UC Santa Barbara

UC Santa Barbara Previously Published Works

Title

Probabilistic sequence alignment of stratigraphic records

Permalink

https://escholarship.org/uc/item/7m02079z

Journal

Paleoceanography and Paleoclimatology, 29(10)

ISSN

2572-4517

Authors

Lin, Luan
Khider, Deborah
Lisiecki, Lorraine E
et al.

Publication Date

2014-10-01

DOI

10.1002/2014pa002713

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

AGU PUBLICATIONS

Paleoceanography

RESEARCH ARTICLE
10.1002/2014PA002713

Key Points:
- Hidden Markov model estimates uncertainty in benthic oxygen isotope alignments
- Radiocarbon-based estimates of sedimentation rate variability
- Bayesian statistics: alignment confidence band, tests of lead-lag relationships

Supporting Information:
- Readme
- Table S1
- Table S2
- Table S3
- Text S1
- Text S2
- Text S3
- Figures S1–S7

Correspondence to:
L. Lin,
luan.lin@mssm.edu

Citation:
Lin, L., D. Khider, L. E. Lisiecki, and C. E. Lawrence (2014), Probabilistic sequence alignment of stratigraphic records, *Paleoceanography*, *29*, 976–989, doi:10.1002/2014PA002713.

Received 22 AUG 2014
Accepted 17 SEP 2014
Accepted article online 27 SEP 2014
Published online 27 OCT 2014

Probabilistic sequence alignment of stratigraphic records

Luan Lin[1], Deborah Khider[2], Lorraine E. Lisiecki[2], and Charles E. Lawrence[1]

[1]Division of Applied Mathematics, Brown University, Providence, Rhode Island, USA, [2]Department of Earth Science, University of California, Santa Barbara, California, USA

Abstract The assessment of age uncertainty in stratigraphically aligned records is a pressing need in paleoceanographic research. The alignment of ocean sediment cores is used to develop mutually consistent age models for climate proxies and is often based on the δ^{18}O of calcite from benthic foraminifera, which records a global ice volume and deep water temperature signal. To date, δ^{18}O alignment has been performed by manual, qualitative comparison or by deterministic algorithms. Here we present a hidden Markov model (HMM) probabilistic algorithm to find 95% confidence bands for δ^{18}O alignment. This model considers the probability of every possible alignment based on its fit to the δ^{18}O data and transition probabilities for sedimentation rate changes obtained from radiocarbon-based estimates for 37 cores. Uncertainty is assessed using a stochastic back trace recursion to sample alignments in exact proportion to their probability. We applied the algorithm to align 35 late Pleistocene records to a global benthic δ^{18}O stack and found that the mean width of 95% confidence intervals varies between 3 and 23 kyr depending on the resolution and noisiness of the record's δ^{18}O signal. Confidence bands within individual cores also vary greatly, ranging from ~0 to >40 kyr. These alignment uncertainty estimates will allow researchers to examine the robustness of their conclusions, including the statistical evaluation of lead-lag relationships between events observed in different cores.

1. Introduction

The stratigraphic alignment of ocean sediment cores plays a vital role in paleoceanographic research because it is used to develop mutually consistent age models for climate proxies measured in these cores. Although stratigraphic alignment algorithms have several different applications, here we focus on one of the most common, which is the creation of a core age model by aligning benthic δ^{18}O data from that core, measured as a function of depth below the sediment surface, to the benthic δ^{18}O data of another record for which age estimates are available. Benthic δ^{18}O is a measure of the oxygen isotope ratio in the calcium carbonate shells of benthic foraminifera. Because a large fraction of benthic δ^{18}O variance derives from the global signal of ice volume [*Adkins et al.*, 2002; *Elderfield et al.*, 2012], similar benthic δ^{18}O signals can be observed in nearly all marine cores [*Lisiecki and Raymo*, 2005].

The main purpose of benthic δ^{18}O alignment is often to develop age estimates for one core that are consistent with the age estimates of a second "target" record. These alignments can then be used to evaluate the relative leads and lags of climate records observed in different ocean sediment cores with, perhaps, greater confidence than is available for the absolute ages of these climate records [*Imbrie et al.*, 1992; *Lisiecki et al.*, 2008; *Lawrence et al.*, 2013]. The technique can be particularly powerful when many cores can be compared because all have been aligned to the same benthic δ^{18}O target, such as the LR04 benthic δ^{18}O stack (average) of 57 different globally distributed records. To date, δ^{18}O alignment has been performed either by manual, qualitative comparison [e.g., *Prell et al.*, 1986; *Shackleton et al.*, 2000] or by deterministic algorithms [e.g., *Martinson et al.*, 1987; *Lisiecki and Lisiecki*, 2002; *Huybers and Wunsch*, 2004].

Despite the widespread use of stratigraphic alignments, very little previous work has been done to evaluate the uncertainties associated with benthic δ^{18}O alignment. One study compared four different manual alignment strategies and found a standard deviation of 2500 year between the different alignment results [*Martinson et al.*, 1987]. Most commonly, only a single alignment is developed for a core and the uncertainty associated with the alignment is usually either not mentioned [e.g., *Hodell et al.*, 2003; *Raymo et al.*, 2004] or assumed to be approximately 4 kyr [*Huybers and Wunsch*, 2004; *Lisiecki et al.*, 2008]. For example, the deterministic "Match" algorithm of *Lisiecki and Lisiecki* [2002] outputs only a single parameter-dependent solution without any information about the uncertainty associated with this alignment. Additionally, the

©2014. American Geophysical Union. All Rights Reserved.

algorithm's penalty weightings, which help prevent overfitting the δ^{18}O signals by penalizing unrealistic changes in sedimentation rates, must be selected by trial and error, relying on user judgment to determine whether the resulting alignment is reasonable.

To address these limitations, we present a probabilistic stratigraphic alignment algorithm (HMM-Match), show how to derive Bayesian confidence bands from this algorithm, and illustrate its application. Whereas Match applies deterministic constraints on sedimentation rate changes based on subjective user-defined penalties, HMM-Match evaluates alignments probabilistically based on observations of sedimentation rate variability derived from independent radiocarbon data. This change improves alignment accuracy.

We begin by describing the probabilistic model and how it differs from the deterministic algorithm of *Lisiecki and Lisiecki* [2002] in section 2. In section 3 we analyze the algorithm's results for the alignment of 35 δ^{18}O records to the LR04 stack [*Lisiecki and Raymo*, 2005], including confidence intervals for δ^{18}O and age estimates as a function of time compared to the LR04 stack. Then, in section 4 we discuss in more detail the assumptions associated with this algorithm and its potential applications. For example, confidence bands returned by our new algorithm can help researchers identify possible errors in the predicted alignment (e.g., due to core disturbances) and test the robustness of lead-lag relationships with respect to alignment uncertainty. The supporting information include a full mathematical description of the algorithm as well as a less technical primer on the geophysical assumptions and statistical framework for HMM-Match. Web service and open sources software for HMM-Match is available at http://ccmbweb.ccv.brown.edu/cgi-bin/geo_align.pl.

2. Methods

2.1. Background

Because direct measurement of age in marine sediment records spanning 10^5–10^6 years is not possible, these ages must be inferred from other data, such as benthic δ^{18}O. Benthic δ^{18}O ratios corresponding to the same time should be similar because a large fraction of benthic δ^{18}O variance derives from the global signal of ice volume. However, depth in core does not necessarily correspond to this common time because sedimentation rates vary between cores and within individual cores. Therefore, we reconstruct sediment core ages by aligning the δ^{18}O measurements of two records to account for unobserved changes in sedimentation rate. Analogous alignment problems exist in many fields as unobserved events delete or insert symbols, sounds or measurements, or warp time through compression or expansion of a record.

The Match algorithm [*Lisiecki and Lisiecki*, 2002] is a well-accepted technique for stratigraphic alignments [e.g., *Channell et al*., 2009; *Lang and Wolff*, 2011] that draws a deterministic inference using dynamic programming optimization [*Bellman*, 1961]. Although dynamic programming has been used extensively to find alignments that maximize measures of sequence similarity [*Sankoff and Kruskal*, 1999; *Smith and Waterman*, 1981; *Lisiecki and Lisiecki*, 2002], it is not suitable for assessing the uncertainty of an alignment.

In contrast, probabilistic alignment algorithms provide a principled and direct means to address uncertainty in a sequence alignment and have been employed in many fields including speech recognition [*Levinson et al*., 1983], music classification [*Pardo et al*., 2004], and extensively in computational biology [*Durbin et al*., 1998; *Zhu et al*., 1998; *Webb et al*., 2002]. The central concept behind such a probabilistic model is that the inference of unknowns (here the age for each depth in the core) is inherently uncertain; thus, they return probability distributions for the unknowns rather than a single result. When only two individual sequences are available, pair probabilistic alignment algorithms are employed [*Webb et al*., 2002], with pair hidden Markov model (HMM) being the most well established of these methods [*Durbin et al*., 1998]. Here we develop a probabilistic model for pairwise stratigraphic alignment, called HMM-Match, by adapting the Match algorithm into the form of a pair hidden Markov model (HMM) [*Rabiner*, 1989; *Durbin et al*., 1998].

2.2. Method Overview

The goal of a stratigraphic alignment is to identify points in two records that were deposited at the same time by accounting for the differences in sedimentation rates in the two records. The basic strategy of our approach is to build a probabilistic model that first models the unobserved differences in sedimentation rates and then models the noise in the proxy data, δ^{18}O in this application. Then the calculus of probability is used to reverse the direction of logical inference to go from the observed proxy data to the unobserved alignment of the records. Because this model covers the whole process and not just the variation in the sampled

©2014. American Geophysical Union. All Rights Reserved.

data, it is a Bayesian model, and thus, all statistical inferences including confidence bands and tests that we report are Bayesian inferences. In Bayesian statistical terms, the model of the physical sedimentation processes is known as the prior model, and the model describing the noisy data is called the likelihood. Here we use this model to align individual input records to the LR04 stack, which is composed of δ^{18}O values with corresponding assigned ages [*Lisiecki and Raymo*, 2005]. The alignment is described by random variables, A_i, that specify the unknown age in the stack corresponding to the depth of the ith data point in the input for all depths in the record $i = 1, 2, \ldots, n$. Thus, we define an alignment $\mathbf{A} = (A_1, A_2, \ldots, A_n)$ as a vector of age assignments for every depth in the input core at which δ^{18}O was measured. We use probability theory in these assignments to account for the fact that the alignment process is uncertain. The mathematical nature of the model requires that we be fully explicit about its assumptions and means that our conclusions are a mathematical consequence of these assumptions and the data.

Bayes rule is employed to accomplish the reversal of logical inference. It combines the probability that the given δ^{18}O data would be observed for a particular alignment (i.e., the probability for a given residual misfit between the aligned record and the target/stack), with the probability that a given alignment would occur independent of any δ^{18}O data (i.e., the probability of a given sequence of sedimentation rate changes), the prior model. This is written as follows:

$$P(\text{alignment}^*|\text{data}, \phi, \theta) = \frac{P(\text{data}|\text{alignment}^*, \theta)P(\text{alignment}^*|\phi)}{\sum_{\text{alignment}} P(\text{data}|\text{alignment}, \theta)P(\text{alignment}|\phi)} \quad (1)$$

where ϕ are the parameters concerning the alignment and θ are the parameters used in the model of the noisy δ^{18}O data. The summation in the denominator of equation (1) is over all possible alignments, as described in the supporting information.

The notation $P(\text{data}|\text{alignment}^*, \theta)$ gives the probability of observing the input core's δ^{18}O given a particular alignment A. In HMM jargon, this term is called the emission model which we specifically call the "δ^{18}O residual model" in section 2.3. The second term in the numerator, $P(\text{alignment}^*|\phi)$ specifies the model for the probability for the alignment vector **A** given the parameters we use to constrain sedimentation rate change. This is a transition model, which we will refer to as the "sedimentation rate model" (sections 2.4 and 2.5). $P(\text{alignment}^*|\text{data}, \phi, \theta)$ gives the posterior probability of any given alignment. Because there are far too many alignments to calculate all their probabilities, we use a stochastic back trace algorithm to draw samples from the distribution of all alignments in exact proportion to their probability and use these to obtain alignment confidence bands (section 2.6). We refer the reader to Texts S1–S3 in the supporting information for the full mathematical model and the algorithms used to complete the sums and draw the samples.

2.3. Emission Model: δ^{18}O Residuals

The emission model accounts for the fact that even if we knew the correct alignment, the two sequences of δ^{18}O ratios cannot be expected to match exactly due to measurement error and spatial variability. Observations of δ^{18}O residuals between individual aligned records and the LR04 stack (Figure S1) indicate that the difference between the δ^{18}O values of the input core and the target at any given time follow a normal probability distribution. Additionally, the overall mean δ^{18}O proxy value in some input records is slightly shifted up or down from those in the stack due to spatial variability in the ocean as well as interlaboratory calibration offsets [*Ostermann and Curry*, 2000]. Thus, our model for δ^{18}O residuals has two parameters: σ^2 which is the variance of the δ^{18}O residual distribution and μ which accounts for potential shifts in the mean of input records. To account for differences in input records, we estimate σ^2 and μ for each input record using the Baum-Welch expectation maximization algorithm [*Rabiner*, 1989; *Durbin et al.*, 1998] (see Texts S1–S3 in the supporting information).

2.4. Point-Based Alignment

Although the original Match algorithm uses alignment intervals, we found that a point-based alignment strategy was more suitable for proper implementation of the δ^{18}O residual model in HMM-Match. Match's interval alignment strategy divides both the input record and the target stack into a user-defined number of same size depth (or age) intervals that contain varying numbers of data points and then aligns records by matching intervals using a finite number of user-defined ratios (typically ~17). However, because δ^{18}O records

©2014. American Geophysical Union. All Rights Reserved.

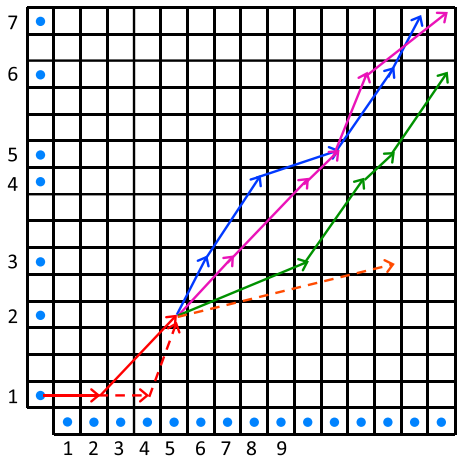

Figure 1. Example alignment. Arrows of different colors and line styles illustrate several possible ways to align an input δ^{18}O signal to the target stack. Grid squares represent intervals to be aligned, which are exactly 1 kyr in the target stack. Numbered dots on the left and bottom edge show the locations of individual δ^{18}O measurements, which are unevenly spaced in the input core.

often contain unevenly spaced data, some of the intervals used by Match may have no data points assigned to them. To address this missing data, Match uses linear interpolation to assign values to such intervals and thus creates data where there is none. While Match's interpolation strategy seems to have little impact on an optimal alignment algorithm, it leads to substantial error in the assignment of uncertainty to an alignment. Because after linear interpolation both low- and high-resolution records appear to have a full complement of data, the use of linear interpolation yields confidence limits for low-resolution records that are as tight as those for high-resolution records.

We addressed this shortcoming by developing a point-based alignment algorithm which abandons the use of intervals in the input records and, instead, aligns each individual data point in the input record to an interval in the target stack. Within this point-based approach, sedimentation rates are allowed to change between each depth with a δ^{18}O measurement using a large number of relative alignment ratios from 0.25 (1:4) to 4 (4:1). To make this more concrete, consider the specific example shown in Figure 1. We see that input data point 2 is three depth units deeper than data point 1. If input data point 1 is aligned to the stack at time 2 and data point 2 is aligned to the stack at time 5, then the alignment ratio is 1:1. If, however, depth point 1 was aligned to the stack at time 4, then alignment of data point 2 to stack at time 5 produces a ratio of 3:1. Note that during the preprocessing procedure, we normalized the input and target sequences to the same scale assuming that they are approximately end-to-end aligned. Changes in the input core's sedimentation rate are constrained using a model of the transition probabilities as described in the following section.

Note that although we align individual points from the input core, we must still define an interval size for the alignment target in order to create a finite number of possible alignments. The target interval size is selected to be approximately one fourth of the average resolution of the input record (e.g., 0.25 kyr for an input core with ~1 kyr resolution) to allow for the full range of possible compression ratios (Figure S3). The depths of all the input points $\mathbf{d} = (d_1, d_2 \ldots, d_N)$ have been measured, are taken as known, and may be unevenly spaced. Before alignment, HMM-Match converts each input depth to a depth unit index by dividing the input core into the same number of intervals as in the target. In Figure 1, the first three depth unit indices would be $d_1 = 1$, $d_2 = 4$, $d_3 = 6$. This conversion involves rounding depths by at most one eighth of their average sample spacing. For example, in a core with 4 cm sample spacing, the conversion from depth to depth unit index would shift the depths of observations by at most 0.5 cm.

2.5. Transition Model: Sedimentation Rates

The sedimentation rate model, $P(\text{alignment}^*|\phi)$, estimates the probability of a given alignment (independent of δ^{18}O data) based on the prior knowledge that changes in sedimentation rates are limited and that larger differences in sedimentation rates are generally less likely than smaller differences. The Match algorithm addresses this issue with user-defined parameters that penalize for extremely large/small sedimentation rates as well as rapid changes in sedimentation rate. In contrast, HMM-Match uses data-derived parameters to estimate the probability of a given sequence of sedimentation rate changes (see section 3). The sedimentation rate model, called a transition model, calculates the probability of transitions between different alignment ratios (~1:1, expansion, or contraction) between each data point in the input core.

Because the number of possible ways to align two records is immense, probabilities within the sedimentation rate model cannot be estimated by enumerating all possible alignments and separately calculating the probability of each. Fortunately, these probabilities can be found exactly without complete enumeration due to the property of conditional independence that is fundamental to Markov chains. A sequence follows a

©2014. American Geophysical Union. All Rights Reserved.

Markov chain if all the history leading to each point in the solution space of a problem can be summarized in a function at each point in the space. Match capitalizes on this property by using a dynamic programming algorithm to divide the alignment optimization problem into progressively smaller subproblems [*Bellman*, 1961]. HMMs capture this property in a similar manner by accumulating sums over probabilistically weighted subproblem histories instead of subproblem optimizations.

Thus, although the probability of a given age assignment A_i for the ith input data point does depend on all others, this dependence can be fully accounted for by the age assignments of A_{i-1} and A_{i-2} (see the supporting information for details), which provide information about whether point $i-1$ was aligned using a ratio of ~1:1, >1 (expansion), or <1 (contraction). Because these probabilistic relationships concern switches in state (i.e., changes in alignment ratios) from the $i-1$th point to the ith point, they are called transition probabilities. This Markov property is specified mathematically as follows:

$$P(A_i|A_1, A_2, \ldots, A_{i-1}) = P(A_i|A_{i-1}, A_{i-2}), i = 3, 4, \ldots, n$$

Thus, HMM-Match solves for all possible alignments by calculating the probability of every possible transition starting with the first two data points in the input core, progressively incrementing by 1 and ending with n (the last data point in the core). During this process, the partial sums for each step are stored in memory for later use in the stochastic back trace algorithm (see section 2.6).

An additional problem is how to align the start and end of each core. Although it is necessary to initially select a segment of the stack to which a record will be aligned, the HMM-Match algorithm includes components to evaluate different potential starting and ending for an alignment and estimates their associated uncertainties. Specifically, we adopt an approach commonly used for DNA sequence alignments [*Durbin et al.*, 1998] to account for the fact that DNA bases can be inserted or deleted through the evolution of a current-day sequence. Here HMM-Match assumes that the input record may be somewhat longer or shorter than the chosen segment of the alignment target by allowing a small number of input data points at the top or bottom of the input core to be seen as deletions or insertions rather than points being generated from the target stack.

2.6. Confidence Limits and Median Solution

To arrive at an alignment solution, both Match and HMM-Match employ two stages. In Match the first stage finds the minimum possible penalty score using an iterative process that goes step-by-step through the alignment matrix and stores in memory the partial minimum at each step leading to the global minimum for the full alignment. In an analogous first stage, HMM-Match goes step-by-step accumulating in memory partial sums until the whole alignment matrix is filled. The last entry gives the sum in the denominator of equation (1).

In the second stage, both Match and HMM-Match employ back trace algorithms that reverse direction of "travel" through the alignment matrix. Because it works from the oldest sediment back up to the youngest, the algorithm captures information from downstream (older) points. In addition, conditional independence allows each back trace step to capture the influence of all upstream (younger) points using the partial minima (in Match) or partial sums (in HMM-Match) stored in memory from the first stage. As a consequence, Match's deterministic back trace algorithm identifies the alignment that is guaranteed to minimize its penalty score. Analogously, HMM-Match, which employs a stochastic back trace algorithm, draws samples in exact proportion to their probability. These sampled alignments are independent of one another because each requires only results from the first stage and, thus, knows nothing about the previously sampled alignments. Therefore, this technique yields an independent representative sample from the space of all alignments. The stochastic back trace algorithm is also fast compared to the forward step, which allows it to efficiently draw a large number of samples.

In this application we use 1000 stochastic back trace samples to obtain alignment confidence limits and to generate a single "best" alignment. Each sample alignment aligns every point in the input to a point in the stack. Then, confidence limits are generated by identifying the range of stack ages that encompass 95% of the sampled alignments for each point (i.e., from the 25th youngest age to the 975th oldest). The tabulation of these confidence limits for all points in the input core yields an estimated confidence band for the full alignment. Because the ages assigned to points in the stack are taken as known, these confidence limits quantify alignment uncertainty but not uncertainty associated with age model of the alignment target.

A single best estimate can be obtained from this probability distribution using statistical decision theory that relies on the choice of a loss function. Match, like many other popular methods, finds the result that

©2014. American Geophysical Union. All Rights Reserved.

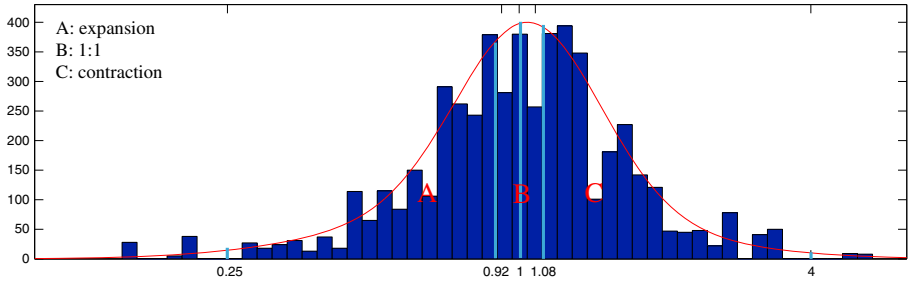

Figure 2. Histogram of sedimentation rate ratios calculated from 37 radiocarbon-dated cores (see Figure S4 and Table S1). The corresponding mixture log-Gaussian fit is plotted in red. These rates are partitioned into three groups: A = expansion, B = 1:1, and C = contraction.

optimizes a score. The most probable alignment of HMM-Match is its optimal-score alignment. However, since optimal-score alignments do not have to be near the center of a probability distribution, this solution (called a modal estimate) can occur at the edges of the probability distribution of the results [*Carvalho and Lawrence*, 2008] (e.g., Figure S2). In this study, we found that a modal estimate will sometimes fall outside of its own confidence bands (see Texts S1–S3 in the supporting information). To address this problem, we instead employed a median estimator using the theory developed by *Carvalho and Lawrence* [2008] as described in Texts S1–S3 in the supporting information. The median alignment is defined as the alignment that minimizes the expected sum of absolute differences from the space of all alignments and is an estimator of the central tendency of the posterior probability distribution of alignments (see the supporting information). This central tendency keeps the median alignment well inside the confidence limits.

3. Radiocarbon-Based Estimates of Marine Sedimentation Rates

In the Match algorithm, the cumulative alignment score includes sedimentation rate and rate change penalties that limit expansion and contraction of the aligned record. These two penalties are reflected in the transition parameters, ϕ, of the sedimentation rate model of HMM-Match. The transition parameters give the probability distribution of aligned ages for the next data point in the input given the distribution of aligned ages of the current data point in the stack. In order to assure that our model produces realistic alignments, we estimate these transition rates using sedimentation rate changes reconstructed from independent radiocarbon data (Figure S4). Specifically, we estimate the probabilities of transitions over ranges for expansion (0.25 to 0.9220), 1:1 (0.9220 to 1.0850), and compression (1.0850 to 4) using the frequency of these transitions in the radiocarbon-based estimates of sedimentation rate change. Probabilities corresponding to particular levels of expansion or contraction are set proportional to their observed frequency in radiocarbon-based estimates of sedimentation rate as approximated by a mixture of two log-Gaussian functions (Figures 2 and S7).

We estimated sedimentation rate variability using radiocarbon dates from planktonic foraminifera in 37 cores with mean sedimentation rates of at least 8 cm/kyr (Table S1 in the supporting information). These cores are globally distributed but exclude cores from the high-latitude (>40°N) North Atlantic where surface reservoir age variations may distort radiocarbon age estimates [*Stern and Lisiecki*, 2013]. Because 16 of these cores have age models extending to at least 40 ka, the distribution of observed sedimentation rates should include variability associated with glacial-interglacial climate change. We specifically analyzed only sedimentation rate estimates based on closely spaced ^{14}C measurements (0.5–4 kyr) so that our estimate of sedimentation rate variability is suitable for alignment of δ^{18}O data with similar temporal resolution. Importantly, our sedimentation rate estimates are independent of any assumptions about the time variation of benthic δ^{18}O because they are based only on radiocarbon data without using any δ^{18}O alignments.

A radiocarbon age model for each core (Figure S4) was developed using the Bayesian age modeling program Bchron [*Haslett and Parnell*, 2008; *Parnell et al.*, 2008] and the Marine09 calibration [*Reimer et al.*, 2009] with constant 405 ^{14}C year reservoir ages (as described in *Peterson et al.* [2014]). Bchron is a Bayesian algorithm that estimates age-depth functions consistent with a core's radiocarbon ages while taking into account the probability distributions of calibrated radiocarbon ages, excluding radiocarbon outliers and allowing for a parameterized number of sedimentation rate changes throughout the core [*Haslett and Parnell*, 2008; *Parnell et al.*, 2008].

©2014. American Geophysical Union. All Rights Reserved.

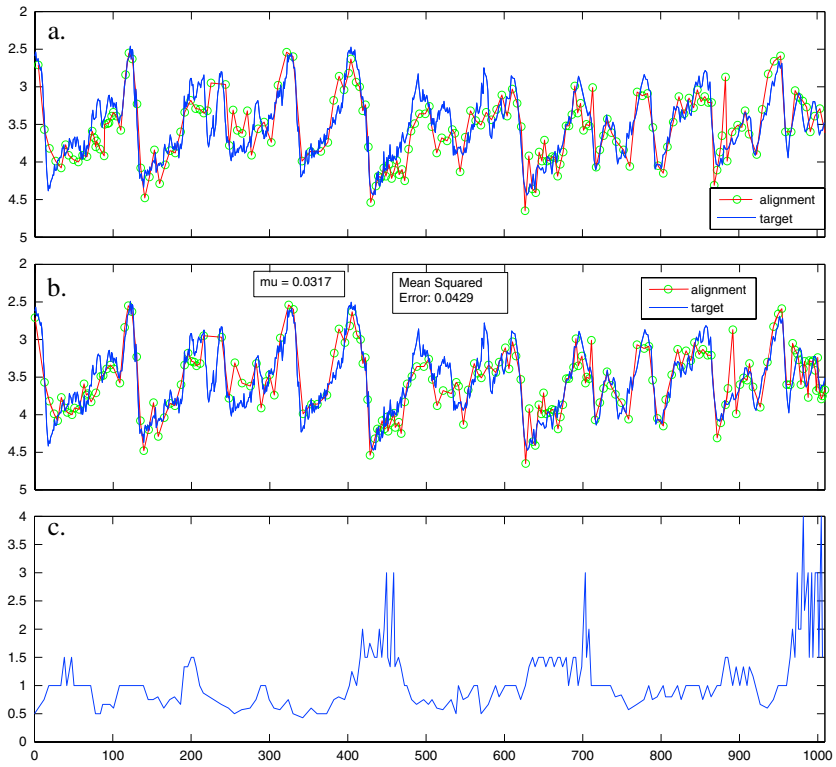

Figure 3. Alignment of DSDP Site 552 to LR04 stack. (a) Benthic δ^{18}O signals (in ‰ Vienna Peedee belemnite (VPDB)) aligned using the Match algorithm of *Lisiecki and Lisiecki* [2002], (b) median alignment using HMM-Match, and (c) corresponding accumulation rates from median alignment in normalized units, where 1 indicates the mean sedimentation rate of the input core; μ is the estimated mean deviation of δ^{18}O from stack. The 95% confidence bands calculated by HMM-Match for alignment are shown in Figure 4.

We estimated sedimentation rates using the mode of Bchron-modeled ages for each ^{14}C measurement and assumed constant sedimentation rates between radiocarbon measurements. By using the Bchron mode age based on complete down-core age models rather than the most probable age for each individual ^{14}C measurement, we avoided the physically impossible age reversals that are sometimes produced by most probable ^{14}C ages. Finally, we divided the sedimentation rates by each core's mean sedimentation rate to convert them into dimensionless sedimentation rate "ratios" that could be compared across different locations (Figure 2). This yielded a total length of sediment equivalent to 544 kyr across the 37 cores.

4. Results

4.1. Alignments

In total, we aligned the benthic δ^{18}O records from 35 cores (Table S1 in the supporting information) to the most recent 1 Myr of the LR04 stack [*Lisiecki and Raymo*, 2005]. Note that these are different cores than the 37 with ^{14}C data used to characterize expected sedimentation rate variability. Here we present the alignment results for two example cores to illustrate the performance of our algorithm, one with low resolution and the other with relatively high resolution. Figures for the remaining 33 input records are presented in the supporting information. Generally, the HMM-Match median alignment is quite similar to the optimal Match alignment, but the HMM-Match alignment returns 95% confidence bands and does not require the user to subjectively select penalty weightings. In section 4.2 we evaluate how the alignment confidence bands for all 35 cores vary as a function of the resolution and signal-to-noise ratio of the δ^{18}O input record.

4.1.1. Example 1: Deep Sea Drilling Program Site 552

The benthic δ^{18}O record from Deep Sea Drilling Program (DSDP) Site 552 [*Shackleton and Hall*, 1984] is a relatively low-resolution record, with an average sample spacing of 5.1 kyr for the last 1 Myr. The optimal alignment of Site 552 δ^{18}O to the LR04 stack using Match [*Lisiecki and Lisiecki*, 2002] (Figure 3a) is quite

©2014. American Geophysical Union. All Rights Reserved.

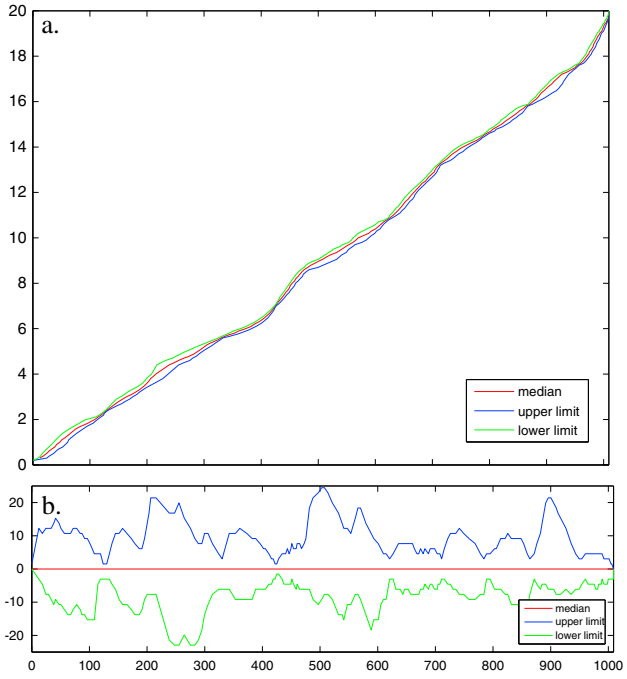

Figure 4. The 95% confidence band and median alignment calculated by HMM-Match for the alignment of DSDP Site 552 to the LR04 stack: (a) plotted as depth versus age and (b) plotted relative to the median age estimate.

similar to the median alignment of our HMM-Match (Figures 3b and 3c). However, the probabilistic alignment algorithm used here can also assess the uncertainty of this alignment as the 95% confidence band around the median alignment (Figure 4a). To illustrate these confidence bands more clearly, we also plot them as deviations from the median-predicted age (Figure 4b). The average width of these bands is 15.6 kyr, but they vary substantially with estimated age. In regions with wide confidence limits, the estimated ages should be interpreted with caution. Notice, for example, at an estimated age of 225 kyr the confidence band is ~40 kyr, which is nearly a full obliquity cycle.

4.1.2. Example 2: Ocean Drilling Program Site 1146

Late Pleistocene benthic δ^{18}O data from Ocean Drilling Program (ODP) Site 1146 [*Clements et al.*, 2008] are higher resolution than Site 552, with an average sample spacing of 1.8 kyr for the last 1 Myr. Again, there is a good agreement between the HMM-Match median fit and the optimal Match alignment (Figure 5). Not surprisingly, the higher-resolution δ^{18}O data yield tighter confidence bands, which have a mean width of 5.2 kyr (Figure 5d). Throughout the record the confidence limits are generally ±5 kyr or less, except in the interval between 800 and 875 kyr.

4.2. Confidence Band Relationships

Examination of the 35 alignments shows that the mean width of confidence bands increases linearly with the resolution of the input records (Figure 6a). We also find that the width of the confidence bands increases approximately linearly with δ^{18}O mean square residual errors (i.e., the degree of misfit between an individual δ^{18}O record and the stack), except for a few outliers (Figure 6b). For most of the outlier cores (e.g., ODP Sites 925, 982, and 984) it appears that relatively high resolution (mean sample spacing of ~2 kyr) is sufficient to keep alignment uncertainty low despite high levels of noise (see alignment figures in the supporting information). In some cases, this "noise" probably reflects millennial-scale climate variability [e.g., *Oppo et al.*, 1998] that differs between cores due to core location and/or differing amounts of smoothing by sediment bioturbation.

The alignments produced by both Match and HMM-Match result from the balance of minimizing the departures of the input proxy values from those in the stack against limiting unrealistic changes in sedimentation rates. Thus, neither algorithm fully adheres to both of these two competing interests. Figure 7b, which shows a histogram of the sedimentation rates employed by HMM-Match over the 35 alignments, indicates that the alignment ratio of 1:1 is used more frequently than expected based on our radiocarbon estimates. A number of possible reasons for the prevalence of 1:1 are discussed in section 5.1; however, overall, this result indicates that the algorithm successfully avoids overtuning the δ^{18}O signals. Additionally, sensitivity tests using sedimentation rate models that vary by a factor of 2 (Figures 7a and 7c) show that the algorithm is only weakly sensitive to potential biases in the radiocarbon-based sedimentation rate estimates, e.g., due to site selection or the analysis of relatively young, less-compacted sediment.

4.3. Stacked Results

The median δ^{18}O values of the 35 records aligned here yield very similar results to the LR04 stack (Figure 8a), which verifies that the HMM-Match and Match algorithms align the orbital and suborbital-scale features of

©2014. American Geophysical Union. All Rights Reserved.

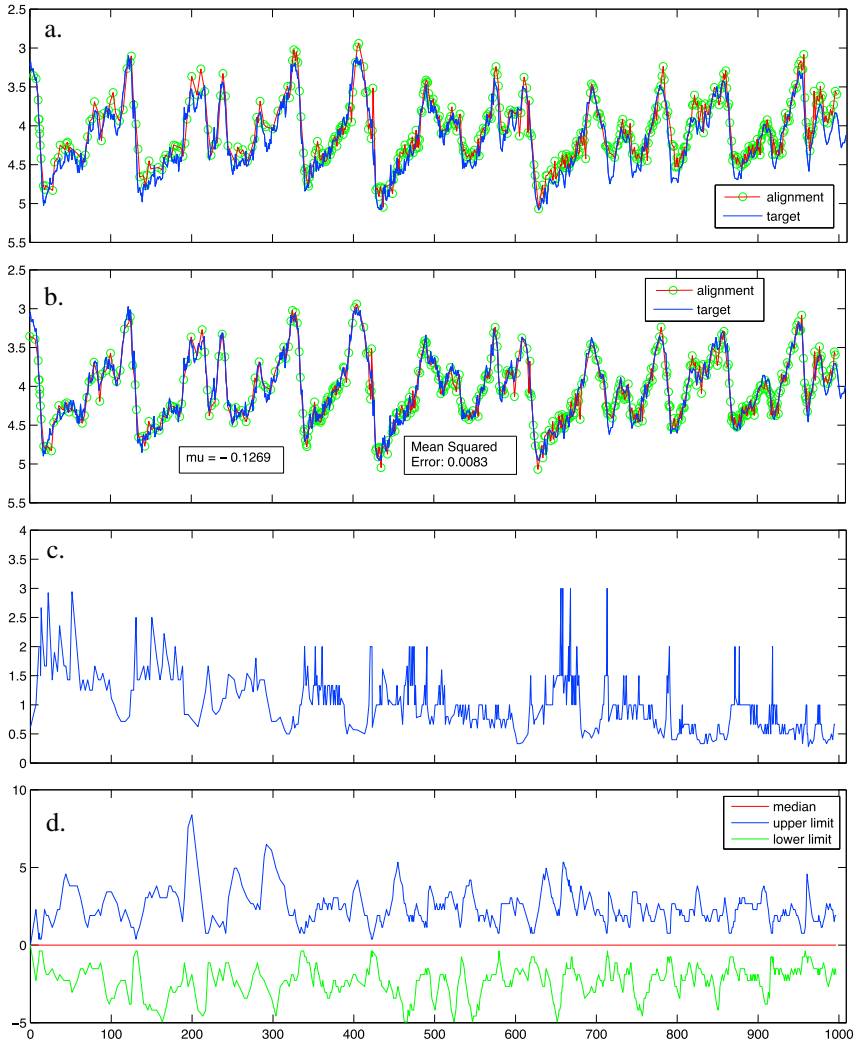

Figure 5. Alignment of ODP Site 1146 to LR04 stack. (a) Benthic δ^{18}O signals (in ‰ VPDB) aligned using the Match algorithm of *Lisiecki and Lisiecki* [2002]. (b) Median alignment using HMM-Match and (c) corresponding accumulation rates from median alignment in normalized units, where 1 indicates the mean sedimentation rate of the input core; μ is the estimated mean deviation of δ^{18}O from stack. (d) The 95% confidence band around the median alignment.

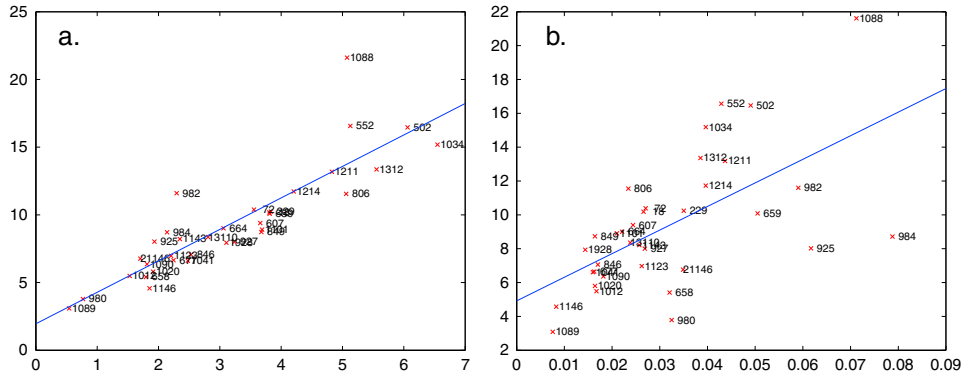

Figure 6. Scatterplots showing the mean width of alignment confidence bands for 35 late Pleistocene benthic δ^{18}O records versus (a) the resolution of the record (measured as mean sample spacing in kiloyear) and (b) the mean square error in residual δ^{18}O (in ‰ VPDB). The number next to each point corresponds to the core name for the δ^{18}O record; see the supporting information for alignment figures and citations.

©2014. American Geophysical Union. All Rights Reserved.

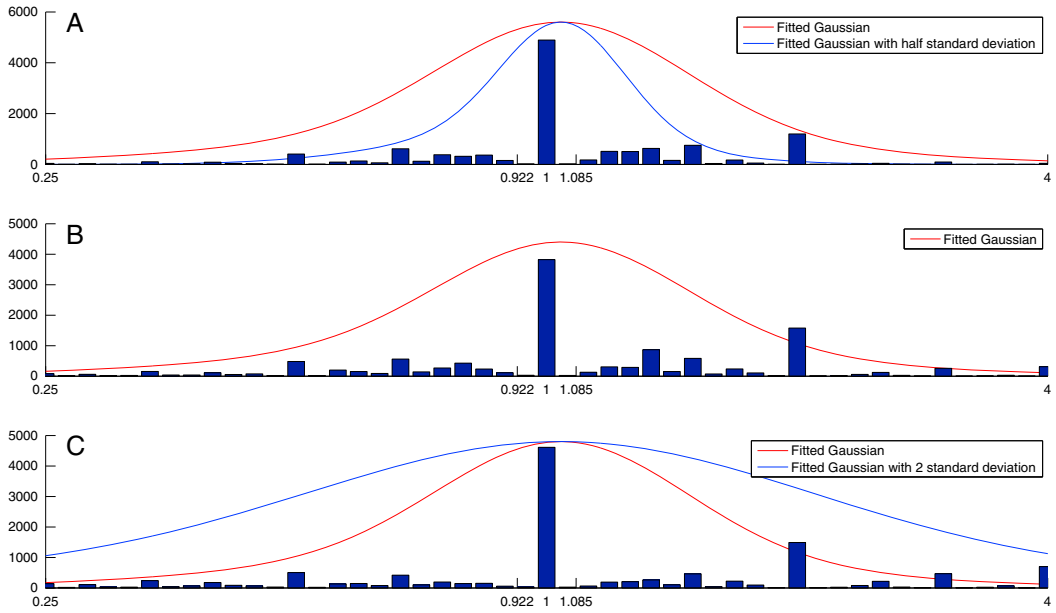

Figure 7. Sensitivity tests comparing the log-Gaussian distribution used as input parameters to the sedimentation rate (transition) model (lines) with histograms of alignment ratios output by the model for all 35 cores aligned to the LR04 stack. The red line in all plots shows the log-Gaussian distribution for alignment ratios estimated from radiocarbon data (shown in Figure 2). During sensitivity tests, the standard deviation of this distribution was adjusted to (a) half of the estimated value (blue), (b) the actual estimated value (no adjustment), and (c) twice the estimated value (blue). These extreme changes had relatively little effect on the alignment ratios output by the model (histograms), indicating that our results are not highly sensitive to uncertainty in this model parameter.

δ^{18}O records in similar ways. The HMM-Match stack is smoother than the LR04 stack, and its glacial cycles have slightly smaller amplitude, but this likely results from differences in the techniques used to combine the data from the 35 records rather than differences in the alignments themselves. The LR04 stack is the average of all individual δ^{18}O observations within each time interval, whereas for the HMM-Match alignments we show the median of δ^{18}O values (with data from each core interpolated to even 1 kyr spacing). Inherently, the median should be less sensitive to noise than the mean and, hence, is expected to yield smoother changes in δ^{18}O. Differences between the median and LR04 stack may also occur because the LR04 stack contains data from 12 additional cores within the last 1 Myr. Confidence intervals for the HMM-Match stack (Figure 8a) reflect the combined effects of measurement error and spatial variability in benthic δ^{18}O.

The mean width of alignment confidence intervals for the 35 cores varies as a function of time (Figure 8b), demonstrating that alignment uncertainty is consistently lower for certain portions of the δ^{18}O record. The smallest alignment uncertainties (3–6 kyr) occur during times of rapid δ^{18}O change (Figure S5), i.e., glacial terminations as well as times of rapid cooling and/or ice growth. When δ^{18}O is changing rapidly, small changes in alignment will produce large δ^{18}O residuals between the input and the alignment target. In contrast, alignment uncertainty is much larger (10–16 kyr) when δ^{18}O is changing slowly or nearly constant for 10–20 kyr; this is because different alignments in these portions of the records all produce similar δ^{18}O residuals. Although glacial terminations typically have lower average alignment uncertainty, it is important to note that these confidence bands do not include uncertainty associated with the assumption of globally synchronous benthic δ^{18}O change as discussed below.

5. Discussion

5.1. Sedimentation Rates

Although we assume that marine sedimentation rates and changes in these rates are reasonably estimated by radiocarbon data from 37 cores, sensitivity tests demonstrate that the algorithm's results are only very weakly sensitive to this assumption (Figure 7). The predominance of 1:1 alignment ratios in HMM-Match

©2014. American Geophysical Union. All Rights Reserved.

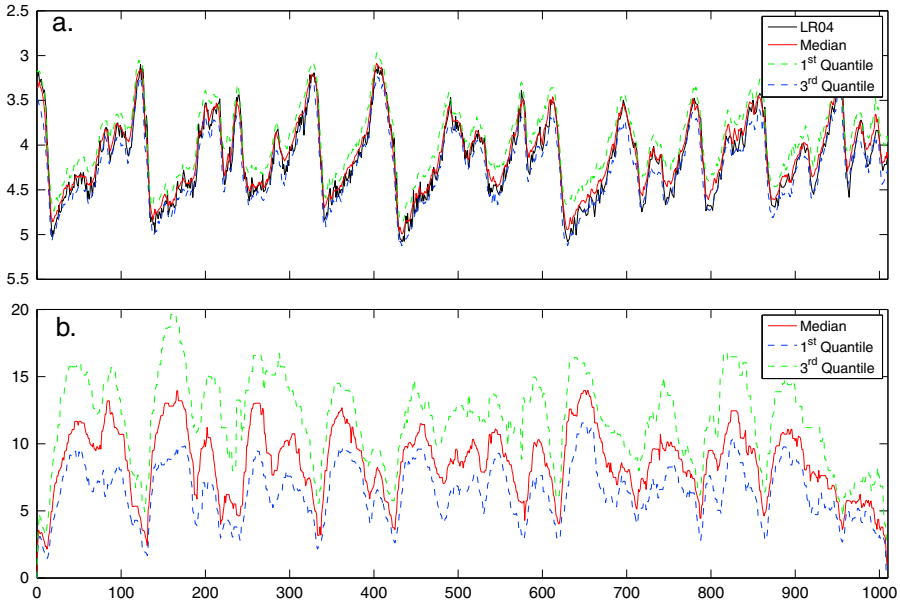

Figure 8. Stacked results for all 35 cores aligned by HMM-Match. We plot the median values (red) and first and third quantiles (blue and green dotted lines) for (a) aligned δ^{18}O values (in ‰ VPDB) and (b) the widths of alignment confidence intervals as a function of time.

results compared to our radiocarbon-based model clearly demonstrates that the algorithm avoids overtuning the δ^{18}O records (i.e., the algorithm does not artificially distort sedimentation rates in an effort to align noise). Several factors may explain why ratios of 0.8–1.2, which are common in the radiocarbon-based estimates, may frequently appear as 1:1 in our alignments. First, small millennial-scale variations in sedimentation rates will have very little effect on benthic δ^{18}O records that are smoothed by bioturbation and have a typical sample spacing of 2–4 kyr. Therefore, these small sedimentation rate changes are unlikely to be reconstructed by the process of δ^{18}O alignment. Second, sedimentation rate changes may be underestimated when cores are aligned to the LR04 stack because the stack's age model was designed to minimize changes in global average sedimentation rate [*Lisiecki and Raymo*, 2005]. Third, apparent sedimentation rate variability in the radiocarbon data may be artificially enhanced by noise, surface reservoir age changes, and the necessity of analyzing only the youngest ($<$40 kyr), least consolidated portion of the sedimentary record.

Another pattern evident from the histograms in Figure 7 is that our discretization of time and depth impacts the distribution of alignment ratios, particularly during compression. Ratios of 4:3, 2:1, 3:2, and 4:1 are clearly more frequent than neighboring ratios. Nonetheless, by applying these ratios to different portions of an input record, the algorithm has sufficient degrees of freedom to approximate all possible δ^{18}O alignments using a finite number of realizations, analogous to the original Match algorithm [*Lisiecki and Lisiecki*, 2002].

5.2. Absolute Age Uncertainty

The algorithm excludes consideration of uncertainty in the age model of the LR04 stack because this uncertainty has only been estimated qualitatively [*Lisiecki and Raymo*, 2005]. Thus, the alignment confidence limits given here are lower bounds on the absolute ages assigned to points in input records. Our current work is a first step toward the development of a probabilistic δ^{18}O stack, which accounts for all sources of uncertainty (i.e., from the alignment of individual δ^{18}O records and the development of orbitally tuned age estimates).

Two tasks remain in our efforts to fully assess the uncertainty in ages derived from the stack: uncertainty stemming from the multiple alignments of input records required to probabilistically rebuild the stack and uncertainty in the assignment of ages based on orbital parameters to intervals in the stack. The bulk of the uncertainty in the assignment of absolute ages using benthic δ^{18}O proxies is an expected stem from the second task. Because measured benthic δ^{18}O values after probabilistic alignment closely track the LR04 mean values (Figure 8), the application of probabilistic alignment methods to reestimate average benthic δ^{18}O stack values will likely produce δ^{18}O values quite similar to the LR04 stack. However, the LR04 age model

©2014. American Geophysical Union. All Rights Reserved.

is constructed by aligning the stacked δ^{18}O values to a simple model driven by orbital parameters [*Imbrie and Imbrie*, 1980], and orbital parameters have been shown to predict the stack's δ^{18}O values with large residual errors [*Ruggieri et al*., 2009]. Because our results demonstrate that alignment confidence bands increase substantially with residual errors, the orbital tuning of δ^{18}O is almost certainly associated with large age uncertainties.

5.3. Relative Age Uncertainty

Although the uncertainty in the absolute ages of the stack may be considerable, often researchers are more interested in evaluating the relative ages of events in two aligned records (i.e., leads and lags). HMM-Match can provide a statistical evaluation of these lead-lag relationships. Because HMM-Match provides mutually consistent, independent alignments of input points to stack points, confidence bands for the alignments of two different records to the LR04 stack can be used to calculate the probability that a point in one of these records leads/lags a point in the other without the use of the LR04 age assignments. Specifically, if the $(1-\alpha)$ % upper Bayesian confidence band for the age of point *i* in record one does not overlap the $(1-\alpha)$ % lower band for the age of point *k* in record two, then the probability that point *i* leads point *k* is $p \leq \left(1-\left(\alpha/2\right)^2\right)$. Here *p* is the Bayesian analog of the frequentist's *p* value obtained in significance testing. This statistical inference requires only the assumption that younger sediment in the LR04 stack is deposited on top of older sediment. With additional information about the uncertainty of relative ages in the target record, we would also be able to draw inferences on magnitudes of these lead-lag relationships without relying on the absolute ages in the LR04 stack.

An additional factor to consider with respect to evaluating potential leads and lags between cores is that our model assumes that changes in δ^{18}O are synchronous throughout the ocean. However, the timing of δ^{18}O change between the Atlantic and Pacific may differ by as much as 4 kyr during glacial terminations [*Skinner and Shackleton*, 2005; *Lisiecki and Raymo*, 2009]. This source of uncertainty is excluded from our algorithm because insufficient data are available to quantify possible differences in the timing of δ^{18}O change between specific sites. Thus, although we find that glacial terminations typically have the smallest alignment uncertainties (Figure 8b), these uncertainty estimates do not account for possible diachronous δ^{18}O change during terminations. One should also consider this additional source of uncertainty when evaluating potential leads and lags between cores, particularly those sampling different water masses.

5.4. Other Applications

The HMM-Match algorithm presented here could also be applied to the alignment of other paleoclimate proxies measured in marine sediment cores as long as at least one of the two records is nearly complete at a specified level of resolution and residuals between the aligned proxy records are approximately normal. In applications where data are not normally distributed, it is frequently possible to use a power transformation to obtain approximate normality. With regard to alignment target resolution, we advise that no more than 10% of the chosen time intervals are missing data and that the missing data do not cluster together. If linear interpolation is used to fill in only a small number of missing target data points, the resulting confidence bands will be shrunk slightly from their proper values. Thus, this algorithm could also be used to align δ^{18}O (or other proxy data) from one sediment core to another individual core, rather than a stack, as long as the target core is of sufficiently high resolution.

HMM-Match is specifically designed to align marine sediment cores because its transition model uses information about sedimentation rate variability in the deep ocean. This model would likely also work well for the alignment of some lacustrine cores. However, HMM-Match could potentially be used to align other climate archives, such as corals, speleothems, or ice cores, if the parameters of its transition model were adjusted appropriately.

6. Conclusions

The assessment of uncertainty is of growing importance in paleoclimate research. We present a method for the assessing age uncertainty based on the alignment of benthic δ^{18}O records to the LR04 stack [*Lisiecki and Raymo*, 2005]. This algorithm is designed specifically to quantify uncertainty associated with the alignment process and thus does not address uncertainty in the absolute age of the target record or possible spatial

©2014. American Geophysical Union. All Rights Reserved.

differences in the timing of δ^{18}O changes. Nevertheless, the confidence bands it returns provide the means to calculate the probability that a point in one record leads (lags) a point in another record independent of uncertainty in the absolute ages assigned in the target, e.g., the LR04 stack.

We find that age uncertainties, measured as the mean width of 95% confidence bands, typically range from 5 to 15 kyr. Tighter confidence bands are associated with cores that have higher-resolution records and smaller residual mean squared errors relative to the target, highlighting the importance in future studies of high-resolution, high-quality data. Our findings of considerable variation in the widths of these bands over the range of estimated ages also demonstrate the importance of examining the full plot of estimated ages because some portions of cores have confidence bands as large as 20–30 kyr. During these intervals of large alignment uncertainty, proxy responses should be interpreted with caution. In general, confidence bands are smaller in portions of the record where δ^{18}O changes rapidly, such as glacial terminations, and larger where δ^{18}O changes more slowly. Lastly, we describe a method by which these confidence bands can be used to obtain Bayesian tests of the lead/lag relationships between cores.

Acknowledgments

Web service and open sources software for HMM-Match is available at http://ccmbweb.ccv.brown.edu/cgi-bin/geo_align.pl. Data supporting Figure 2 are available in Table S3 in the supporting information, calculated using radioarbon data from references in Table S1 in the supporting information. Match alignments to the LR04 stack are available at http://lorraine-lisiecki.com/stack.html. We appreciate William Thompson's support in the development of the software and his work on preparing the web service and open sources software and Stuart Geman for his assistance in the formulation of the HMM algorithm. Joseph Stern provided Bchron radiocarbon age models. Support provided by NSF-OCE 1025438.

References

Adkins, J. F., K. McIntyre, and D. P. Schrag (2002), The salinity, temperature, and d18O of the glacial deep ocean, *Science*, *298*, 1769–1773.

Bellman, R. E. (1961), *Dynamic Programming*, Princeton Univ. Press, Princeton, New Jersey.

Carvalho, L. E., and C. E. Lawrence (2008), Centroid estimation in discrete high-dimensional spaces with application in biology, *Proc. Natl. Acad. Sci. U.S.A.*, *105*, 3209–3214.

Channell, J. E. T., C. Xuan, and D. A. Hodell (2009), Stacking paleointensity and oxygen isotope data for the last 1.5 Myr (PISO-1500), *Earth Planet. Sci. Lett.*, *283*, 14–23.

Clements, S. C., W. L. Prell, Y. Sun, Z. Liu, and G. Chen (2008), Southern Hemisphere forcing of Pliocene δ^{18}O and the evolution of Indo-Asian monsoons, *Paleoceanography*, *23*, PA4210, doi:10.1029/2008PA001638.

Durbin, R., S. Eddy, A. Krogh, and G. Mitchison (1998), *Biological Sequence Analysis*, vol. 356, Cambridge Univ. Press, Cambridge, U. K.

Elderfield, H., P. Ferretti, M. Greaves, S. Crowhurst, I. N. McCave, D. Hodell, and A. M. Piotrowski (2012), Evolution of ocean temperature and ice volume through the mid-Pleistocene climate transition, *Science*, *337*, 704–709.

Haslett, J., and A. Parnell (2008), A simple monotone process with application to radiocarbon-dated depth chronologies, *J. R. Stat. Soc. Ser. C (Appl. Statist.)*, *57*(4), 399–418.

Hodell, D. A., C. D. Charles, J. H. Curtis, P. G. Mortyn, U. S. Ninnemann, and K. A. Venz (2003), Data report: Oxygen isotope stratigraphy of ODP Leg 177 Sites 1088, 1089, 1090, 1093, and 1094, *Proc. Ocean Drill. Program Sci. Results*, *177*, 1–26.

Huybers, P., and C. Wunsch (2004), A depth-derived Pleistocene age model: Uncertainty estimates, sedimentation variability, and nonlinear climate change, *Paleoceanography*, *19*, PA1028, doi:10.1029/2002PA000857.

Imbrie, J., and J. Z. Imbrie (1980), Modeling the climatic response to orbital variations, *Science*, *207*, 943–953.

Imbrie, J., et al. (1992), On the structure and origin of major glaciation cycles. 1. Linear responses to Milankovitch forcing, *Paleoceanography*, *7*, 701–738, doi:10.1029/92PA02253.

Lang, N., and E. W. Wolff (2011), Interglacial and glacial variability from the last 800 ka in marine, ice and terrestrial archives, *Clim. Past*, *7*, 361–380.

Lawrence, K. T., D. M. Sigman, T. D. Herbert, C. A. Riihimaki, C. T. Bolton, A. Martinez-Garcia, A. Rosell-Mele, and G. H. Haug (2013), Time-transgressive North Atlantic productivity changes upon Northern Hemisphere glaciation, *Paleoceanography*, *28*, 740–751, doi:10.1002/2013PA002546.

Levinson, S., L. Rabiner, and M. Sondhi (1983), An introduction to the application of the theory of probabilistic functions of a Markov process of automatics speech recognition, *Bell Syst. Tech. J.*, *62*, 1035–1074.

Lisiecki, L., and P. Lisiecki (2002), Application of dynamic programming to the correlation of paleoclimate records, *Paleoceanography*, *17*(D4), 1049, doi:10.1029/2001PA000733.

Lisiecki, L., and M. A. Raymo (2005), Pleistocene stack of 57 globally distributed benthic δ^{18}O records, *Paleoceanography*, *20*, PA1003, doi:10.1029/2004PA001071.

Lisiecki, L. E., and M. E. Raymo (2009), Diachronous benthic δ^{18}O responses during late Pleistocene terminations, *Paleoceanography*, *24*, PA3210, doi:10.1029/2009PA001732.

Lisiecki, L. E., M. E. Raymo, and W. B. Curry (2008), Atlantic overturning responses to Late Pleistocene climate forcings, *Nature*, *456*, 85–88.

Martinson, D. G., N. G. Pisias, J. D. Hays, J. Imbrie, T. C. Moore, and N. J. Shackleton (1987), Age dating and the orbital theory of the ice ages: Development of a high-resolution 0 to 300,000-year chronostratigraphy, *Quat. Res.*, *27*, 1–30.

Oppo, D. W., J. F. McManus, and J. L. Cullen (1998), Abrupt climate events 500,000 to 340,000 year ago: Evidence from subpolar North Atlantic sediments, *Science*, *279*, 1335–1338.

Ostermann, D. R., and W. B. Curry (2000), Calibration of stable isotopic data: An enriched delta O-18 standard used for source gas mixing detection and correction, *Paleoceanography*, *15*, 353–360, doi:10.1029/1999PA000411.

Pardo, B., J. Shifrin, and W. Birmingham (2004), Name that tune: A pilot study in finding a melody from a sung query, *J. Am. Soc. Inf. Sci. Technol.*, *55*, 283–300.

Parnell, A. C., J. Haslett, J. R. M. Allen, C. E. Buck, and B. Huntley (2008), A flexible approach to assessing synchroneity of past events using Bayesian reconstructions of sedimentation history, *Quat. Sci. Rev.*, *27*, 1872–1885, doi:10.1016/j.quascirev.2008.07.009.

Peterson, C. D., L. E. Lisiecki, and J. V. Stern (2014), Deglacial whole-ocean δ^{13}C change estimated from 480 benthic foraminiferal records, *Paleoceanography*, *29*, 549–563, doi:10.1002/2013PA002552.

Prell, W. L., J. Imbrie, D. G. Martinson, J. J. Morley, N. G. Pisias, N. J. Shackleton, and H. F. Streeter (1986), Graphic correlation of oxygen isotope stratigraphy: Application to Late Quaternary, *Paleoceanography*, *1*, 137–162, doi:10.1029/PA001i002p00137.

Rabiner, L. R. (1989), A tutorial on hidden Markov models and selected application in speech recognition, *Proc. IEEE*, *77*(2), 257–286, doi:10.1109/5.18626.

©2014. American Geophysical Union. All Rights Reserved.

Raymo, M. E., D. W. Oppo, B. P. Flower, D. A. Hodell, J. F. McManus, K. A. Venz, K. F. Kleiven, and K. McIntyre (2004), Stability of North Atlantic water masses in face of pronounced climate variability during the Pleistocene, *Paleoceanography*, *19*, PA2008, doi:10.1029/2003PA000921.

Reimer, P. J., et al. (2009), IntCal09 and Marine09 radiocarbon age calibration curves, 0-50,000 years cal BP, *Radiocarbon*, *51*, 1111–1150.

Sankoff, D., and J. Kruskal (Eds) (1999), *Time Warps, String Edits, and Macromodecules: The Theory and Practice of Sequence Comparison*, CSLI Publ., Stanford Univ., Stanford, Calif.

Shackleton, N. J., and M. A. Hall (1984), Oxygen and carbon isotope stratigraphy of DSDP Hole 552A: Plio-Pleistocene glacial history, *Initial Rep. Deep Sea Drill. Proj.*, *81*, 599–609.

Shackleton, N. J., M. A. Hall, and E. Vincent (2000), Phase relationships between millennial-scale events 64,000–24,000 years ago, *Paleoceanography*, *15*, 565–569, doi:10.1029/2000PA000513.

Skinner, L. C., and N. J. Shackleton (2005), An Atlantic lead over Pacific deep-water change across Termination I: Implications for the application of the marine isotope stage stratigraphy, *Quat. Sci. Rev.*, *24*, 571–580.

Smith, T. F., and M. S. Waterman (1981), Identification of common molecular subsequences, *J. Mol. Biol.*, *147*, 195–197.

Stern, J. V., and L. E. Lisiecki (2013), North Atlantic circulation and reservoir age changes over the past 41,000 years, *Geophys. Res. Lett.*, *40*, doi:10.1002/grl.50679.

Webb, B. J., J. Liu, and C. E. Lawrence (2002), BALSA: Bayesian algorithm for local sequence alignment, *Nucleic Acids Res.*, *30*, 1268–1277.

Zhu, J., J. Liu, and C. E. Lawrence (1998), Bayesian adaptive sequence alignment algorithms, *Bioinforatics*, *14*, 1–15.

©2014. American Geophysical Union. All Rights Reserved.