = (T_+ - T_-) \sigma_t u_1 + u_2 \sigma_T + T_-. \quad (\text{B8})$$

Eqs (B6) and (B8) are required for calculating the Jacobian term. This is achieved by considering those parameters which are unchanged by the birth or death move. This leads to a 2×2 determinant:

$$|\mathbf{J}| = \left| \frac{\partial(T_*, t_*)}{\partial(u_1, u_2)} \right| \\ = \begin{vmatrix} (T_+ - T_-) \sigma_t & \sigma_T \\ \sigma_t (t_+ - t_-) & 0 \end{vmatrix} \\ = -\sigma_T \sigma_t (t_+ - t_-), \quad (\text{B9})$$

where the four terms have been derived by differentiating eqs (B6) and (B8). The acceptance term for a birth move becomes:

$$\alpha = \min \left\{ 1, \frac{p(k+1)}{p(k)} \frac{(k+1)}{L} \frac{(t_+ - t_*)(t_* - t_-)}{(t_+ - t_-)} \right. \\ \times \frac{p(\mathbf{T}' | \wp, k)}{p(\mathbf{T} | \wp, k)} \cdot \frac{p[\mathbf{d} | \mathbf{m}', \wp, (k+1)]}{p[\mathbf{d} | \mathbf{m}, \wp, k]} \\ \left. \times \frac{L d_{k+1}}{b_k (k+1)} \cdot (-\sigma_T \sigma_t) (t_+ - t_-) \right\}, \quad (\text{B10})$$

where the the first ratio term always takes the value 1. The corresponding acceptance term for a death step, has the same form but with appropriate relabelling such that the current model is dimension k , moving via death to a model of dimension $(k-1)$. The ratio terms also need to be inverted (Green 1995).