Bayesian Layer-Counting in Ice-Cores - Reconstructing the Time Scale

J. J. Wheatley[1], P. G. Blackwell[1], N. J. Abram[2], E. W. Wolff[3]

Bayesian Young Statisticians Meeting (BAYSM), Milan June, 5-6, 2013
Paper no. 43

[1] School of Mathematics and Statistics, The University of Sheffield, UK
j.wheatley@shef.ac.uk
[2] Research School of Earth Sciences, Australian National University, Australia
[3] British Antarctic Survey, Cambridge, UK

Abstract

The concentrations of various chemicals, particles and gasses in ice-cores hold a continuous record of climatic and environmental information dating back hundreds of thousands of years. In order to interpret these data we must first learn about their underlying, unobserved time scale. We present a fully Bayesian univariate approach to obtain a marginal posterior distribution for the time of year, as well as the date, at each depth.

Markov chain Monte Carlo; time-axis uncertainty

1 Introduction

Environmental signals are measured from ice-cores either by cutting them into sections – indexed by depth – and analysing the melt-water to provide a piecewise average or via Continuous Flow Analysis (CFA), see [1]. Some high resolution signals, those with multiple measurements per year, have annual cycles which show as quasi-periodic seasonality in the depth series. Layer-counting uses this periodicity to count back in time, year by year, and is currently achieved by eye, at considerable effort, see [1]. Using a simple, flexible model for one such signal we use a Markov chain Monte Carlo approach to reconstruct the underlying periodic process. The latent chronology is sampled directly in a way that allows the number of cycles in the reconstruction to be changed without the need for dimension-changing algorithms such as Reversible Jump. We allow for the dependence in observation error and the lack of stationarity by modelling means, amplitudes and errors as continuous functions of depth.

2 The Model

The signal is modelled as:

$$x_i = \alpha_i \sin(2\pi\tau_i) + \beta_i$$

where τ_i is the latent time-scale of interest at depth $i \in (1, 2, \ldots, n)$. The reconstruction of the signal is described by the parameters: $\boldsymbol{\theta} = \{\boldsymbol{\tau}, \boldsymbol{\alpha}, \boldsymbol{\beta}\}$. Figure 1 shows the model fit to a short stretch of ammonium signal, around 11 cycles, from the NGRIP ice-core (Greenland) [2], measured at 1mm intervals via CFA. The reconstruction is shown as a dotted black line where the data is missing, otherwise it matches exactly to the signal.

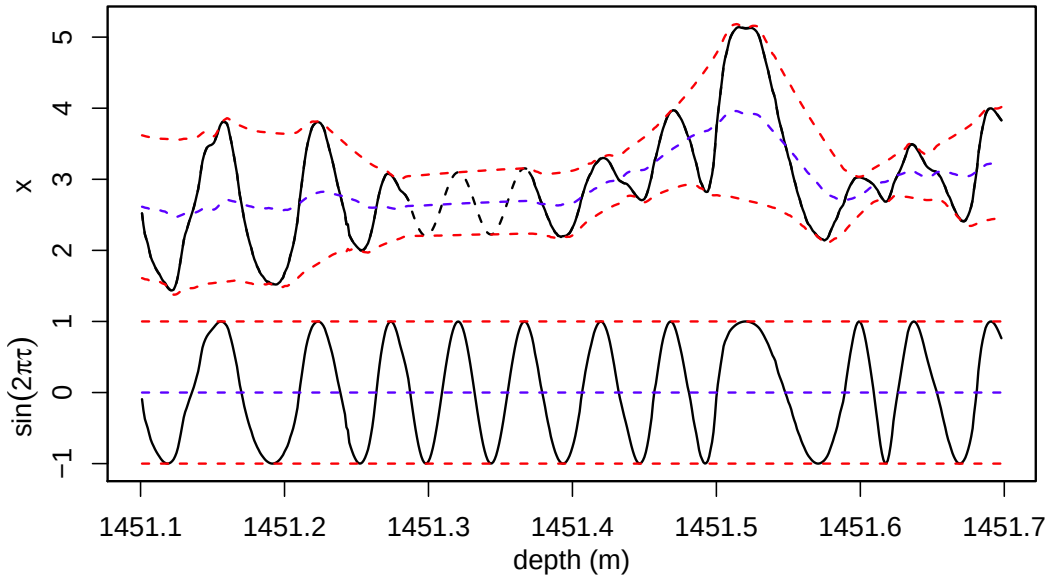

Figure 1: Posterior mean reconstruction of the NGRIP ammonium signal. (top) the mean reconstruction, with $\boldsymbol{\beta}$ as a dotted blue line and $\boldsymbol{\beta} \pm \boldsymbol{\alpha}$ as a dotted red line. (bottom) $\sin(2\pi\boldsymbol{\tau})$.

2.1 Priors

The elapsed times over each depth increment are independently Gamma distributed with shape ψ and rate λ,

$$\tau_i - \tau_{i-1} \sim G(\psi, \lambda),$$

$\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ model the amplitude and mean level of the signal, these are intended to be slow-moving processes and their prior takes the form of two independent Gaussian random walks, for $i \in (2, 3, \ldots, n)$:

$$\alpha_i \sim N(\alpha_{i-1}, \sigma_\alpha^2) \quad \text{and} \quad \beta_i \sim N(\beta_{i-1}, \sigma_\beta^2).$$

3 Implementation

$\boldsymbol{\theta}$ is updated in intervals, $I = \{i \mid s < i < f\}$, chosen uniformly at random. The reconstruction within I is updated conditionally on that outside of I.

3.1 Maintaining cycle count

The conditional distribution of $\boldsymbol{\alpha}$ in I given $\boldsymbol{\alpha}$ outside of I is

$$\boldsymbol{\alpha}_I \mid \boldsymbol{\alpha}_{-I} \sim N(\boldsymbol{\mu}_\alpha, \Sigma_\alpha),$$

where

$$\mu_{\alpha,j} = \alpha_s + \frac{j(\alpha_f - \alpha_s)}{m+1}, \quad j \in (1, 2, \ldots, m),$$

and

$$\Sigma_{\alpha,jk} = \sigma_\alpha^2 \left(\min(j,k) - \frac{jk}{m+1} \right), \quad j,k \in (1, 2, \ldots, m),$$

I containing m data points. Similarly for $\boldsymbol{\beta}_I$. The signal in I is

$$\boldsymbol{x}_I = S\boldsymbol{\alpha}_I + \boldsymbol{\beta}_I,$$

where S is a matrix containing $\sin(2\pi\boldsymbol{\tau}_I)$ along the diagonal. Thus

$$\boldsymbol{x}_I \mid \boldsymbol{\tau}_I, \boldsymbol{\alpha}_{-I}, \boldsymbol{\beta}_{-I} \sim N(S\boldsymbol{\mu}_\alpha + \boldsymbol{\mu}_\beta,\ S\Sigma_\alpha S^T + \Sigma_\beta).$$

$\boldsymbol{\tau}'_I$ is proposed from its prior, conditioned on $\boldsymbol{\tau}_{-I}$, by sampling $\boldsymbol{u}$ – $m+1$ perturbations from a Dirichlet distribution with constant shape ψ – and setting

$$\tau'_{I,j} = \tau_s + (\tau_f - \tau_s) \sum_{k=1}^{j} u_k, \quad j \in (1, 2, \ldots, m).$$

This proposal has acceptance probability

$$\frac{p(\boldsymbol{x}_I | \boldsymbol{\tau}'_I, \boldsymbol{\alpha}_{-I}, \boldsymbol{\beta}_{-I})}{p(\boldsymbol{x}_I | \boldsymbol{\tau}_I, \boldsymbol{\alpha}_{-I}, \boldsymbol{\beta}_{-I})}.$$

If $\boldsymbol{\tau}'_I$ is accepted, S is set to S', and $(S\boldsymbol{\alpha}_I)'$ is drawn from

$$S\boldsymbol{\alpha}_I \mid \boldsymbol{x}_I, \boldsymbol{\tau}'_I, \boldsymbol{\alpha}_{-I}, \boldsymbol{\beta}_{-I} \sim N(\boldsymbol{\mu}_{S\alpha}, \Sigma_{S\alpha}),$$

where

$$\boldsymbol{\mu}_{S\alpha} = S\boldsymbol{\mu}_\alpha + S\Sigma_\alpha S^T (S\Sigma_\alpha S^T + \Sigma_\beta)^{-1} (\boldsymbol{x}_I - S\boldsymbol{\mu}_\alpha - \boldsymbol{\mu}_\beta)$$

and

$$\Sigma_{S\alpha} = S\Sigma_\alpha S^T - S\Sigma_\alpha S^T (S\Sigma_\alpha S^T + \Sigma_\beta)^{-1} S\Sigma_\alpha S^T.$$

Then $\boldsymbol{\alpha}'$ is set to $S^{-1}(S\boldsymbol{\alpha}_I)'$, and $\boldsymbol{\beta}'_I$ to $\boldsymbol{x}_I - (S\boldsymbol{\alpha}_I)'$.

3.2 Changing cycle count

To add a cycle into I the proposal, $\boldsymbol{\tau}'_I$, is conditioned to run between τ_s and $\tau_f + 1$. This adds a term to the acceptance probability,

$$\frac{p(\boldsymbol{\tau}'_I)q(\boldsymbol{\tau}_I)}{p(\boldsymbol{\tau}_I)q(\boldsymbol{\tau}'_I)} = \frac{(\Delta+1)^{(m+1)(\psi-1)}e^{-\lambda(\Delta+1)}}{\Delta^{(m+1)(\psi-1)}e^{-\lambda\Delta}}$$

where $\Delta = \tau_f - \tau_s$, which compares Δ and $\Delta+1$ with respect to the $G((m+1)\psi, \lambda)$ distribution. If this step is accepted, $\boldsymbol{\tau}$ to the right of I are incremented by 1. Cycles can be removed in a similar manner.

3.3 Hyper-parameters

σ_α and σ_β are given uninformative inverse-gamma priors, and updated via Gibbs steps. λ is given an uninformative Gamma prior, and updated via a Gibbs step. ψ is updated via a Metropolis-Hastings step with a flat, improper, prior.

4 Conclusions

Our approach automates the layer-counting process, providing information about the time of year, as well as the date, at each depth. The updates can be easily adapted for missing values – the reconstruction filling the gaps as seen in figure 1. A different approach to this problem can be found in [3]; the method presented here has the advantages that it is fully Bayesian and provides a more detailed chronology.

References

[1] K. K. Andersen, A. Svensson, S. J. Johnsen, S. O. Rasmussen, M. Bigler, R. Röthlisberger, U. Ruth, M. L. Siggaard-Andersen, J. Peder Steffensen, D. Dahl-Jensen, B. M. Vinther, and H. B. Clausen. **The Greenland ice core chronology 2005, 15 42 ka. Part 1: constructing the time scale**. *Quaternary Sci. Rev.*; 2006; 25; pp. 32463257.

[2] D. Dahl-Jensen, N. S. Gundestrup, H. Miller, O. Watanabe, S. J. Johnsen, J. P. Steffensen, H. B. Clausen, A. Svensson, and L. B. Larsen. **The North-GRIP deep drilling programme**. *Ann. Glaciol.*; 2002; 35; pp. 1-4.

[3] J. J. Wheatley, P. G. Blackwell, N. J. Abram, J. R. McConnell, E. R. Thomas, and E. W. Wolff. **Automated ice-core layer-counting with strong univariate signals**. *Climate of the Past*; 2012; 8; pp. 1869-1879.