

A Bayesian Hierarchical Chronology Model for Time Series Analysis of Paleoenvironmental Data

Aaron Springford*

Abstract

Time series data consist of observations collected sequentially in time. The analysis of time series is well developed in cases where the sampling times are known and evenly spaced. However, in some cases we may not observe the times directly and are forced to treat them as latent. Common examples include paleoenvironmental studies based on core samples, such as paleolimnology. Paleolimnological core data consist of sediment samples collected at sequential depths. A subset of the samples are then dated using radio-isotopes or other methods. A chronology model that relates the observed time proxy (depth) to time is required before any time series analysis can be undertaken. Chronology models used in the literature tend to favour piecewise linear formulations, usually by interpolating between depths with available dating estimates. I present a chronology model based instead on modelling the accumulation of the core. A Bayesian hierarchical approach results in probabilistic chronology estimates that are robust to outliers common to radio-isotope analysis, and allows uncertainty in timing to be applied directly to subsequent time series analysis.

Key Words: Time series analysis, Chronology model, Bayesian

1. Introduction

Time series data consist of observations collected sequentially in time. The basic structure of the data is

$$x(t(1)), x(t(2)), \dots, x(t(N))$$

where $t(i) \leq t(j)$ for $i < j$. The representation $x()$ refers solely to data structure rather than any sort of process or function x . In most cases, the $\{t(i), i = 1, 2, \dots, N\}$ are known, and in many cases the time increments $d_i = t(i) - t(i - 1)$ are equal for all $i = 2, 3, \dots, N$. However, it is sometimes the case that the time increments are not all the same, especially in observational data where the observer doesn't control the observation times. Common examples include time series that are evenly sampled but that contain missing values, or event data such as nerve impulses (Brillinger 1984). In other cases, we may not observe the times directly and are forced to treat them as latent. Common examples are paleoenvironmental studies based on extracted cores, which are composed of material that has accrued over time. For example, paleolimnology studies often collect cores from lakebed sediments. The cores are sectioned into depth slices, and the material from the slices are analyzed for proxies of past environmental conditions (Last et al. 2001). Although the slices are often (but not necessarily) uniform in depth, the corresponding time increments are unknown and likely non-uniform. Other common examples of core data include ocean sediments, peat, rock and terrestrial soil, trees, and ice.

Paleoenvironmental research is focused on describing phenomena that vary with time and often have periodic dynamics. Thus, time series modelling and spectral analysis are the statistical tools necessary for description and inference. Techniques for time series modelling and spectrum estimation are well-established for series that are evenly sampled at known intervals. Common models for time series include autoregressive moving average

*Queen's University at Kingston, Department of Mathematics and Statistics, Jeffery Hall, University Avenue, Kingston, Ontario, Canada, K7L 3N6. Email: aaron.springford@queensu.ca

(ARMA) models and state-space models such as Kalman filtering (Brockwell and Davis 2009). When producing spectral estimates, averaging estimates from overlapped segments or from multiple orthogonal tapers reduces variance, and the use of a fast Fourier transform reduces computation time (Percival and Walden 1993).

Core data face particular challenges when it comes to obtaining time series model parameter estimates and spectral estimates. First, the time series data tend to be limited in length due to high sampling costs associated with sectioning and analyzing the core for proxies of interest, leading to estimates with large variance compared to longer series of the same process. Second, the data are irregularly sampled in time. The irregular sampling limits the types of models that can be used, and can also increase computational burden. For example, spectral estimates are more expensive to compute because the fast Fourier transform relies on uniform sampling. Third, the data are sampled at unknown times; the relationship between depth in the core and time is uncertain. Errors in timing result in errors in estimates of model parameters and difficulties estimating the spectrum. In particular, detection of periodic components that manifest as peaks in the spectrum, and the estimation of the frequencies of peaks can suffer when times are unknown. A quantitative way of incorporating this uncertainty into the analysis is required (Blaauw and Heegaard 2012). The balance of this paper is focused on addressing the unknown relationship between depth in the core and age of a sample using a Bayesian hierarchical chronology model.

1.1 Modelling chronologies

Chronology modelling involves estimating times (ages) using at a minimum some sort of age proxy. Times are obviously necessary for the analysis of time series. However, chronology models can vary in their interpretation of the available data, and as a result provide different estimates of the chronology. I propose a short list of characteristics that a chronology model for core data should possess:

1. The ages are monotonically increasing with depth in the core.
2. The ages are robust to errors in dating estimates.
3. The model makes full use of the information contained in the distribution of dating estimates (i.e. the likelihood of the data).
4. Uncertainty in the ages can be propagated to spectrum estimates or estimates of time series model parameters.

Feature (1) is fundamentally a result of the deposition process assumed to have taken place in the core – older layers of the core were deposited before newer ones, and thus must be deeper in the core. Although this assumption might be violated in practice, such violations can be assumed to be evident to trained experts when the core is examined.

Feature (2) is desirable because several factors can influence the dating estimates that are reported by labs. For example, if the core consists of sediments and the dating estimate is inferred from a ^{14}C isotope ratio, then the material chosen for ageing (e.g. a piece of organic material such as a chip of wood) might correspond to a different age than the bulk material at the same depth. This type of error is exacerbated when the variance of the dating estimate as reported by the lab is underestimated.

Feature (3) is desirable because of the somewhat unconventional distributions that can arise when radio-isotope dates, in particular ^{14}C dates, are calibrated. ^{14}C dates must be calibrated because the ambient concentration of ^{14}C is not constant over history (Damon et al. 1978; Stuiver et al. 1998b). In some cases, a particular current-day ^{14}C isotope ratio

maps to multiple possible deposition dates. As a result, the calibrated age distribution for a ^{14}C sample is often multimodal, and the expected value of the calibrated distribution might have a very small or zero probability. In such cases, it seems farcical to use a point value for the calibrated age estimate. Rather, the information contained in the full likelihood is more representative of the distribution of possible dates and is thus more appropriate.

Feature (4) is desirable because the chronology model does not exist in isolation but rather in the larger context of a time series analysis. As mentioned above, if uncertainties in the timing of events are not propagated through to the time series analysis, then the total amount of uncertainty in time series estimates will be too small, and we will be overconfident in our conclusions.

Previous chronology models can be categorized into three general approaches. The first approach is to assume that time increments are piecewise linear and focus on interpolating between available dating estimates. In the simplest case, this amounts to “connecting the dots” between dating estimates, where the “dots” are often the mean or median of the calibrated age likelihoods. The second approach is to fit a function through the dating estimates. The choice of function family is often arbitrary. For example, in the analysis of the motivating dataset taken from Yu et al. (2006), a two-piece linear function was chosen by the authors. Polynomial functions or spline functions are also possible choices. Another related approach is *wiggle match dating* (Pearson 1986; Blaauw et al. 2003), where reversals in uncalibrated ^{14}C dates are used to tune the chronology of the core. The third approach is to model the accumulation of the core using what amounts to a state-space model in which the observations are the dating likelihoods and the hidden process is the deposition of the core (e.g. Bronk Ramsey 2008; Haslett and Parnell 2008). Additional data or information thought to influence the deposition process can also be included if available. Previous approaches of this type have met criticism for allowing unrealistic accumulation rates (Blaauw and Heegaard 2012). In this paper, I present a chronology modelling approach that is based on accumulation of the core, but that avoids unrealistic accumulation rates by assuming that the time increments share a common distribution.

2. Motivating dataset and methods

I take my motivating dataset from a study of climate on the Tibetan Plateau, in particular the strength of summer and winter monsoons in the region (Yu et al. 2006). The data come from a peat core that was sectioned into 649 one-centimetre intervals. The core covers much of the Holocene, roughly 11,000 years. Proxy records from each section were obtained, as well as ^{14}C isotope ratios from a subset of eleven sections. I converted these raw ^{14}C age estimates into likelihoods using the calibration software *Calib6.11* (Stuiver et al. 1998a). The proxy records included peat greyscale and humification, which are thought to be proxies for the strength of summer monsoons, and ash content which is thought to be a proxy for the strength of winter monsoons (Yu et al. 2006).

2.1 Bayesian hierarchical chronology model

In the general case, a chronology model relates a time proxy (usually depth) to time. Core data such as those in the motivating dataset are comprised of depth sections, a subset of which are dated using radio-isotopes or other methods. The data consist of the section depths $\{d_i : i = 0, 1, \dots, N\}$, assumed known, and the ages $\{a_j : j = 1, 2, \dots, M\}$ with corresponding ageing errors $\{\sigma_j : j = 1, 2, \dots, M\}$, from which the likelihoods $\{f(a_j, \sigma_j) : j = 1, 2, \dots, M\}$ for the calibrated ages $\{a_j^c : j = 1, 2, \dots, M\}$ are obtained. From the section depths, the depth increments are the first differences

$\{v_k : k = 1, 2, \dots, N - 1\}$, $v_k = d_k - d_{k-1}$, each with a corresponding (latent) time increment u_k . The focus of the chronology model is to estimate the latent time increments u to obtain times as the cumulative sums $t_k = \sum_{l=1}^k u_l$, $k = 1, 2, \dots, N - 1$. Here I assume that $t_0 = 0$, but t_0 could alternatively be included as an additional parameter. To relate the ageing data with the latent times, let ϕ_j be the index k of t corresponding to the age a_j^c . Modelling the time increments and forcing the condition $u_k \geq 0$ for all k restricts the time-depth relationship to be monotonically increasing.

I developed the prior distribution on the time increments using the following logic: Imagine the accumulation of a depth increment v_k to be the accumulation of α_k smaller increments, each of which has a waiting time that follows an Exponential distribution with rate parameter λ_k (a Poisson process). Then the waiting time u_k for the entire increment v_k will follow a Gamma(α_k, λ_k) distribution. If we standardize the time increments $w_k = u_k/v_k$, then $w_k \sim \text{Gamma}(\alpha_k, \lambda_k/v_k) = \text{Gamma}(\alpha_k, \beta_k)$.

Based on this development, I assume that the standardized time increments follow the prior

$$w_k \sim \text{Gamma}(\alpha, \beta) \text{ for all } k \tag{1}$$

Note that this prior is common to all of the standardized time increments (only one (α, β) rather than k separate prior distributions). The likelihood is the probability distribution for the calibrated ^{14}C ages a_j^c , which are assumed to follow

$$a_j^c \sim f(a_j, \sigma_j) \tag{2}$$

where $a_j^c = t_{\phi_j} = \sum_{l=1}^{\phi_j} u_l = \sum_{l=1}^{\phi_j} w_l v_l$. These likelihoods were obtained using the calibration software *Calib6.11* (Stuiver et al. 1998a), and tend to be highly variable in form from sample to sample; calibrated likelihoods associated with younger ages tend to be better estimated, but are often multimodal.

The hyperpriors for α and β were chosen to be conjugate to the Gamma distribution in equation 1. The gamma distribution as used here is defined

$$\text{Gamma}(w_k | \alpha, \beta) = \beta^\alpha \frac{1}{\Gamma(\alpha)} w_k^{\alpha-1} e^{-\beta w_k}.$$

Given α , the conditional distribution for β is

$$P(\beta | w_1, \dots, w_{N-1}) \propto \begin{cases} \exp\left(-\beta \sum_{k=1}^{N-1} w_k\right) \beta^{\alpha(N-1)} & w_k > 0 \\ 0 & \text{otherwise} \end{cases}.$$

The conjugate hyperprior for β is also a Gamma distribution with hyperparameters $\alpha_1 > 0$ and $\beta_1 > 0$, resulting in a Gamma posterior for β with hyperparameters

$$\begin{aligned} \alpha'_1 &= \alpha(N - 1) + \alpha_1 \\ \beta'_1 &= \frac{\beta_1}{1 + \beta_1 \sum_{k=1}^{N-1} w_k}. \end{aligned} \tag{3}$$

Given β , the conditional distribution for α is

$$P(\alpha | w_1, \dots, w_{N-1}) \propto \begin{cases} \frac{\beta^{\alpha(N-1)}}{\Gamma(\alpha)} \left(\prod_{k=1}^{N-1} w_k\right)^{\alpha-1} & w_k > 0 \\ 0 & \text{otherwise} \end{cases}.$$

The conjugate hyperprior for α is

$$\pi(\alpha | \alpha_2, \beta_2, \gamma_2) \propto \begin{cases} \frac{\alpha_2^{\alpha-1} \beta_2^{\alpha \gamma_2}}{\Gamma(\alpha)^{\beta_2}} & \alpha > 0 \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

with hyperparameters $\alpha_2, \beta_2, \gamma_2 > 0$. The posterior for α has hyperparameters

$$\begin{aligned}\alpha'_2 &= \alpha_2 \prod_{k=1}^{N-1} w_k \\ \beta'_2 &= \beta_2 + (N - 1) \\ \gamma'_2 &= \gamma_2 + (N - 1).\end{aligned}\tag{5}$$

Notice that the posterior for α requires integration of the kernel in equation 4, which has no closed form and must be evaluated numerically.

Numerical issues arise if α becomes large. For example, $\alpha = 100 \implies \Gamma(\alpha) = 9.3 \times 10^{155}$. When posterior weight is given to values of $\alpha > 100$, this computing limitation makes the use of a Gamma prior on w difficult or impossible. However, if α is large the Gamma prior is approximately Normal. For fixed β , large values of α correspond to large numbers of increment layers, and each increment layer's waiting time is assumed to follow an Exponential distribution. Denote the waiting time for each increment layer $\delta \sim Exp(\lambda)$. Then by the Central Limit Theorem,

$$\sum_{i=1}^d \delta_i \xrightarrow{D} N(\alpha/\beta, \alpha/\beta^2)$$

as $d \rightarrow \infty$, where \xrightarrow{D} means converges in distribution. Holding β constant, $\alpha \rightarrow \infty$ as $d \rightarrow \infty$. Thus, when posterior weight is given to values of $\alpha > 100$ (and computational issues arise), a Normal prior for the standardized latent time increments becomes suitable

$$w_k \sim N(\mu, 1/\tau^2) \text{ for all } k\tag{6}$$

where μ is the mean and τ^2 is the information (reciprocal of variance) of the distribution. Equation 6 replaces the Gamma prior (equation 1).

Hyperpriors for the Normal prior (equation 6) are on μ and τ^2 . Given τ^2 , the conditional distribution for μ is

$$P(\mu|w_1, \dots, w_{N-1}, \tau^2) \propto \exp\left(-\frac{(N-1)\tau^2}{2} \left(\mu - \frac{1}{N-1} \sum_{k=1}^{N-1} w_k\right)^2\right).$$

The conjugate hyperprior is

$$\pi(\mu|\mu_1, \tau_1^2) = \frac{\tau_1}{\sqrt{2\pi}} \exp\left(-\frac{\tau_1^2}{2}(\mu - \mu_1)^2\right)\tag{7}$$

with hyperparameters μ_1 and τ_1^2 . The posterior for μ has hyperparameters

$$\begin{aligned}\mu'_1 &= \frac{\mu_1 \tau_1^2 + \tau^2 \sum_{k=1}^{N-1} w_k}{\tau_1^2 + (N-1)\tau^2} \\ \tau_1^{2'} &= \tau_1^2 + (N-1)\tau^2\end{aligned}\tag{8}$$

Given μ , the conditional distribution for τ^2 is

$$P(\tau^2|w_1, \dots, w_{N-1}, \mu) \propto \tau^{N-1} \exp\left(-\frac{\tau^2 \sum_{k=1}^{N-1} (w_k - \mu)^2}{2}\right).$$

The conjugate hyperprior is a Gamma distribution, however this hyperprior choice is not recommended in general because it is difficult to make weakly informative with respect to the posterior (Gelman 2006). A better weakly informative hyperprior is the uniform distribution on τ^{-1}

$$\pi(\tau^2|\sigma_L, \sigma_U) = \begin{cases} \frac{1}{\sigma_U - \sigma_L} & \sigma_L \leq \tau^{-1} \leq \sigma_U \\ 0 & \text{otherwise} \end{cases}$$

with hyperparameters σ_L and σ_U , $0 \leq \sigma_L < \sigma_U$.

For this hyperprior, the posterior is

$$\pi(\tau^2|\sigma_L, \sigma_U, \mu, w_1, \dots, w_k) \propto \begin{cases} \tau^{N-1} \exp\left(-\frac{\tau^2 \sum_{k=1}^{N-1} (w_k - \mu)^2}{2}\right) & \sigma_L \leq \tau^{-1} \leq \sigma_U \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

which can be sampled numerically.

2.2 Sampling chronology model posterior parameter distributions

The conditional definition of the chronology model in the previous section makes computation of the posterior distribution for model parameters using a Gibbs sampler straightforward. I used a Gibbs sampler with parameter-specific update methods depending on the availability of an analytic posterior (Gelman et al. 2003). For those parameters with analytically integrable posterior distributions, parameter updates are draws from the conditional posterior directly since conditional posterior parameters are available using an ordinary Gibbs step (Geman and Geman 1984). For the latent standardized time increments w_1, \dots, w_{N-1} , I used a Metropolis update step (Metropolis et al. 1953). For the α hyperparameter or the τ^2 hyperparameter (depending on the choice of prior on w), I used a slice sampler (Neal 2003).

I began the chain near the expected posterior values for the model parameters, and convergence to the posterior was rapid. I burned in the chain for 20,000 samples. Following burn-in, I sampled 5,000,000 iterations of the Markov chain, thinning to 1,000 to save storage space due to strong autocorrelation and poor mixing in the chain. The poor mixing is presumably due to strong anti-correlation between the w parameters. Although computation was relatively quick (on the order of minutes) adaptive tuning of the Metropolis updates and/or a reparameterization of the model might result in improvements to sampling efficiency and speed of computation.

3. Results

The posterior distributions for all model parameters are well-identified (figure 1). In particular, it is evident that the posterior for w is not Normal - the age likelihoods have had a noticeable effect on the prior. The resulting uncertainties in the chronology are evident when the time series for the three proxies are plotted (figure 2). In particular, times towards the end of the series and between approximately 2800 and 5800 years before present have high associated uncertainty, whereas times between 5800 and 6700 years before present are strongly constrained by the likelihoods of the ^{14}C data at 305 and 321 centimetres depth.

The next step in a time series analysis of these data would be to obtain estimates of time series model parameters and/or spectra using the times from the 1000 sampled chronologies. The advantage of the Bayesian approach here is obvious since posterior parameter estimates and/or posterior spectrum estimates can be calculated directly from the posterior

Figure 1: Posterior distributions, plotted as probability densities, for the mean age increment μ , the age increment standard error τ^{-1} , the age increments w (left to right, top panels). Posterior age versus depth trajectories based on 1000 posterior samples (bottom panel). Also plotted with the posterior age versus depth trajectories are violin plots (in orange) of the likelihood from calibrated ^{14}C dating.

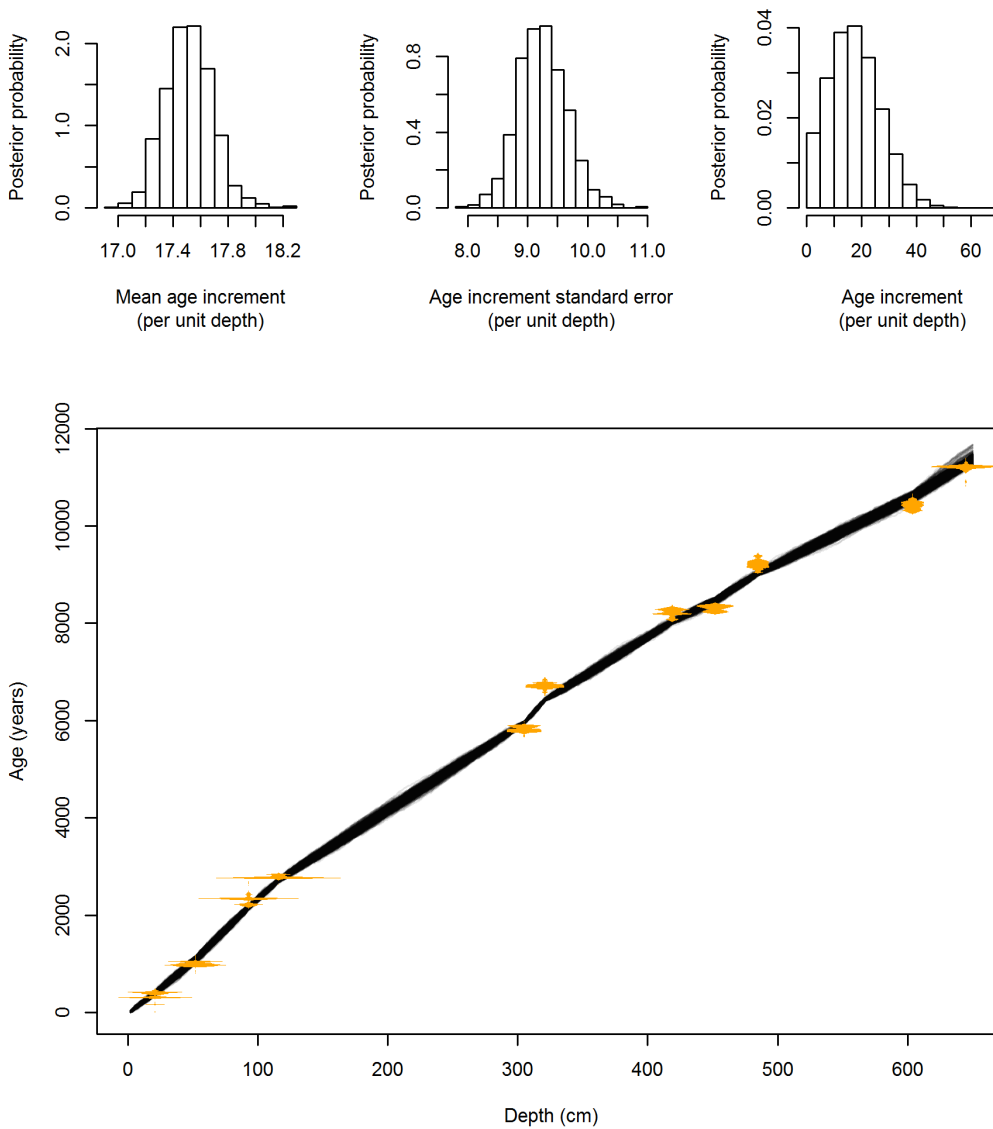
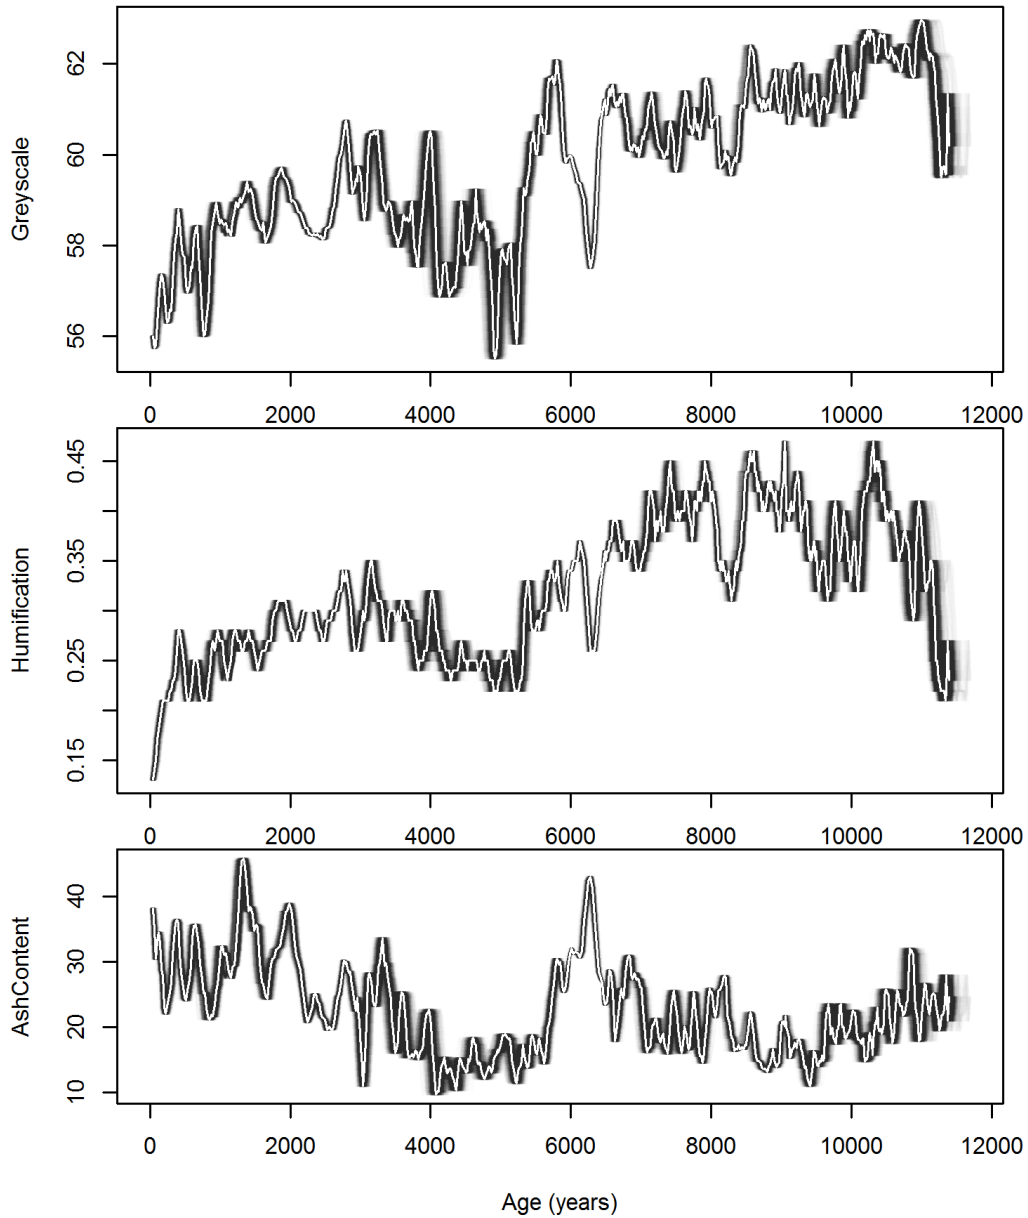


Figure 2: One thousand posterior time series of equal probability (light grey lines, over-plotted) and the mean posterior time series (white line). The top panel are the time series for the Greyscale measurements, the middle panel Humification measurements, and the bottom panel Ash Content.



chronology samples, assuming independence between the environmental proxies and the core chronology.

4. Discussion

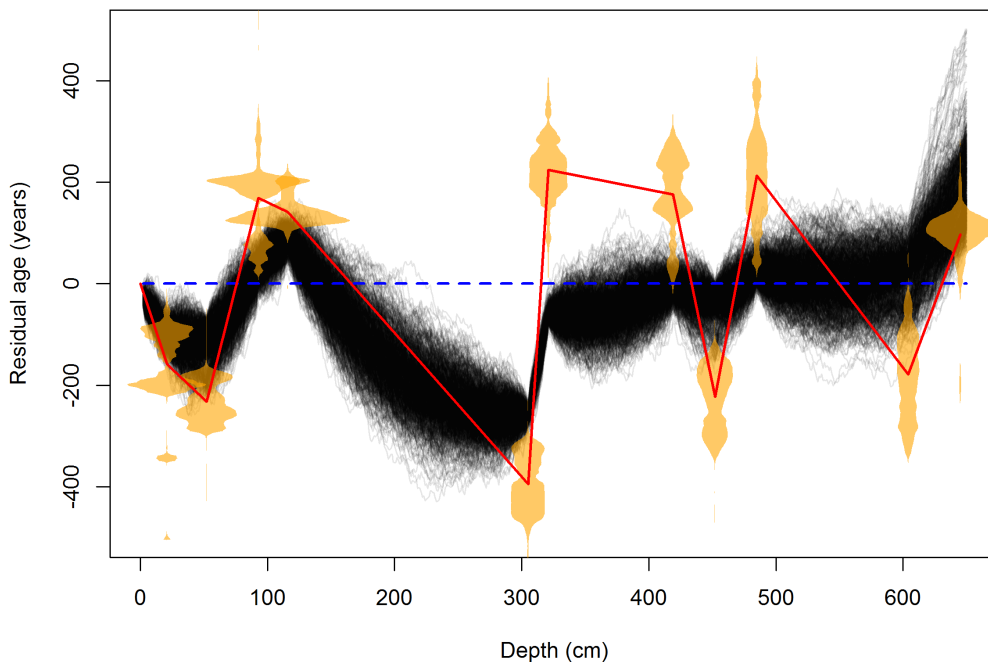
The chronology model presented here was constructed with the four key features discussed in section 1.1 in mind: (1) Monotone chronology; (2) Robust to outliers; (3) Uses the complete dating likelihood; (4) Uncertainty can be propagated to time series analysis. By modelling the time it takes to deposit a standardized depth increment, and restricting these time increments to be positive, we are able to guarantee a monotonically increasing time series (feature 1). Note that some other approaches, in particular simple piecewise linear approaches, cannot guarantee monotonically increasing times and thus can be problematic for subsequent time series analysis. The Bayesian approach includes the entire likelihood of the dating estimates directly (feature 3), rather than a summary such as the expected value or median value of the dates. The approach also results in a posterior probability distribution for each of the latent ages which can be used in subsequent time series analysis of the environmental proxies (feature 4).

Chronology methods used in the literature tend to favour piecewise linear formulations (e.g. Blaauw and Christen 2005; Blaauw et al. 2007). However, a very similar model to the one proposed here was devised recently by Bronk Ramsey (2008). The similarity is in modelling the accumulation of the core. However, Bronk Ramsey presents the problem in a slightly different way. Rather than model the depth increments, the focus is on a series of times. Interpolation between times is done uniformly or using a Poisson number of small depth increments, the thickness of which must be set by the user. Another approach equivalent to Bronk Ramsey's where the number of accumulation increments is Poisson, but the increment size is random (and thus the age versus depth relationship is piecewise linear) was devised by Haslett and Parnell (2008). However, it does not allow for information on accumulation rate to be included, nor does it model accumulation rates as a group, resulting in unrealistic accumulation rates in some cases (Blaauw and Heegaard 2012).

The problem of unrealistic accumulation rates is handled automatically in the Bayesian hierarchical model presented here. According to Bayes' rule, the posterior distribution for the time increments is a combination of the prior distribution and the dating likelihoods. Note that only the prior distribution family is defined *a priori* in the hierarchical model; the prior distribution for the time increments is also a combination of hyperprior and likelihood. Thus, the hierarchical model is essentially a form of regularization with shrinkage of the individual time increments automatically controlled by the estimated prior distribution (Gelman et al. 2003; Carlin and Louis 2008).

Estimators that exhibit shrinkage behaviour are well known to have reduced sensitivity to outliers (Box 1980; Gelman et al. 2003). In the context of the current model, the shrinkage behaviour is evident when comparing the chronology obtained from the Bayesian hierarchical model to the chronology obtained from a cubic polynomial fit to the mean likelihoods and the chronology obtained from a piecewise linear connect-the-dots fit to the mean likelihoods (figure 3). Examining the chronologies, it is clear that the Bayesian hierarchical model chronology shrinks the chronology trajectory from the individual likelihoods towards the average chronology represented by the cubic polynomial fit. Considering the situation in which one of the likelihoods represents an outlier, the average chronology will move towards the outlier but only slightly due to the influence of the rest of the data, and the shrinkage behaviour will constrain the posterior chronology accordingly. Thus, the Bayesian hierarchical model will have limited sensitivity to outliers (feature 2), which can be evaluated on a case-by-case basis.

Figure 3: Posterior age versus depth trajectories based on 1000 posterior samples (black), plotted as residuals from a cubic polynomial fit to the mean likelihoods of the calibrated ^{14}C dates. Also plotted are the violin plots (orange) of the likelihoods from calibrated ^{14}C dating, a dashed zero-line (blue) representing the cubic polynomial fit, and a piecewise linear fit to the means of the calibrated likelihoods. The Bayesian hierarchical trajectories can be seen as a compromise between the cubic polynomial fit and the piecewise linear fit.



Using standard Markov chain Monte Carlo (MCMC) techniques to sample the posterior resulted in excessive computation time due to poor mixing of the chain for the time increments w . The w were sampled using a Metropolis step, and the Metropolis algorithm requires some tuning of the proposal distribution to achieve efficient sampling. Improvements to sampling efficiency are possible using adaptive Metropolis algorithms such as Haario et al. (2001) that tune the proposal distribution as sampling progresses, and could prove useful due to the strong anticorrelations expected for w .

Although not considered here due to the choice of motivating dataset, the analysis of varved data would require modification to the likelihood to allow varve layer counts to be included. Varved core data exhibit periodic banding, often associated with processes that operate on an annual cycle, and this data should be included in the likelihood. If, for example, the core was varved but sectioned at independent (possibly regular) intervals, then the number of varve layer counts in each section could be included as binomial likelihoods with the true number of varve layers and the probability of detection being parameters. Although an enticingly simple starting point, there are at least two potential issues with this choice of likelihood for the varve layer data: (1) The identification of neighbouring layers is likely not the product of two independent Bernoulli trials – if a single layer is missed then the following layer is almost surely not; (2) The true number of varve layers and the probability of detection are likely both mediated by the deposition rate. The sensitivity of results to the simple binomial assumption needs to be explored before extending the

Bayesian hierarchical approach to varved core records.

5. Conclusion

Time series analysis is complicated when the times themselves are unknown, as is the case for paleoenvironmental core data. In such cases, a chronology model that relates the unobserved times to depth in the core is a necessary first step for time series analysis. In this paper, I present a Bayesian hierarchical approach that has several advantages compared to previous approaches. A preliminary analysis shows that the posterior distribution for model parameters is easily sampled, and that sampling efficiency gains are no doubt possible. The availability of the posterior distribution makes subsequent time series analysis straightforward. Finally, the model can be extended to cases where additional information in the form of varved records is available with a modification to the data likelihood.

6. Acknowledgement

This work was funded in part by a scholarship from the National Science and Engineering Research Council of Canada.

References

- Blaauw, M., Bakker, R., Christen, J., Hall, V., and Plicht, J. (2007), "A Bayesian framework for age modeling of radiocarbon-dated peat deposits: case studies from the Netherlands," *Radiocarbon*, 49, 357–368.
- Blaauw, M. and Christen, J. (2005), "Radiocarbon peat chronologies and environmental change," *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 54, 805–816.
- Blaauw, M. and Heegaard, E. (2012), *Estimation of Age-Depth Relationships*, Springer, chap. 12, Developments in Paleoenvironmental Research.
- Blaauw, M., Heuvelink, G., Mauquoy, D., van der Plicht, J., and van Geel, B. (2003), "A numerical approach to ^{14}C wiggle-match dating of organic deposits: best fits and confidence intervals," *Quaternary Science Reviews*, 22, 1485–1500.
- Box, G. E. (1980), "Sampling and Bayes' inference in scientific modelling and robustness," *Journal of the Royal Statistical Society. Series A (General)*, 383–430.
- Brillinger, D. (1984), "Statistical inference for irregularly observed processes," in *Time Series Analysis of Irregularly Observed Data*, ed. Parzen, E., Springer, New York, pp. 38–57.
- Brockwell, P. and Davis, R. (2009), *Time Series: Theory and Methods*, Springer Series in Statistics, Springer.
- Bronk Ramsey, C. (2008), "Deposition models for chronological records," *Quaternary Science Reviews*, 27, 42–60.
- Carlin, B. and Louis, T. (2008), *Bayesian Methods for Data Analysis, Third Edition*, Texts in Statistical Science, Taylor & Francis.

- Damon, P. E., Lerman, J. C., and Long, A. (1978), "Temporal Fluctuations of Atmospheric ^{14}C : Causal Factors and Implications," *Annual Review of Earth and Planetary Sciences*, 6, 457.
- Gelman, A. (2006), "Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper)," *Bayesian analysis*, 1, 515–534.
- Gelman, A., Carlin, J., Stern, H., and Rubin, D. (2003), *Bayesian data analysis*, Chapman & Hall/CRC.
- Geman, S. and Geman, D. (1984), "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 721–741.
- Haario, H., Saksman, E., and Tamminen, J. (2001), "An adaptive Metropolis algorithm," *Bernoulli*, 223–242.
- Haslett, J. and Parnell, A. (2008), "A simple monotone process with application to radiocarbon-dated depth chronologies," *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 57, 399–418.
- Last, W. M., Smol, J. P., and Birks, H. J. B. (2001), *Tracking environmental change using lake sediments*, vol. 1-004, Dordrecht: Kluwer Academic Publishers.
- Metropolis, N., Rosenbluth, A., Rosenbluth, M., Teller, A., and Teller, E. (1953), "Equation of state calculations by fast computing machines," *The Journal of Chemical Physics*, 21, 1087.
- Neal, R. (2003), "Slice sampling," *Annals of Statistics*, 705–741.
- Pearson, G. (1986), "Precise calendrical dating of known growth-period samples using a 'curve fitting' technique." *Radiocarbon*, 28, 292–299.
- Percival, D. and Walden, A. (1993), *Spectral Analysis for Physical Applications*, Cambridge University Press.
- Stuiver, M., Reimer, P., and Braziunas, T. (1998a), "High-precision radiocarbon age calibration for terrestrial and marine samples," *Radiocarbon*, 40, 1127–1151.
- Stuiver, M., Reimer, P. J., Bard, E., Beck, J. W., Burr, G., Hughen, K. A., Kromer, B., McCormac, G., Plicht, J. v. d., and Spurk, M. (1998b), "INTCAL98 Radiocarbon Age Calibration, 24,000-0 cal BP" .
- Yu, X., Zhou, W., Franzen, L., Xian, F., Cheng, P., and Jull, A. (2006), "High-resolution peat records for Holocene monsoon history in the eastern Tibetan Plateau," *Science in China Series D: Earth Sciences*, 49, 615–621.