

ELASTIC TIE-POINTING—TRANSFERRING CHRONOLOGIES BETWEEN RECORDS VIA A GAUSSIAN PROCESS

Timothy J Heaton¹ • Edouard Bard² • Konrad A Hughen³

ABSTRACT. We consider a general methodology for the transferral of chronologies from a master reference record containing direct dating information to an undated record of interest that does not. Transferral is achieved through the identification, by an expert, of a series of tie-points within both records that are believed to correspond to approximately contemporaneous events. Through tying of the 2 records together at these points, the reference chronology is elastically deformed onto the undated record. The method consists of 3 steps: creation of an age-depth model for the reference record using its direct dating information; selection of the tie-points and translation of their age estimates from the reference to the undated record; and finally, creation of an age-depth model for the undated record using these uncertain tie-point age estimates. Our method takes full account of the uncertainties involved in all stages of the process to create a final chronology within the undated record that allows joint age estimates to be found together with their credible intervals. To achieve computational practicality, we employ a Gaussian process to create our age-depth models. Calculations can then be performed exactly without resort to extremely slow Monte Carlo methods involving multiple independent model fits that would be required by other age-depth models.

1. INTRODUCTION

When working with an archaeological or paleoenvironmental record, it is important that it has a reliable chronological timescale in order that the information it contains can be placed correctly in context. In many instances, however, creating this timescale directly from the record itself may not be possible, and alternative approaches may be required where age information is borrowed from an accepted reference record by identification of shared events observed within both records. Specifically, suppose that we are presented with a depositional record of interest that does not possess any direct dating information of its own but is able to be matched or tied at certain depths to a reference record that does have age information. In this article, we aim to transfer the dating information from the reference record through these identified tie-points to create a full chronology for the undated record.

The above problem is sometimes called tuning and has been studied by several other authors who present a range of methods (e.g. Martinson et al. 1982; Lisiecki and Lisiecki 2002; Haam and Huybers 2010). However, as far as we are aware, none of these provide a mathematically rigorous approach to quantifying the sources of uncertainties involved and their impact on the final chronology. This article presents a transferral method that we believe provides a clear and justifiable handling of potential uncertainties throughout the transfer process. This is achieved by the careful consideration of the various sources of error and dependencies that will always be present in such transfers but have hitherto often been ignored. In addressing how these uncertainties should be passed through the process, we hope to improve the quality of the resultant chronology and, as shown in Section 3.2, enable the user to access much more information about that chronology than simply individual estimates. In addition, our method is quick to implement and intuitively simple to understand.

Our algorithm contains 3 intuitive and natural steps. First, the dating information within the reference record is used to create an age-depth model for that deposit. Second, a set of tie-points is identified in both records. These tie-points are typically thought to correspond to approximately contempora-

¹School of Mathematics and Statistics, University of Sheffield, Sheffield S3 7RH, United Kingdom. Corresponding author. Email: t.heaton@sheffield.ac.uk.

²CEREGE, Aix-Marseille University, CNRS, IRD, Collège de France, Technopole de l'Arbois BP 80, 13545 Aix en Provence Cedex 4, France.

³Department of Marine Chemistry & Geochemistry, Woods Hole Oceanographic Institution, Woods Hole, Massachusetts 02543, USA.

neous events, but, as explained later, this is not strictly required. Within the reference record, estimates can be found for the ages of these tie-points from the previous modeling step. Finally, these ages are tied to the corresponding points on the undated record, leaving us with a set of depths at which we have transferred age estimates. These transferred ages are used to create the chronology for the undated record, making sure to take account of their uncertainty as outputs of the reference age-depth model. We implement our approach in R (R Core Team 2013) and code is available on request.

The method's main novelty lies in recognition that, since they arise from the reference record's age-depth model, the ages transferred to the undated record at the tie-points are uncertain. These values cannot simply be treated as though they were observed data but must be integrated out carefully to create the final chronology. While in theory this integration can be done for any form of age-depth model, practical considerations mean that we use a Gaussian process (GP). This allows the integration to be performed quickly and exactly. Other age-depth models would require numerical Monte Carlo integration whereby the undated chronology would have to be independently refitted many times. Since each such fit may take several minutes to run, this would be computationally infeasible. A fuller explanation of this can be found in Section 3.

Using a GP also allows us to incorporate and express dependency information simply, which we will show can be important in both the transfer and later inference. Such dependence will always arise in the uncertainties of age estimates in an age-depth model—neighboring depths will have similar ages. The dependence arising from the reference record's model should be passed through the tie-points and used in creating the final chronology, and this final chronology will itself possess dependencies we would like to quantify. Our approach is capable of this as well as permitting dependence within the observed dates of the reference record that form its initial age-depth model, for example, if they are varve counted.

The intention of this article is to present a general methodology as opposed to a specific case study. As a consequence, to aid comprehension of the underlying ideas, we illustrate our work with a stylistic example to avoid the distraction of describing how ties between 2 records may be identified in practice (which will typically be problem-specific). We discuss briefly the application of our method on real data in sections 5 and 6, but for a practical example including justification for tie selection, please see the accompanying Bard et al. (2013, this issue) article providing timescales for both the Pakistan and Iberian margins. The method has also been used to update the Cariaco Basin timescales from Hughen et al. (2006)—the ties used were as in 2006 but with an updated reference record chronology (Edwards et al., submitted). These 3 timescales with their corresponding uncertainties have gone into creation of the IntCal13 and SHCal13 international calibration curve estimates (Reimer et al. 2013; Hogg et al. 2013, this issue).

We set out our work as follows. In Section 2, we introduce the problem we are aiming to solve and present some useful notation. Next, as the concept of a GP may be new, Section 3 gives a short introduction of its application and explains why it offers benefits for this problem when compared with the other age-depth models currently available such as OxCal (Bronk Ramsey 2008), Bacon (Blaauw and Christen 2011), and BChron (Haslett and Parnell 2008). Here, we also give an illustration of the importance of dependence in age-depth modeling and how it can significantly improve later inference. Section 4 presents the main 3-step algorithm used to transfer the timescale from the reference record to the new record formally setting out how the reference age-depth model is fitted; the procedure by which the tie-points are chosen and information passed across; and finally, the creation of the new chronology.

While we use a stylized example throughout this article, we do not mean this to suggest that selecting ties between records is trivial, and the success of our method depends crucially upon reliable identification of such a series of tie-points. Since such selection is often problem-specific, we do not provide rules as to how these points should be chosen. However, in Section 5 we present some general comments based on our own experiences and discuss the potential difficulties of choosing these matchings. As selection of tie-points can be difficult, any successful method should be robust. In Section 6, we investigate the effect that changes in both tie-point selection and the certainty on their contemporaneity have on the final chronology and demonstrate that our approach is robust to such issues. Finally, we conclude with Section 7 discussing potential further work. Also included are 2 technical appendices. Appendix A summarizes the mathematics behind a GP, while Appendix B gives the specific details of how our model is fitted and the uncertainty in the tie-point ages is integrated out.

2. THE PROBLEM

2.1 Transferring Timescales

Suppose, as illustrated in Figure 1, we are presented with a depositional record of interest for which we wish to create a chronology. This record does not in itself possess any direct quantitative information on its timescale but can nevertheless be matched to a reference record, which, at a set of depths d_1, \dots, d_n , does have dating information. Specifically, we assume that this matching is made by the expert identification of several points in both records (denoted by p_j and \tilde{p}_j for $j = 1, \dots, K$), which are either thought to be approximately contemporaneous or separated by a known time period. Typically, these elastic ties may be found by comparison of some signal simultaneously recorded in both records, which might change in value according to exposure to linked external events (e.g. we might use $\delta^{18}\text{O}$ to identify shared climatic changes). We aim to transfer the dating information from the reference record through these tie-points to create a complete chronology for the undated record.

The reader should be aware that the matching displayed in Figure 1 with a set of clear, separated peaks has been simplified greatly to aid comprehension of the principal ideas. In reality, the identification of the tie-points is itself a complex problem since we would typically expect substantial “noise” to be present within the signals used to match between reference and undated records. This noise may not solely be random measurement fluctuation but could also represent uncorrelated structures linked to genuine, but local, events not shared by both records. To complicate matters more, the size of these uncorrelated artifacts may be as large as those of the shared signal. More detail on these issues can be found in Section 5 together with some recommendations from the authors. The main aim of this work is not, however, to present arguments on matchings but instead for the general exposition of the transferral method. As such, in the rest of the article, we assume that such a matching has been determined by an external expert accompanied with uncertainties as defined in Section 2.2.

2.2 Formalization and Notation

We start by setting up the problem and describing our assumptions. For ease of understanding, we will define the true calendar age at depth d within the directly dated reference record as $\theta_R(d)$, while the true calendar age at depth x within the undated record will be denoted by $\theta_N(x)$.




Figure 1 An illustration of the general problem. The top plot shows our undated record for which we wish to find age estimates at depths x_1, \dots, x_m . The bottom plots shows the reference record at which we have observed age estimates at depths d_1, \dots, d_n . The peaks represent a stylized signal recording shared external events, allowing us to tie the 2 records together at a set of depths p and \tilde{p} .

Reference Record

Within our reference record, we suppose that we have imprecise age observations T_1, \dots, T_n at a set of depths d_1, \dots, d_n according to the model

$$T_i = \theta_R(d_i) + \varepsilon_i^R, \quad i = 1, \dots, n$$

Here, the noise ε_i^R is assumed normally distributed but is not required to be independent and may possess some dependence structure. Thus, as well as the typical ^{14}C or U-Th dating, we also permit varve counting and wiggle-matching (see Reimer et al. 2009) as dating tools.

Undated Record of Interest

Within the undated record, we aim to create a complete chronology via an age-depth model. For any chosen set of depths x_1, \dots, x_m , we desire not only a point estimate of the ages $\theta_N(x_i)$ but also a measure of the uncertainty surrounding those estimates and the dependence between them.

Tie-Points

We suppose we are able to manually identify K matching tie-points between the 2 records at depths p_1, \dots, p_K within the reference and $\tilde{p}_1, \dots, \tilde{p}_K$ within the undated record. These ties may typically represent significant events that have been observed by both records and can be identified by changes in the records themselves. We expect these elastic ties to thus represent either contemporaneous points in both records or points that are separated by some known period. We do not require that these ages are exactly contemporaneous but allow the possibility that some uncertainty exists so long as its size can be quantified, for example, if we believe that the event recorded at depth p_i in the reference record is within ± 100 yr of the event recorded at depth \tilde{p}_i in the undated record. More precisely, we model

$$\theta_R(p_i) = \theta_N(\tilde{p}_i) + \alpha_i + \varepsilon_i^M$$

where α_i is the expected offset between the 2 recorded events (zero if the events are thought contemporaneous) and ε_i^M denotes the matching error. These matching errors may themselves be dependent, although, for simplicity, in this article we assume independence with each $\varepsilon_i^M \sim N(0, \sigma_i^2)$ where σ_i is specified by the expert performing that matching.

3. A GAUSSIAN PROCESS AND OTHER AGE-DEPTH MODELS

3.1 Choosing an Age-Depth Model

There are a wide range of increasingly complex age-depth modeling techniques available that could in principle be employed in our methodology, including OxCal (Bronk Ramsey 2008), Bacon (Blaauw and Christen 2011), and BChron (Haslett and Parnell 2008). A good summary of these current techniques can be found in the review article by Parnell et al. (2011). However, in this article, we do not consider their use since they would make the method computationally infeasible for most users. Instead, we employ a GP allowing fast transfers.

The reason for this speed differential lies in Equation B1 of Appendix B, which, after simplification, becomes

$$\pi(\theta_N | \mathbf{x}, \mathbf{T}, \mathbf{d}, \mathbf{p}, \tilde{\mathbf{p}}) = \int \pi(\theta_N | \theta_R(\mathbf{p}), \mathbf{x}, \tilde{\mathbf{p}}) \pi(\theta_R(\mathbf{p}) | \mathbf{p}, \mathbf{T}, \mathbf{d}) d\theta_R(\mathbf{p})$$

Here, $\pi(\theta_R(\mathbf{p}) | \mathbf{p}, \mathbf{T}, \mathbf{d})$ represents the age distribution of the tie-point ages identified in the reference record's age-depth model and $\pi(\theta_N | \theta_R(\mathbf{p}), \mathbf{x}, \tilde{\mathbf{p}})$ represents the age distribution of the undated record's age-depth model given the ages of the ties in the reference. In our aim to find the LHS, we therefore need to integrate out the uncertainty in the reference record ages $\pi(\theta_R(\mathbf{p}) | \dots)$ over the undated age-depth model $\pi(\theta_N | \theta_R(\mathbf{p}), \dots)$. This can be done in 2 ways, exactly if we use a GP or by numerical Monte Carlo integration with the other techniques mentioned above.

A Gaussian Process

With a GP, both $\pi(\theta_R(\mathbf{p}) | \mathbf{p}, \mathbf{T}, \mathbf{d})$ and $\pi(\theta_N | \theta_R(\mathbf{p}), \mathbf{x}, \tilde{\mathbf{p}})$ follow multivariate normal distributions. As a consequence, this integral can be calculated exactly as described in Appendix B. Such calculation is immediate and hence the entire transfer can be performed extremely quickly.

While allowing rapid calculation, a GP still allows us easy access to, and use of, the covariance information inherent in the model. Fitting any age-depth model will introduce a significant dependence structure in its resultant age estimates—accuracy of the estimate at one point will depend upon accuracy at nearby ages. The age-depth model in the reference record will have such dependence and this should be passed through the transfer via Equation B1. The final chronology will also have dependence, which we would like to access. A GP permits both of these to happen. We discuss the importance of such dependence in Section 3.2.

Other Age-Depth Models

If the age-depth model used for either the reference or undated record results in a joint distribution for the ages that is not normally distributed, as is the case for OxCal, Bacon, and BChron, then this integral cannot be found exactly but instead must be estimated by Monte Carlo. This would require the following iterative process:

- Randomly sampling from the reference age-depth model a set of possible values for the tie-point ages $\theta_R(\mathbf{p})|\mathbf{p}, \mathbf{T}, \mathbf{d}$.
- Transferring these sampled reference ages onto the undated record, treating them as the “hypothetically true” ages of the ties in the reference record and adding the matching uncertainty.
- Fitting an age-depth model $\theta_N|\theta_R(\mathbf{p})$ to the undated record using these sampled reference ages.

We would need to repeat this process a large number of times, on each pass sampling a different set of possible values from the reference record, transferring those sampled ages to the undated record, and fitting a new age-depth model on the undated record. We would complete the estimation by averaging the multiple age-depth models formed for the undated record. Since each fit of the more complex alternative age models may typically take several minutes, and we would need to loop through the above sufficiently many times to fully account for the uncertainty in the final distribution, this rapidly becomes computationally impractical. In addition, it is not clear if the current implementation of these alternative models permit their input data to have any covarying age information. As such, reference records dated by varve counting or wiggle-matching may not be usable.

3.2 The Importance of Dependence

To illustrate the potential importance of including dependence in age-depth modeling, we provide 2 examples. In the first, we consider its ability to significantly improve inference reliant on joint age estimates, for example, if one wished to know the time elapsed between 2 depths in the undated record. The second shows the effect that passing the covariance information arising in the initial age-depth model formed within the reference record through the transfer can have on the final chronology.

3.2.1 Joint Age Estimation

In creating a chronology for the Pakistan Margin (Bard et al. 2013, this issue), there was particular interest in the ages at 2 nearby depths, 8.795 and 8.805 m, since they also had radiocarbon measurements to be included into the IntCal project (Reimer et al. 2013, this issue). Our model predicts the age at these depths to be 18.34 ± 0.35 and 18.36 ± 0.35 kBP, respectively. Here, kBP refers to kyr before 1950 and we have reported the credible intervals as $\pm 2\sigma$.

The model provides much more information in its dependence structure than simply these 2 intervals, however. In Figure 2, we plot the joint posterior credible intervals for the ages at these depths including and excluding the covariance information available from the age-depth model. In these plots, we have excluded the density that does not satisfy the required order constraints according to the Bayesian procedure specified in Section 3.3. If we do not include any covariance information, then we would believe there is an approximately 95% posterior probability that the 2 joint depths lie in the red bordered area—the age at depth 8.795 m is completely independent of the age at depth 8.805 m other than respecting the required ordering. Alternatively, using the covariance information in the age-depth model then we can narrow this to the blue bordered region. There is a lot of dependence between the 2 ages as shown by the very narrow shape of this region. They are still uncertain but must vary together in such a way that the difference between the 2 age estimates must be small. Suppose furthermore that we are interested in the time period that has elapsed between the 2 depths. This is shown on the right-hand side of Figure 2. If the covariance from the age-depth model is used, we can be much more precise about the age difference.




Figure 2 The effect of covariance within age-depth models. The left plot shows the approximate 95% posterior credible interval for the joint estimate of 2 depths both if covariance information is used (blue bordered area) and not used (red bordered area). The right-hand plots show the estimated difference in ages between the 2 depths, the top plot if covariance is ignored and the bottom if it is incorporated.

3.2.2 Passing Dependence through the Transfer

Our second illustration uses our stylized example data set of Figure 1. In fitting our age-depth model to the reference record based upon the external age estimates at depths d_1, \dots, d_7 , we will create dependence between the age estimates at the 8 tie-points that should be transferred through to the undated record and used when creating its chronology. We show the final chronologies obtained for the undated record in Figure 5. The blue and red dashed lines show the posterior credible intervals according to whether such information has been passed or not. While the posterior mean age will always be the same with or without covariance transfer, the uncertainty on the estimated age at any individual depth can be visibly different.

3.3 A Gaussian Process—Benefits and Weaknesses

As explained, the main reason for using a GP as our age-depth model is its ability to make the transfer computationally feasible while still incorporating covariance information. However, it should also be a realistic model. GPs have not been specifically designed to fit age-depth sedimentation models but are instead general techniques for non-parametric regression. As such, they do not assume a fixed form for the deposition/growth rate but allow the data to “speak” for themselves. If required, they provide considerable flexibility in the level of prior information about the age-depth model. In performing our transfers, we have made the assumption that the sedimentation rate is, unless the data provide evidence to the contrary, expected to be constant over time (at an unknown rate) but with the potential for small continuous variations that will show up as undulations in the fitted model. The locations and scale of these undulations will be learned from the observed data and, intuitively, will smooth observed fluctuations towards the estimated constant sedimentation rate. For a short summary of basic GP-model fitting and covariance expression, see Appendix A or Rasmussen and Williams (2006) if more information is required.

There is, however, no ideal age-depth model, and a GP does have potential weaknesses. Firstly, it may be the case that the tie-points are visually selected as obvious correlations between variations of climate proxies, and hence potentially correspond to extreme climate events. While the sedimentation rate is unknown, if its variation could be related to climate change, more significant sedimentation rate changes may occur at tie-points and a continuous (albeit slow) change in sedimentation may not be appropriate. In addition, use of a standard GP does not strictly enforce monotonicity in the resultant age-depth estimates. While this is clearly unrealistic in a sedimentation model, this potential lack of monotonicity is required for simple covariance information transferral through tie-pointing. Monotonicity can, however, be added back into any final age-depth model using an ad hoc Bayesian procedure whereby only the posterior distribution that satisfies the required ordering constraints is considered. In our experience of fitting such models to age-depth data, monotonicity is also rarely violated in practice.

4. METHODOLOGY

The methodology has 3 distinct steps:

1. Create an age-depth model for the reference record able to provide estimates at any chosen depth together with a measure of the uncertainty in these estimates;
2. Select tie-points and transfer age information over to the undated record;
3. Create an age-depth model for the undated record integrating out the uncertainty in the transferred ages created in step 1.

We explain each step in more detail below, although for a full, technical justification of the algorithm we refer the reader to Appendix B.

The first and third steps can either be fully automated or allow the user some input into the form of the age-depth model. The overall success of the method does, however, rely crucially on the second stage, which requires an external expert to identify the tie-points to use and provide potential information regarding their uncertainty as to the contemporaneity of the match. There will clearly be a degree of subjectivity in these choices. We look in Section 6 at the sensitivity of our method to both leaving out tie-points entirely and changing the certainty as to contemporaneity of the ties.

4.1 Creating an Age-Depth Model for the Reference Record

Inputs for this Stage

We commence by creating an age-depth model for the reference record as shown in Figure 3. For this, we assume that, within the reference record, we have a series of depths d_1, \dots, d_n with associated age estimates T_1, \dots, T_n arising from the model

$$T_i = \theta_R(d_i) + \varepsilon_i^R, \quad i = 1, \dots, n$$

Here, the ε_i^R can arise from any multivariate normal distribution and significantly, as we are using a GP, are allowed to be dependent, as might be the case if the estimates are found via varve counting or wiggle-matching. We do, however, require that the level of uncertainty and the dependence structure is known.




Figure 3 Creating an age-depth model for the reference record. The solid dots represent the age estimates observed within the reference record and associated 2σ error bars. The solid blue line shows the posterior mean for the GP-based age-depth model with the 2σ credible interval shown in red. At the selected tie-points p_1, \dots, p_n , we can use the model to estimate the corresponding ages $\theta_1, \dots, \theta_n$.

Outputs from this Stage

Having fitted a GP to the data according to the methodology in Appendix A, we can, given any depths \mathbf{p} within the reference record, create joint age estimates $\hat{\theta}_R(\mathbf{p})$ together with measures of uncertainty on these estimates and dependence encoded in the form of a covariance matrix $\Sigma_R(\mathbf{p})$.

4.2 Selecting Elastic Tie-Points and Transferring Information

Inputs for this Stage

Our next step relies upon the input of an expert to identify the points that will be used to tie the 2 records together. This may typically be done via the comparison of a continuous signal in both records thought to be affected by shared external events—e.g. $\delta^{18}\text{O}$ as a proxy for temperature. In this case, matchings may be chosen at the “obvious” correlation points believed to correspond to extreme and common events (global or at least regional), which should affect all systems simultaneously. It is possible there may be time lags between the response of the 2 systems to such events. We are able to incorporate these time lags and potential uncertainty in the synchronicity of the matchings if it can be quantified by the expert.

We require a set of matched depths p_i in the reference record and \tilde{p}_i in the undated record, which are believed to be tied via their respective ages according to

$$\theta_R(p_i) = \theta_N(\tilde{p}_i) + \alpha_i + \varepsilon_i^M \quad (1)$$

i.e. the expected age at depth p_i in the reference record is α_i years greater than the age at depth \tilde{p}_i in the undated record. Here, each ε_i^M can again be dependent, although, for simplicity, we assume they are independent $N(0, \sigma_i^2)$ with σ_i^2 representing the uncertainty in the expert's matching. Such a matching is illustrated in Figure 4 where the series of peaks observed in each record are tied together in order.




Figure 4 Selecting tie-points between records. Within the blue undated record, we identify a set of depths $\tilde{p}_1, \dots, \tilde{p}_k$, which can be matched to a corresponding set of depths within the red reference record p_1, \dots, p_k .

Outputs from this Stage

The aim of this stage is to use the age-depth model on the reference record to estimate the ages at the reference tie-points and then transfer this information over to the undated record.

4.3 Creating an Age-Depth Model for the New Record

Suppose hypothetically that we were to know exactly the true age of the tie-points in the reference record $\theta_R(\mathbf{p})$. According to Equation 1, these ages could simply be seen as noisy observations of the ages at depths $\tilde{\mathbf{p}}$ in the undated record since

$$y_i = \theta_R(p_i) - \alpha_i = \theta_N(\tilde{p}_i) + \varepsilon_i^M$$

where $\varepsilon_i^M \sim N(0, \varepsilon_i^2)$ denotes the matching uncertainty. We could then fit another GP to the (\tilde{p}_i, y_i) pairs to obtain an estimated chronology for the undated record and we would be done.

However, we do not know the true ages of the depths \mathbf{p} in the reference record but instead only have the uncertain estimates $\hat{\theta}_R(\mathbf{p})$ provided by the initial reference model. As a consequence, we need to first consider what are the possible “true” reference ages $\theta_R(\mathbf{p}) \sim N(\hat{\theta}_R(\mathbf{p}), \Sigma_R(\mathbf{p}))$ and for each find what individual age-depth model would be created for the undated record. We then have to average all of these separate models to integrate out the reference model uncertainty. Typically, this would be done by Monte Carlo integration as described in Section 3.3, making the method extremely slow. However, a GP permits it to be done exactly and immediately; details are given in Appendix B. We are then left with a complete chronology for the undated record.

Outputs from this Stage

Given any set of depths, we are able not only to give a point estimate for the corresponding ages, but also provide full covariance information specifying the dependence. The final chronology obtained in this way for our illustrative example is shown in Figure 5. As explained in Section 3.3, the joint estimates at a given set of points are not strictly enforced to be monotonic in the above procedure. We can, however, create such monotonicity by the ad hoc application of a Bayesian procedure using a monotonic prior. This corresponds to truncating the joint distribution for the depths of interest to that region that does satisfy the required monotonic constraints. Hence, to obtain a sample of possible ages, we sample from the full joint distribution and simply reject those samples that are not monotonic.




Figure 5 Creating an age-depth model for the undated record. In black, we show the posterior mean for the final chronology of the undated record. In blue, we show the 95% credible interval if covariance information from the reference age-depth model is passed through to the undated record. In red, we show the credible interval if it is not. The posterior mean will not be affected by the passing of covariance.

5. TIE-POINTING CONSIDERATIONS

5.1 Selecting Tie-Points

Throughout this article, we have presupposed that identification of the points tying the reference and undated record has already taken place. This may have given the impression that such identification is a simple step in the process with which we do not need to be concerned. In reality, such tie-pointing is a much more complex and difficult problem. It is also crucially important in determining the quality of the final chronology. We do not wish to give a fixed set of rules as to how such tie-pointing should be performed since we believe it should be problem-specific, but, in this section, we discuss the issues involved and describe the experiences and recommendations of the authors. Our advice relates to the matchings created for the Pakistan and Iberian margins (Bard et al. 2013, this issue)

and the update to the Cariaco Basin of Hughen et al. (2006). Here, paleoclimatic proxies were used to select tie-points, aiming to identify shared climatic changes through their effects on these signals. We talk predominantly in this paleoclimatic context but believe our suggestions are transferable to other frameworks.

5.1.1 Difficulties in Choosing Ties

The stylized illustrations in this article reduced the selection of common tie-points to the matching of unequivocal peaks separated by a complete absence of other signal. In practice, however, there will be several issues that complicate identification. Firstly, in both records, the signal used for matching is likely to contain polluting artifacts that relate to uncorrelated structures that are not shared between the records. In our paleoclimatic examples, these were linked to genuine but local climatic changes, additional effects on paleoclimatic proxies, and traditional noise in the proxy measurements themselves.

A second complicating issue is that the amplitude of signal changes used for matching may not be equivalent in the reference and undated records, e.g. the relative amplitude of temperature records may be different in different locations. In fact, in some cases, the change may not even be homothetic. Further difficulties can arise if considering proxies that are related in a nonlinear way to the underlying parameter of interest—while $\delta^{18}\text{O}$ of Greenland ice is monotonic, it is not a linear function of temperature. This is compounded if using proxies that represent different (though related) parameters in the climate system such as the comparisons of the Cariaco grayscale record with Greenland $\delta^{18}\text{O}$ used by Hughen et al. (2004). The first is a proxy for tradewind-induced coastal upwelling off Venezuela, while the second is, as stated above, a nonlinear function of Greenland temperature.

5.1.2 Why Tie-Point Rather Than Tune the Entire Record Automatically?

Several attempts have been made to automate the tuning process (e.g. Lisiecki and Lisiecki 2002). Such methods usually work by deforming the entirety of one record on to another with the optimal deformation selected via maximization of an overall correlation-type coefficient. While such methods could be considered more objective, in many records maximizing via a simple correlation coefficient is not appropriate. In addition, these methods do not appear to quantify the uncertainties involved in transferral.

While, in records of the same proxy from similar locations, a straightforward measure of correlation may be appropriate, in many instances one may be matching using diverse paleoclimatic proxies. These are likely to have the difficulties discussed above including a lack of linearity in response and use of different proxies in each record. In such cases, it would be an error to match all signal features via a standard correlation measure; we would instead need to develop a proxy-specific measure that could quantify how the response to the underlying parameter of interest in the selected proxy in one record would be mirrored in the proxy of the other. Furthermore, the large number of polluting and unrelated artifacts mean that often we should determine only those that relate to shared events and discard the rest. Such selection is difficult to automate, especially since the polluting artifacts may create effects on our signal as large as those that relate to the common events.

For these reasons, we believe it is often preferable to identify tie-points marking sharp transitions (assumed to be synchronous) rather than trying to maximize an overall correlation coefficient. When possible, it is also useful to consider multiple proxies in the same undated archives to check validity (see e.g. Bard et al. 2013, this issue).

5.1.3 Presenting Ties

Since tie selection has potential for subjectivity, the choice made should be clearly presented and justified in any transfer using the method presented in this article. This should be done in such a way that replication is possible. Tie-point selection is crucial to the accuracy of the final chronology; it is therefore important that the proxy records (both reference and undated) used for their identification are provided for the reader so that a judgment on the quality of the tuning can be made and to allow refitting should they wish to improve or change the selection of ties.

5.2 A Practical Example—Transferring a Timescale from Hulu Cave to the Arabian Sea

To provide a concrete example of tie selection, we present the approach taken in Bard et al. (2013, this issue) for their Pakistan Margin sediments. Here, the reference record was the high-resolution $\delta^{18}\text{O}$ record of the Hulu Cave in China dated by numerous and precise U-Th ages first published by Wang et al. (2001) with an augmented resolution provided by Edwards et al. (submitted). It is a record of the Asian monsoon that is clearly linked to the Heinrich and Dansgaard-Oeschger (H-DO) events observed in the North Atlantic. This Hulu $\delta^{18}\text{O}$ is a monotonic function of the monsoon rainfall intensity, but the exact relationship may be nonlinear and could even be variable through time.

To identify the ties in the Arabian Sea, we identified proxies linked in a semi-quantitative way to monsoon intensity. There is a full literature on this subject showing that proxies for marine productivity (e.g. organic carbon) and its consequences at mid-depth (e.g. denitrification and acidification) respond clearly to H-DO events. See Schulz et al. (1998) and Schulte et al. (1999) for details on organic carbon; Altabet et al. (2002) and Pichevin et al. (2007) for denitrification; and Böning and Bard (2009) for mid-depth acidification.

In the Pakistan Margin deep-sea core MD042876, we first measured multiple such proxies including abundance of organic carbon and specific molecules; nitrogen and its isotopes ($\delta^{15}\text{N}$); major, minor, and trace elements; and the mineralogy of carbonates (calcite and aragonite). All these records exhibited very clear signatures of the H-DO events (Pichevin et al. 2007; Böning and Bard 2009).

To tie this stratigraphy to the Hulu $\delta^{18}\text{O}$ record, we decided to use the % of total organic carbon as the “master” undated record. Having created a set of tie-points with this proxy, we confirmed our matches by cross-checking with both the nitrogen isotopes ($\delta^{15}\text{N}$) and the CaCO_3 record (measured at a higher resolution: every 5 mm over the 25-m core). The chosen ties were found to agree extremely well across the various proxies within the undated record (Bard et al. 2013, this issue).

5.2.1 Creating the Pakistan Margin, Iberian Margin, and Cariaco Chronologies Using OxCal-Dated Tie-Points

This article has recommended using GPs to create the age-depth models for both the reference record and the undated record for the reasons given in Section 3.3. However, this was not quite the approach taken in creation of the Pakistan and Iberian margin (Bard et al. 2013, this issue) and updated Cariaco (Hughen et al. 2006) chronologies. For these records, the Hulu Cave $\delta^{18}\text{O}$ reference record already had a master chronology created using high-resolution U-Th dates and the OxCal model (Edwards et al., submitted). It was felt that, rather than recreate a new GP-based Hulu timescale, this OxCal-created chronology should be kept as the single age-depth model for the Hulu reference to maintain consistency and comparability between different projects. Consequently, the independent age estimates of the tie-points in Hulu available from OxCal were modeled as independent normals during transferral. A GP was then used for creation of the final chronology in the

undated records based upon these age estimates. This still allowed the uncertainty in the tie-point ages to be integrated out exactly and Monte Carlo integration was not required.

5.3 Other Tie-Pointing Issues

Our process of elastic tie-pointing also presents other issues worthy of discussion, discussed below.

Implicit Incorporation of Prior Beliefs

When selecting ties, it is likely that the expert is implicitly using some prior beliefs about the nature of the sedimentation process to direct their search. For example, if sedimentation in a record is believed to be approximately 1 m per kyr, they may look in the neighborhood around 13 m for an event corresponding to the end of the Younger Dryas. This could in some informal sense be seen as a Bayesian-type prior on the underlying chronology. Evidence for such a prior would typically be strong (one is likely to have considerable external information on sedimentation rate). In applying our method, one would then view the transferral as modifying this prior in light of the observed locations of the events and the assumption of a smooth chronology.

Combining External Dating Information and Tie-Pointing

It is possible that the undated record may have some direct dating information. For example, it may be known that one end arises from the present day. Such mixed information can be incorporated into our approach by treating the directly dated depths as tie-points independent from those in the reference record. Similarly, in principle, one could transfer dating information from 2 distinct reference records by treating the ties from each reference record independently.

Modeling Sedimentation Rate

Our GP models the relationship between age and depth directly rather than considering sedimentation unlike OxCal, Bacon, and BChron. However, the 2 domains are equivalent and one can transfer from one to another. It is not clear that working with sedimentation offers particular benefits over working directly with the chronology and some measure of smoothness. Formally, however, as sedimentation is found by differentiation of the age-depth model, in our approach it will implicitly be modeled by another Gaussian process with a constant mean. If desired, one could infer the estimated sedimentation rate from our output.

Mixing Tie-Points

It is possible that during tie-pointing, a user may wrongly associate a point corresponding to event A in one record to a point corresponding to event B in the other. Such a mistake would obviously lead to an erroneous final chronology, especially as such an error may propagate through other ties. If it is unclear as to which tie-points should be matched to one another (i.e. depth p_i in the reference could match to either depth \tilde{p}_i^* or \tilde{p}_i^\dagger in the undated record), one could create the transferred chronologies for each possibility and then weight the 2 outputs according to the initial belief in the likelihood of each match. However, we would suggest that any such unclear tie not be considered for use in transferral.

6. ROBUSTNESS TO TIE-POINT SELECTION

6.1 Testing Robustness

Tie-point selection is critical to determining the success of our method and, as explained in Section 5, may be complex. We would ideally like a method of chronology transferral that is rela-

tively robust to small changes in these tie-points and mistakes that could be made. There are 2 main errors that could occur in tie-point selection: a false tie could be identified (or conversely a true tie could be missed); alternatively, the level of certainty in the contemporaneity of a tie could be misspecified. We investigate the effect of both these potential errors on the Iberian Margin chronology of Bard et al. (2013, this issue).

6.2 Changing Matching Certainty

Given a set of ties at depths p_i and \tilde{p}_i in the reference and undated record, respectively, we assume they correspond approximately to contemporaneous events in the 2 records according to

$$\theta_R(p_i) = \theta_N(\tilde{p}_i) + \varepsilon_i^M$$

where $\theta_R(p_i)$ is the age at depth p_i in the reference record and $\theta_N(\tilde{p}_i)$ is the age at depth \tilde{p}_i in the undated record. The $\varepsilon_i^M \sim N(0, \sigma_i^2)$ denotes the uncertainty we have in the contemporaneity in the 2 recorded events with σ_i specified by the expert providing the data—for example, that they believe the event recorded at p_i in the reference occurred within $2\sigma_i$ yr of the event at \tilde{p}_i in the undated record.

For the ties identified in the Pakistan and Iberian margins (Bard et al. 2013, this issue) and the update to the Cariaco chronology of Hughen et al. (2006), a value of $\sigma_i = 100$ yr was chosen for each point based upon the judgments of the experts involved. Here, we explore the effect of being less and more sure of tie-point contemporaneity by finding the resultant chronology for 3 possible σ_i values:

- Small: $\sigma = 20$ yr - a high level of certainty on contemporaneity suggesting a strong belief the identified events occurred at almost exactly the same time.
- Medium: $\sigma = 100$ yr - the level chosen for creation of the Pakistan and Iberian margins (Bard et al. 2013, this issue) and the updated Cariaco chronology (Hughen et al. 2006).
- Large: $\sigma = 500$ yr - a lower level of certainty on contemporaneity corresponding to a belief that the events could occur at more distant times.

The results are shown in Figure 6. There is little difference seen between $\sigma = 20$ and 100 yr on the chronology created; both are relatively wiggly and pass close to the tie-point data. Conversely, with $\sigma = 500$ yr, the chronology becomes more linear and does not pass as closely to the data. This is perhaps to be expected since as uncertainty in contemporaneity increases, the transferred tie-point age estimates provide less information.

More interestingly, we also observe that increasing the uncertainty in contemporaneity appears to lead to more certain final chronologies at depths away from the tie-points. This can be seen in the inset plot of Figure 6, where the black ($\sigma = 20$) dotted credible intervals are much wider than the red ($\sigma = 500$). This is perhaps unexpected. For a formal explanation of this phenomenon, one needs to consider the maximum likelihood estimates found for the Gaussian process in the undated record described in Appendix B.4. With $\sigma = 500$, the maximum likelihood value determined for λ is much increased, meaning that the chronology can only vary significantly from a linear model over much longer depth intervals than when $\sigma = 20$ or 100. However, an intuitive justification is perhaps more beneficial.

6.2.1 Intuitive Explanation of Narrower Credible Intervals with Larger Matching Uncertainty

Before we start fitting our model, it is not known how wiggly our chronology is going to be; instead, we aim to infer that from the data itself. If matching uncertainties are small, then the chronology will need to bend more in order to pass sufficiently close to these points. However, if we believe there to




Figure 6 The effect on the Iberian Margin chronology (Bard et al. 2013, this issue) of altering the level of certainty in the contemporaneity of the tie-points. The main plot shows the chronologies over the entire record, while the inset zooms in on a smaller section to show more detail. Three levels of certainty were chosen: high with $\sigma = 20$ yr (shown in black); medium $\sigma = 100$ yr (blue); and low $\sigma = 500$ yr (red). The solid dots represent the Iberian record's tie-points and their associated age estimates from the Hulu Cave chronology (Edwards et al., submitted); the accompanying solid bars show ± 100 yr corresponding to the medium uncertainty level. The dotted lines represent the 95% posterior credible interval for each chronology.

be large matching uncertainties, we could plausibly create a more linear chronology since it is not required to go as near to the individual data points.

In practice, the wiggleness of a chronology is a result of implicit changes in the sedimentation rate. However, our GP approach does not explicitly model sedimentation and does not restrict such changes to the tie-points. Instead, the underlying model is constructed so that if changes to sedimentation rate are seen to occur at any individual depth, then they could plausibly also be thought to occur at any other depth. This could be seen as a potential weakness in the method if it is believed that the identified tie-points correspond to climatic changes that are likely to be the sole/major changes in sedimentation.

As a consequence, if the data suggest a wiggly chronology (due to small matching uncertainties), then, intuitively, the method creates a model where the sedimentation rate can change more. It is not able to restrict these changes to lie only at tie-points but instead allows them to occur between the tie-points too. Hence, we create a final chronology that is less certain as to the ages far from the selected tie-points. Alternatively, if the tie-pointing supports a more linear chronology (as obtained with large matching uncertainties, or small matching uncertainties if the data are linear), then the method infers there are no sedimentation rate changes, and hence models it as close to constant over the entire time period. This gives more certainty on the chronology away from the tie-points.

6.3 Removing a Tie-Point

We also investigated the effect of leaving out tie-points entirely, to mimic the impact of the expert missing a tie-point or creating a spurious one. In such instances, one cannot expect the chronology in the region of the missing tie to give the same age estimates as those that would have been obtained with the true tie, or vice versa in the case of a spurious tie. However, we would hope that outside this local region the nature of the chronology would not be significantly changed.

In Figure 7, we show the results of one such experiment where a tie towards the beginning of the record was left out. The black curves (mean and 2σ credible intervals) represent the original chronology obtained using all of the ties, while in blue we show the chronology obtained if the tie in red is left out. As expected, the chronology in the region of the missing tie is different. This tie indicates a level of local variation that is not strongly supported by the rest of the data. However, the curve in the neighborhood of the unaffected tie-points show no significant differences. We also considered leaving out other points and multiple points at once. Similar results were seen where the chronology was altered in the region of the missing ties but little difference observed outside these regions indicating our method to be relatively robust to such errors.




Figure 7 The effect of failing to select a tie-point on the Iberian Margin chronology. In black, we present the Iberian chronology of Bard et al. (2013, this issue) formed using the tie-points identified by the solid dots and age estimates transferred from Hulu Cave (Edwards et al., submitted). In blue, we show the chronology that would have been created if we had failed to include the second tie-point (shown in red) in the transferral. The inset plot is a zoomed version of the chronology showing the section surrounding the missed tie.

CONCLUSION

In this article, we have presented a novel method for the transfer of timescales from a reference record to an undated record when the 2 records have been matched via a series of externally identified tie-points. Intuitively, we wish to elastically stretch or squash the reference chronology onto the

new record using the selected ties. Our approach consists of a simple 3-step process that is able to incorporate the distinct uncertainties and dependencies created at each stage and justify their accumulation from one step to the next. This will result in higher-quality final chronologies.

In principle, many age-depth models could be utilized in the process, although we select a GP to obtain computational feasibility while still retaining straightforward modeling of the dependence structure. Such dependence is inherent in all chronologies and, as we have shown, its use can improve the quality and usefulness of the resultant estimates. The resultant chronology will not only consist of point estimates but provides joint information for a set of depths including their shared uncertainties.

As shown in Figure 5, the uncertainty in our transferred age-depth models will tend to increase smoothly away from our selected tie-points. However, the age estimate for a particular depth will not depend solely upon the uncertainty of the nearest tie-point. Rather, we are able to borrow strength from the entire set of data points (ties and the original reference data) since we believe there to be a smooth age-depth curve underlying both of the records. As a consequence, although the original observed depths in the reference record or the tie-point matchings may be relatively uncertain, the final age-depth model itself can be estimated more precisely by combining the information from all the data together. The consequence of this will be that, even at depths distant from the tie-points, useful age estimates will still be possible.

There is potential for much further work in this field. Our current approach relies on the input of an expert to identify the set of tie-points, perhaps via their skill at identifying contemporaneous events in the 2 records through the correlations of external proxy signals reflecting shared external events. One potential improvement could be to automate this tie selection process. Further, for some records it may be sensible to try to match the entirety of these signals rather than just via a small discrete set of ties. Initial attempts at such methods have been attempted by Martinson et al. (1982), Lisiecki and Lisiecki (2002), and Haam and Huybers (2010), albeit without quantification of the inherent uncertainties. Success for any automated method would require incorporation of knowledge on the particular paleo-proxies used for matching. Their often complex and differing responses to the climate system mean that a straightforward linear correlation would typically be inappropriate to judge the quality of match and instead require use of a proxy-dependent measure. Determination of such measures would be challenging.

Related to this, development of a method to automate multiproxy tuning would be a further worthwhile direction of research. Typically, multiple signals in each record are available for matching and one would wish to include information from all of these simultaneously in determining any matching. Again, this would require selection of a suitable measure of match across all of these proxies combined. Such measures would require recognition that some proxies may track the underlying climate system more reliably than others, some proxies may miss certain events, and that 2 proxies within any individual record would also be correlated. Incorporating such scientific knowledge would strongly strengthen any chronology transferral method.

Finally, we should state that, even though we have tried in this article to quantify potential errors throughout chronology transferral, any chronology transferral will always contain a degree of subjectivity. If it is performed by individuals who are unaware of the complex functioning of both the climate system and paleo-proxies, then there is the potential for mistakes to be made. Examples where such errors have occurred can be found in Blaauw (2012). We therefore conclude by reiterating our recommendation that any transferral should be fully justified and allow reproduction of results by others.

ACKNOWLEDGMENTS

The authors would like to thank Dr J E Oakley for discussions about how uncertainty should be transferred between the steps of the algorithm; Prof C E Buck, along with 2 anonymous referees, for helpful comments that have considerably improved this paper; finally, Profs P J Reimer and P G Blackwell without whom this work would not have been conceived.

REFERENCES

- Altabet MA, Higginson MJ, Murray DW. 2002. The effect of millennial-scale changes in Arabian Sea denitrification on atmospheric CO₂. *Nature* 415(6868): 159–62.
- Bard E, Ménot G, Rostek F, Licari L, Böning P, Edwards RL, Cheng H, Wang YJ, Heaton TJ. 2013. Radiocarbon calibration/comparison records based on marine sediments from the Pakistan and Iberian margins. *Radiocarbon* 55(4), this issue.
- Blaauw M. 2012. Out of tune: the dangers of aligning proxy archives. *Quaternary Science Reviews* 36:38–49.
- Blaauw M, Christen JA. 2011. Flexible paleoclimate age-depth models using an autoregressive gamma process. *Bayesian Analysis* 6(3):457–74.
- Böning P, Bard E. 2009. Millennial/centennial-scale thermocline ventilation changes in the Indian Ocean as reflected by aragonite preservation and geochemical variations in Arabian Sea sediments. *Geochimica et Cosmochimica Acta* 73(22):6771–88.
- Bronk Ramsey C. 2008. Deposition models for chronological records. *Quaternary Science Reviews* 27(1–2): 42–60.
- Edwards RL, Cheng H, Wang YJ, Yuan DX, Kelly MJ, Kong XG, Wang XF, Burnett A, Smith E. 2013. A refined Hulu and Dongge Cave climate record and the timing of the climate change during the last glacial cycle. *Earth and Planetary Science Letters*, submitted.
- Haam E, Huybers P. 2010. A test for the presence of covariance between time-uncertain series of data with application to the Dongge Cave speleothem and atmospheric radiocarbon records. *Paleoceanography* 25(2): PA2209, doi:10.1029/2008PA001713.
- Haslett J, Parnell AC. 2008. A simple monotone process with application to radiocarbon-dated depth chronologies. *Journal of the Royal Statistical Society, Series C (Applied Statistics)* 57(4):399–418.
- Hogg AG, Hua Q, Blackwell PG, Niu M, Buck CE, Guilderson TP, Heaton TJ, Palmer JG, Reimer PJ, Reimer RW, Turney CSM, Zimmerman SRH. 2013. SHCal13 Southern Hemisphere calibration, 0–50,000 years cal BP. *Radiocarbon* 55(4), this issue.
- Hughen K, Lehman S, Southon J, Overpeck J, Marchal O, Herring C, Turnbull J. 2004. ¹⁴C activity and global carbon cycle changes over the past 50,000 years. *Science* 303(5655):202–7.
- Hughen K, Southon J, Lehman S, Bertrand C, Turnbull J. 2006. Marine-derived ¹⁴C calibration and activity record for the past 50,000 years updated from the Cariaco Basin. *Quaternary Science Reviews* 25(23–24): 3216–27.
- Lisiecki LE, Lisiecki PA. 2002. Application of dynamic programming to the correlation of paleoclimate records. *Paleoceanography* 17(4):1049–60.
- Martinson DG, Menke W, Stoffa P. 1982. An inverse approach to signal correlation. *Journal of Geophysical Research* 87(B6):4807–18.
- Parnell AC, Buck CE, Doan TK. 2011. A review of statistical chronology models for high-resolution, proxy-based Holocene palaeoenvironmental reconstruction. *Quaternary Science Reviews* 30(21–22):2948–60.
- Pichevin L, Bard E, Martinez P, Billy I. 2007. Evidence of ventilation changes in the Arabian Sea during the Late Quaternary: implication for denitrification and nitrous oxide emission. *Global Biogeochemical Cycles* 21(4): GB4008, doi:10.1029/2006GB002852.
- R Core Team. 2013. R: A Language and Environment for Statistical Computing. Vienna: R Foundation for Statistical Computing. URL: <http://www.R-project.org>.
- Rasmussen CE, Williams CKI. 2006. *Gaussian Processes for Machine Learning*. Cambridge: MIT Press.
- Reimer PJ, Baillie MGL, Bard E, Bayliss A, Beck JW, Blackwell PG, Bronk Ramsey C, Buck CE, Burr GS, Edwards RL, Friedrich M, Grootes PM, Guilderson TP, Hajdas I, Heaton TJ, Hogg AG, Hughen KA, Kaiser KF, Kromer B, McCormac FG, Manning SW, Reimer RW, Richards DA, Southon JR, Talamo S, Turney CSM, van der Plicht J, Weyhenmeyer CE. 2009. IntCal09 and Marine09 radiocarbon age calibration curves, 0–50,000 years cal BP. *Radiocarbon* 51(4): 1111–50.
- Reimer PJ, Bard E, Bayliss A, Beck JW, Blackwell PG, Bronk Ramsey C, Buck CE, Cheng H, Edwards RL, Friedrich M, Grootes PM, Guilderson TP, Hafliadason H, Hajdas I, Hatté C, Heaton TJ, Hoffman DL, Hogg AG, Hughen KA, Kaiser KF, Kromer B, Manning SW, Niu M, Reimer RW, Richards DA, Scott EM, Southon JR, Staff RA, Turney CSM, van der Plicht J. 2013. IntCal13 and Marine13 radiocarbon age calibration curves 0–50,000 years cal BP. *Radiocarbon* 55(4), this issue.
- Schulte S, Rostek F, Bard E, Rullkötter J, Marchal O. 1999. Variations of oxygen-minimum and primary productivity recorded in sediments of the Arabian Sea. *Earth and Planetary Science Letters* 173(3):205–21.
- Schulz H, von Rad U, Erlenkeuser H. 1998. Correlation

between Arabian Sea and Greenland climate oscillations of the past 110,000 years. *Nature* 393(6680):54–7.

Wang YJ, Cheng H, Edwards RL, An ZS, Wu JY, Shen

C-C, Dorale JA. 2001. A high-resolution absolute-dated Late Pleistocene monsoon record from Hulu Cave, China. *Science* 294(5550):2345–8.

APPENDIX A: A GAUSSIAN PROCESS

A.1 Definition of a Gaussian Process

As described in Rasmussen and Williams (2006), a (one-dimensional input) zero-mean *Gaussian process* $f(z) \sim \mathcal{GP}(0, k(z, z'))$ is a collection of random variables, any finite number of which have a joint Gaussian distribution. It is completely specified by its covariance function

$$k(z, z') = \mathbf{E}[f(z)f(z')]$$

Here, the random variables represent the values of the function $f(z)$ at our location z . Hence, if we have a set of N locations $\mathbf{z}_\star = (z_1^\star, z_2^\star, \dots, z_N^\star)^T$ where we are interested in the value of $f(\cdot)$, we find that

$$\mathbf{f}_\star = \mathbf{f}(\mathbf{z}_\star) \sim \mathcal{N}(0, K_{\mathbf{z}_\star, \mathbf{z}_\star})$$

where $K_{\mathbf{z}_\star, \mathbf{z}_\star}$ denotes the matrix of the covariances evaluated at all pairs of the locations z_i^\star .

Typically, we are not, however, interested in just drawing values from the prior but will have observed data that provide more information about the function. Specifically, suppose that we observe \mathbf{y} , the function $f(z)$ at a set of M locations $\mathbf{z} = (z_1, z_2, \dots, z_M)^T$ subject to noise according to the model

$$\mathbf{y} = \mathbf{f}(\mathbf{z}) + \boldsymbol{\epsilon}$$

where the noise $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \Sigma)$ may have dependence encoded in the covariance matrix Σ .

We can then write the joint distribution of the observed \mathbf{y} and the function $f(\cdot)$ at our set of points of interest \mathbf{z}_\star as

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{f}_\star \end{bmatrix} = \mathcal{N}\left(0, \begin{bmatrix} K_{\mathbf{z}, \mathbf{z}} + \Sigma & K_{\mathbf{z}, \mathbf{z}_\star} \\ K_{\mathbf{z}_\star, \mathbf{z}} & K_{\mathbf{z}_\star, \mathbf{z}_\star} \end{bmatrix}\right)$$

After some simple algebra shown in Rasmussen and Williams (2006), we can then derive the posterior distribution for our function values at our points of interest \mathbf{z}_\star

$$\mathbf{f}_\star | \mathbf{z}, \mathbf{y}, \mathbf{z}_\star \sim \mathcal{N}(\bar{\mathbf{f}}_\star, \text{cov}(\mathbf{f}_\star))$$

where

$$\bar{\mathbf{f}}_\star = K_{\mathbf{z}, \mathbf{z}_\star}^T [K_{\mathbf{z}, \mathbf{z}} + \Sigma]^{-1} \mathbf{y},$$

$$\text{cov}(\mathbf{f}_\star) = K_{\mathbf{z}_\star, \mathbf{z}_\star} - K_{\mathbf{z}, \mathbf{z}_\star}^T [K_{\mathbf{z}, \mathbf{z}} + \Sigma]^{-1} K_{\mathbf{z}, \mathbf{z}_\star}$$

A.2 Including a Mean Function

It is also possible to include a mean function into our GP according to a set of basis functions with coefficients β to be inferred from the data. Specifically, consider

$$g(z) = f(z) + \mathbf{h}(z)^T \beta, \quad \text{where } f(z) \sim \mathcal{GP}(0, k(z, z'))$$

where $\mathbf{h}(z)$ are a set of fixed basis functions (in our case those representing a straight line, i.e. $\mathbf{h}(z) = (1, z)^T$) and β are additional parameters to be estimated (in our case, the intercept and slope of the fitted line). If we then place a non-informative prior on our parameters β , we find a predictive distribution for

$$\mathbf{g}(\mathbf{z}_\star) | \mathbf{z}, \mathbf{y}, \mathbf{z}_\star \sim \mathcal{N}(\bar{\mathbf{g}}_\star, \text{cov}(\mathbf{g}_\star))$$

where

$$\bar{\mathbf{g}}_\star = \bar{\mathbf{f}}_\star + R^T \bar{\beta}$$

$$\text{cov}(\mathbf{g}_\star) = \text{cov}(\mathbf{f}_\star) + R^T (H[K_{\mathbf{z}, \mathbf{z}} + \Sigma]^{-1} H^T)^{-1} R$$

Here,

$$\bar{\beta} = (H[K_{\mathbf{z}, \mathbf{z}} + \Sigma]^{-1} H^T)^{-1} H[K_{\mathbf{z}, \mathbf{z}} + \Sigma]^{-1} \mathbf{y}$$

$$R = H_\star - H[K_{\mathbf{z}, \mathbf{z}} + \Sigma]^{-1} K_{\mathbf{z}, \mathbf{z}_\star}$$

with the H matrix recording the $\mathbf{h}(z)$ vectors for all observed points \mathbf{z} (similarly H_\star for the points of interest \mathbf{z}_\star). More detail if required can again be found in Rasmussen and Williams (2006), although it should be noted they present a multidimensional input GP and hence use different notation—our z and \mathbf{z} change to \mathbf{x} and X , respectively.

APPENDIX B: OUR MODEL

B.1 Conditioning on the Age of the Tie-Points

We wish to find the posterior distribution for the ages $\theta_N(\mathbf{x})$ at the depths \mathbf{x} given the observed ages \mathbf{T} within the reference record at depths \mathbf{d} . To find this, we condition on the ages at the tie-points in the reference record $\theta_R(\mathbf{p})$. Then,

$$\pi(\theta_N | \mathbf{x}, \mathbf{T}, \mathbf{d}, \mathbf{p}, \tilde{\mathbf{p}}) = \int \pi(\theta_N | \theta_R(\mathbf{p}), \mathbf{x}, \mathbf{T}, \mathbf{d}, \mathbf{p}, \tilde{\mathbf{p}}) \pi(\theta_R(\mathbf{p}) | \mathbf{x}, \mathbf{T}, \mathbf{d}, \mathbf{p}, \tilde{\mathbf{p}}) d\theta_R(\mathbf{p}). \quad (\text{B1})$$

We then make the assumptions that

- The ages θ_N only depend upon \mathbf{T} through the chosen tie-points, i.e. $\theta_R(\mathbf{p})$ are sufficient for the GP model in the new record. This means we can simplify

$$\pi(\theta_N | \theta_R(\mathbf{p}), \mathbf{x}, \mathbf{T}, \mathbf{d}, \mathbf{p}, \tilde{\mathbf{p}}) = \pi(\theta_N | \theta_R(\mathbf{p}), \mathbf{x}, \tilde{\mathbf{p}})$$

with

$$\theta_R(p_i) = \theta_N(\tilde{p}_i) + \varepsilon_i \quad \text{for } i = 1, \dots, K.$$

- The undated tie-points $\tilde{\mathbf{p}}$ provide no information about the reference ages, so

$$\pi(\theta_{\mathbf{R}}(\mathbf{p})|\mathbf{x}, \mathbf{T}, \mathbf{d}, \mathbf{p}, \tilde{\mathbf{p}}) = \pi(\theta_{\mathbf{R}}(\mathbf{p})|\mathbf{p}, \mathbf{T}, \mathbf{d}).$$

B.2 Fitting the Reference Age-Depth Model

We can find $\pi(\theta_{\mathbf{R}}(\mathbf{p})|\mathbf{p}, \mathbf{T}, \mathbf{d})$ using the results given in Appendix A.2, where we find

$$\theta_{\mathbf{R}}(\mathbf{p})|\mathbf{p}, \mathbf{T}, \mathbf{d} \sim \mathcal{N}(A_{\mathbf{R}}\mathbf{T}, \Sigma_{\mathbf{R}})$$

where

$$A_{\mathbf{R}} = K_{\mathbf{d}, \mathbf{p}}^T [K_{\mathbf{d}, \mathbf{d}} + \Sigma_O]^{-1} + R_{\mathbf{R}}^T (H[K_{\mathbf{d}, \mathbf{d}} + \Sigma_O]^{-1} H^T)^{-1} H(K_{\mathbf{d}, \mathbf{d}} + \Sigma_O)^{-1},$$

$$\Sigma_{\mathbf{R}} = K_{\mathbf{p}, \mathbf{p}} - K_{\mathbf{d}, \mathbf{p}}^T [K_{\mathbf{d}, \mathbf{d}} + \Sigma_O]^{-1} K_{\mathbf{d}, \mathbf{p}} + R_{\mathbf{R}}^T (H[K_{\mathbf{d}, \mathbf{d}} + \Sigma_O]^{-1} H^T)^{-1} R_{\mathbf{R}}.$$

Here,

$$R_{\mathbf{R}} = H_{\mathbf{p}} - H_{\mathbf{d}} [K_{\mathbf{d}, \mathbf{d}} + \Sigma_O]^{-1} K_{\mathbf{d}, \mathbf{p}}$$

with the $H_{\mathbf{d}}$ matrix recording the $\mathbf{h}(\mathbf{d})$ vectors for all observed depths \mathbf{d} (similarly $H_{\mathbf{p}}$ for the tie-point depths \mathbf{p}). Here, K is the kernel function relating to the reference model and Σ_O is the uncertainty in the dated observations.

B.3 Fitting the New Age-Depth Model

Given tie-point ages $\theta_{\mathbf{R}}(\mathbf{p})$ in the reference record, we can also use Appendix A.2 to find

$$\theta_{\mathbf{N}}|\theta_{\mathbf{R}}(\mathbf{p}), \mathbf{x}, \tilde{\mathbf{p}} \sim \mathcal{N}(B_{\mathbf{N}}\theta_{\mathbf{R}}(\mathbf{p}), \Sigma_{\mathbf{N}}),$$

where

$$B_{\mathbf{N}} = C_{\tilde{\mathbf{p}}, \mathbf{x}}^T [C_{\tilde{\mathbf{p}}, \tilde{\mathbf{p}}} + \Sigma_M]^{-1} + R_{\mathbf{N}}^T (H[C_{\tilde{\mathbf{p}}, \tilde{\mathbf{p}}} + \Sigma_M]^{-1} H^T)^{-1} H(C_{\tilde{\mathbf{p}}, \tilde{\mathbf{p}}} + \Sigma_M)^{-1},$$

$$\Sigma_{\mathbf{N}} = C_{\mathbf{x}, \mathbf{x}} - C_{\tilde{\mathbf{p}}, \mathbf{x}}^T [C_{\tilde{\mathbf{p}}, \tilde{\mathbf{p}}} + \Sigma_M]^{-1} C_{\tilde{\mathbf{p}}, \mathbf{x}} + R_{\mathbf{N}}^T (H[C_{\tilde{\mathbf{p}}, \tilde{\mathbf{p}}} + \Sigma_M]^{-1} H^T)^{-1} R_{\mathbf{N}}.$$

Here,

$$R_{\mathbf{N}} = H_{\mathbf{x}} - H_{\tilde{\mathbf{p}}} [C_{\tilde{\mathbf{p}}, \tilde{\mathbf{p}}} + \Sigma_M]^{-1} C_{\tilde{\mathbf{p}}, \mathbf{x}}$$

with C the kernel for the GP relating to the undated model and Σ_M the uncertainty in the matching.

All that is now left is to integrate out $\theta_{\mathbf{R}}(\mathbf{p})$ and we are left with a predictive distribution for $\theta_{\mathbf{N}}(\mathbf{x})$, the ages of interest,

$$\theta_{\mathbf{N}}|\mathbf{x}, \mathbf{T}, \mathbf{d}, \mathbf{p}, \tilde{\mathbf{p}} \sim \mathcal{N}(B_{\mathbf{N}}A_{\mathbf{R}}\mathbf{T}, B_{\mathbf{N}}\Sigma_{\mathbf{R}}B_{\mathbf{N}}^T + \Sigma_{\mathbf{N}}).$$

B.4 Choosing the Kernel Function

In order to fit our model, we need to specify the kernel functions C and K for the age-depth GPs in the reference and undated record. For both, we select the Matern family with $\nu = 3/2$, so that the kernel function is

$$k(z, z') = \tau^2 \left(1 + \sqrt{3} \frac{|z - z'|}{l} \right) \exp \left(-\sqrt{3} \frac{|z - z'|}{l} \right)$$

This class of kernel function was chosen since this results in sample paths of the processes that we believed represented realistic potential age-depth models.

The parameters τ and l are selected separately for the reference and undated age-depth models via maximum likelihood. To determine these parameters in the case of the reference model, we use the observed data (\mathbf{T}, \mathbf{d}) , while for the undated record, for simplicity, we use the posterior mean age obtained from the reference model at the tie-points, i.e. $(E[\theta_{\mathbf{R}}(\mathbf{p})], \tilde{\mathbf{p}})$.