

Dynamic time warping of palaeomagnetic secular variation data

Cedric J. Hagen¹,^{ORCID} Brendan T. Reilly,^{1,2} Joseph S. Stoner¹ and Jessica R. Creveling¹

¹College of Earth, Ocean, and Atmospheric Sciences, Oregon State University, Corvallis, OR 97330, USA. E-mail: hagence@oregonstate.edu

²Scripps Institution of Oceanography, University of California, San Diego, La Jolla, CA 92037, USA

Accepted 2020 January 3. Received 2019 December 19; in original form 2019 July 31

SUMMARY

We present and make publicly available a dynamic programming algorithm to simultaneously align the inclination and declination vector directions of sedimentary palaeomagnetic secular variation data. This algorithm generates a library of possible alignments through the systematic variation of assumptions about the relative accumulation rate and shared temporal overlap of two or more time-series. The palaeomagnetist can then evaluate this library of reproducible and objective alignments using available geological constraints, statistical methods and expert knowledge. We apply the algorithm to align previously (visually) correlated medium to high accumulation rate northern North Atlantic Holocene deposits (10^1 – 10^2 cm ka⁻¹) with strong radiocarbon control. The algorithm generates plausible alignments that largely conform with radiocarbon and magnetic acquisition process uncertainty. These alignments illustrate the strengths and limitations of this numerical approach.

Key words: Atlantic Ocean; Magnetostratigraphy; Palaeomagnetic secular variation; Palaeomagnetism; Statistical methods.

1 INTRODUCTION

Palaeomagnetic secular variation (PSV) describes the spatio-temporal variation of Earth's magnetic field at decadal to millennial timescales arising from convection in Earth's liquid outer core (Johnson & McFadden 2007; Lund 2007). Palaeomagnetic archives record the history of changes in geomagnetic field morphology (inclination and declination) and intensity (both relative and absolute). Utilization of PSV for the purpose of stratigraphic correlation has a long and rich history (e.g. Thompson 1973). While relative palaeointensity (RPI) reconstructions can be complicated by lithologic variability (e.g. Schwartz *et al.* 1996; Mazaud 2006) that requires strict quality assessment (Tauxe 1993; Stoner & St-Onge 2007), PSV direction reconstructions are more robust to this variability (e.g. Reilly *et al.* 2018). However, certain criteria should still be met for PSV direction stratigraphic correlation, including:

(i) The vector time-series should faithfully record the geomagnetic field at or soon after the time of deposition, hold well-defined characteristic remanent magnetizations (ChRM) by magnetic mineralogies indicative of a primary magnetization, and display no signs of physical sediment deformation. Ideally, a PSV vector time-series should also be replicated in multiple cores within the basin or on a regional level.

(ii) Accumulation rates and sampling resolution should be high enough to resolve centennial to millennial trends in PSV. These rates and resolution may be higher than that required for palaeoclimate reconstructions, as the magnetic acquisition process can act as a low-pass filter (e.g. Lund & Keigwin 1994; Roberts & Winklhofer 2004; Balbas *et al.* 2018).

(iii) The records should be geographically limited such that they share a common geomagnetic history (see Korte *et al.* 2018 for recent analysis). However, records could be compared over greater distances if systematic relationships within regional PSV data are understood, as has been suggested for centennial to millennial timescale variations (e.g. Valet *et al.* 2008; Nilsson *et al.* 2010, 2011; Stoner *et al.* 2013; Walczak *et al.* 2017).

In most PSV stratigraphic correlation studies, PSV signals are aligned through supervised or 'wobble-matched' correlations, using a range of approaches from a single parameter (e.g. inclination, declination or RPI) to the whole vector. As any correlation is non-unique and possibly biased by the palaeomagnetist, we present a new tool based on dynamic programming to help evaluate possible PSV correlations. Application of dynamic programming to align Quaternary palaeoclimate and RPI palaeomagnetic records has successfully contributed iconic records in each field (e.g. Clark 1985; Lisiecki & Lisiecki 2002; Lisiecki & Raymo 2005; Channell *et al.* 2009; Haam & Huybers; Xuan *et al.* 2016). Here we introduce an algorithm that utilizes a dynamic programming technique (dynamic time warping; DTW) to simultaneously align PSV inclination and declination direction vector time-series. This algorithm generates a library of potential alignments through the systematic variation of assumptions about relative accumulation rate and shared temporal overlap between multiple vector time-series. The palaeomagnetist can then evaluate this alignment library using available geological constraints and expert knowledge, both of which can be difficult or impossible to quantify and fully incorporate in a purely mathematical approach. Because this study focuses on sedimentary palaeomagnetic archives, we refer to assumptions of accumulation rate.

However, the method can theoretically be applied to volcanic or archeologic palaeomagnetic archives with stratigraphy and, accordingly, can reflect different assumptions, such as eruption rate.

To illustrate the strengths and limitations of this method, we apply a DTW algorithm to align three well studied PSV directional vector time-series from medium to high accumulation rate deposits (10^1 – 10^2 cm ka⁻¹) in the northern North Atlantic (Stoner *et al.* 2007, 2013; Fig. 1). These vector time-series fulfill the criteria for PSV stratigraphic correlation outlined above and have previously been used as templates to assess or constrain the chronology of regional records (e.g. Ólafsdóttir *et al.* 2013, 2019; Strunk *et al.* 2018; Caron *et al.* 2019; Reilly *et al.* 2019). We report that this algorithm can successfully generate plausible PSV alignments generally within radiocarbon and magnetic acquisition process uncertainties.

2 MATERIALS AND METHODS

2.1 Data sets used

We focus our discussion on two numerical experiments, where a candidate PSV vector time-series is aligned, or warped, to a target PSV vector time-series. We use three cores from the northern North Atlantic that have been previously studied, documented to record PSV, and have strong independent radiocarbon chronologies. All data sets used are unsmoothed and as reported in their original publications. For our first experiment, we choose PSV directions from the Marion Dufresne II Core MD99-2269 (66.63°N, 20.85°W; 365 m water depth; 2533 cm core length), recovered during the 1999 IMAGES (International Marine Global Changes Study) expedition, as the target magnetic vector time-series ('data set' herein); we choose PSV directions from Core MD99-2322 (67.14°N, 30.83°W; 714 m water depth; 2617 cm core length), recovered during the same expedition (Stoner *et al.* 2007; Fig. 1), as the candidate data set. Core 2269 was raised from Húnaflói, North Iceland Shelf, while Core 2322 was raised from the deepest part of the Kangerlussuaq Trough, Southeast Greenland Shelf. A supervised correlation between the two PSV vector time-series has been shown to be consistent with each core's independent radiocarbon chronology (Stoner *et al.* 2007), and supervised PSV correlations of Core 2269 to Icelandic lake sediments are consistent with tephra stratigraphy (Ólafsdóttir *et al.* 2013). Cores 2269 and 2322 chronologies are constrained by 27 and 20 radiocarbon dates, respectively, from mollusks and mixed benthic foraminifera (Dunhill *et al.* 2004; Stoner *et al.* 2007) and have been subject to detailed tephra analysis (Kristjánsdóttir *et al.* 2007; Jennings *et al.* 2014). Long-term accumulation rates for both cores are high. Core 2269 accumulation rates vary between ~300 and 400 cm ka⁻¹ from 9.5 to 12 ka and ~100–300 cm ka⁻¹ from 9.5 ka to present (Stoner *et al.* 2007). Core 2322 accumulation rates are ~300–600 cm ka⁻¹ from 9 to 12 ka and ~100 to 200 from 9 ka to present (Stoner *et al.* 2007). This experiment tests what we consider ideal PSV data that are high resolution and less likely to contain large offsets between the sediment and magnetic age (e.g. deMenocal *et al.* 1990; Suganuma *et al.* 2010; Stoner *et al.* 2013; Simon *et al.* 2018) or signal attenuation (e.g. Lund & Keigwin 1994; Roberts & Winklhofer 2004; Valet & Fournier 2016; Balbas *et al.* 2018) that can result from depth-dependent magnetic acquisition processes (i.e. post-depositional remanent magnetization (pDRM); Irving & Major 1964; Verosub 1977; Egli & Zhao 2015) or sampling method (e.g. Weeks *et al.* 1993). These data are also considered ideal for alignment as they span almost exactly the same duration and have similar accumulation rates (as discussed later).

In a second numerical experiment, we again use Core 2269 as the target data set but use PSV from the Holocene interval of Integrated Ocean Drilling Program (IODP) Expedition 303 Site U1305 (57.48°N, 48.53°W; 3459 m water depth; focusing on the upper 555 cm) as the candidate data set (Channell *et al.* 2006; Stoner *et al.* 2013, Fig. 1). Site U1305 was drilled on the lee side of the Eirik Ridge, a deep sea contourite deposit, and is noted for containing expanded Pleistocene interglacial sediments compared with glacial sediments (Mazaud *et al.* 2012). The Holocene PSV vector time-series has been studied in detail and displays broadly consistent signals, although offset in time, with other regional records (Stoner *et al.* 2013). Site U1305 chronology is constrained by 17 radiocarbon dates on mixed planktonic foraminifera, with accumulation rates ranging from about 40–80 cm ka⁻¹ (Stoner *et al.* 2013, and this study; Table 1). As Site U1305 is of lower temporal resolution than Core 2269 and was demonstrated to have uncertainties associated with pDRM acquisition processes (Stoner *et al.* 2013), this experiment tests the algorithm performance on good data with realistic uncertainties inherent to sedimentary records of palaeomagnetism. Site U1305 also presents challenges for numerical alignment as it has lower accumulation rates relative to the target data set and spans a subset of the time covered by the target data set (as discussed later). For this exercise, Site U1305 is relocated by calculating its virtual geomagnetic pole (VGP) path and then calculating inclination and declination of that VGP path at the location of Core 2269. This assumes a common geomagnetic field over the 1750 km distance between the two sites, but accounts for the ~9° latitude and ~28° longitude difference (*cf.* Korte *et al.* 2018; Reilly *et al.* 2018).

All radiocarbon dates from Stoner *et al.* (2007, 2013) and reported here (Table 1) were calibrated using the MARINE13 calibration curve (Reimer *et al.* 2013), the MATCAL calibration software (Lougheed & Obrochta 2016), and modeled for age–depth relationships using the Undatable software package with a conservative xfactor of 1 and no bootstrapping (Lougheed & Obrochta 2019). Although marine radiocarbon reservoir effects likely varied at these locations over this time interval (e.g. Eiriksson *et al.* 2011; Wanamaker *et al.* 2012) and have been partially corrected using the known terrestrial ages of tephra layers (Stoner *et al.* 2013), for this exercise we simply use ΔR values of 0 yr to focus on testing whether the stratigraphic correlations derived from PSV are consistent with the relative radiocarbon stratigraphy.

2.2 PSV dynamic time warping

2.2.1 Dynamic time warping

In computer programming, dynamic programming is a technique for optimizing complex problem solving by minimizing the number of calculations. This is accomplished by dividing a complex problem into smaller subproblems and storing the answers for later use to ensure maximum efficiency. DTW is an alignment technique for time-series data sets that uses generalized dynamic programming, which was first applied to spoken word recognition (Sakoe & Chiba 1978). Dynamic programming has been shown to have geological applications (e.g. Clark 1985; Thompson & Clark 1989; Lisiecki & Lisiecki 2002; Lisiecki & Raymo 2005; Channell *et al.* 2009; Haam & Huybers; Hay *et al.* 2019). Dynamic programming had been applied to palaeointensity records using the Match algorithm (Lisiecki & Lisiecki 2002; Channell *et al.* 2009). Furthermore, Thompson & Clark (1989) used a slotting technique to align palaeointensity, palaeomagnetic reversal and PSV time-series with the aid of manual

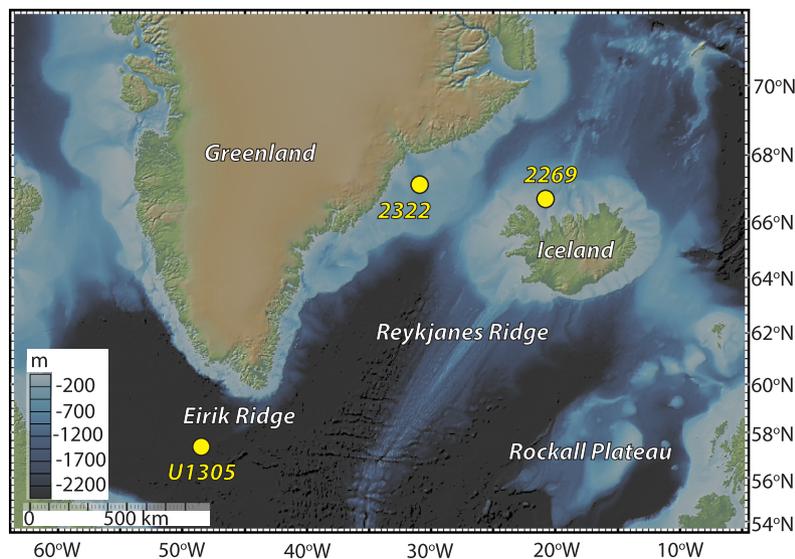


Figure 1. Regional Map with Core Locations. Cores MD99 2269 and 2322, and Integrated Ocean Drilling Program (IODP) Expedition 303 Site U1305 indicated with yellow dots.

Table 1. Additional site U1305 radiocarbon dates.

Depth (mcd)	Reported Age ^{14}C years BP	Min 2σ Cal. years BP	Med. Prob.	Max 2σ Cal. years BP
4.805	7815 ± 25	8196	8298	8352
6.940	9250 ± 25	9942	10120	10171

Notes: Radiocarbon dates to supplement those reported by Stoner *et al.* (2013). All dates are on mixed planktonic foraminifera. Calibration for this study uses $\Delta R = 0$ and the Marine13 curve (Reimer *et al.* 2013).

tuning. At present dynamic programming is not widely utilized in PSV stratigraphic correlation.

Of the aforementioned dynamic programming algorithms that address geological time-series data, we chose to adapt that of Hay *et al.* (2019) to align PSV for the following two reasons. First, this algorithm calculates every possible pairing of strata in the target and candidate data sets, as opposed to aligning stratigraphic intervals (two or more strata, e.g. Thompson & Clark 1989; Lisiecki & Lisiecki 2002). Secondly, the easily modifiable penalty functions allow the algorithm to catalog libraries of possible alignments (see discussion of g and $edge$ parameters below). Together, these features facilitate the alignment of data sets characterized by limited age control, as multiple solutions may fit within the available geochronologic constraints. Next, we detail how we adapted the algorithm presented by Hay *et al.* (2019) to minimize the angular (cosine) difference between palaeomagnetic inclination and declination vectors, as opposed to minimizing the global squared difference in the values of isotopic data.

The cost matrix The application of DTW to PSV time-series centers on the assembly of a cost matrix that quantifies the angular difference between every possible alignment between the target and candidate data sets. While traditional formulations of dynamic programming would calculate the squared difference between the inclination or declination of target and candidate data sets individually, the PSV-specific algorithm presented here calculates the cosine distance (CD) between the two angles to align inclination and declination simultaneously.

Each cell of the n by m cost matrix (C)—where n and m reflect the length of the target and candidate data set, respectively—is

populated as follows:

$$C(n, m) = 1 - \cos \theta_{n,m}, \quad (1)$$

where $(\theta_{n,m})$ represents the angular distance between a specific pair of PSV data from the target and candidate data sets. θ is calculated as follows. First, declination and inclination angles are deconstructed to compute the x , y and z directional components of each PSV vector time-series (target and candidate). Next, the code computes the cross product of the target and candidate directional components:

$$t_1 = [x_t \ y_t \ z_t] \times [x_c \ y_c \ z_c], \quad (2)$$

where subscripts t and c refer to the target and candidate components, respectively. Each resulting column is added in quadrature to calculate t_2 :

$$t_2 = \sqrt{t_1(:, 1)^2 + t_1(:, 2)^2 + t_1(:, 3)^2}. \quad (3)$$

The code then computes the dot product of the transposed target and candidate directional components:

$$t_3 = \begin{bmatrix} x_t & x_c \\ y_t & y_c \\ z_t & z_c \end{bmatrix} \cdot \begin{bmatrix} x_c \\ y_c \\ z_c \end{bmatrix}. \quad (4)$$

Finally, the code computes the arctangent of t_2 and t_3 to determine the resulting angular distance:

$$\theta = \text{atan2}(t_2, t_3). \quad (5)$$

The cosine distance (CD) for a given θ (stratal pairing) is calculated as:

$$\text{CD} = 1 - \cos(\theta) \quad (6)$$

and the CD value of every possible target-candidate stratal pair populates the cost matrix (see eq. 1). Linear programming minimizes the sum of these CDs, effectively charting an optimal alignment path through the cost matrix (the ‘best fit’). To catalogue a library of possible alignment solutions, two variables (*edge* and *g*, which augment the cost matrix) are varied to explore the range of possibilities.

The g parameter Values of *g* can reward ($g < 1$) or penalize ($g > 1$) the stretching or squeezing of the candidate data set by weighting the off-diagonal elements of the cost matrix, effectively rewarding or penalizing diagonal travel through the cost matrix. For data sets with equal sampling resolution (e.g. 1 cm resolution in both vector time-series), the *g* parameter effectively rewards or penalizes alignments where the vector time-series have similar accumulation rates. However, for data sets with different sampling resolution, the *g* parameter will reward or penalize alignments where the vector time-series have similar ratios of accumulation rate and input sampling resolution. For illustrative purposes, in these examples of the algorithm application, *g* varies between 0.97 and 1.02, although the extremes of this range are typically greater than what is necessary to achieve stratigraphically plausible alignments (e.g. Hay *et al.* 2019).

The edge parameter The *edge* parameter penalizes solutions that exceed the time bounds of the target data set, with larger values rewarding alignments that have greater overlap between the two data sets. In these examples of the algorithm application, *edge* varies between 0.01 and 0.2 for illustrative purposes, although when *edge* > 0.15 the algorithm tends to converge upon the same alignment solutions (e.g. Hay *et al.* 2019).

Demonstration of dynamic time warping To more concretely illustrate the dynamic time warping alignment process, we consecutively step through the code output for the alignment of two short, synthetic PSV vector time-series (see the light purple columns of the Fig. 2a matrices and light purple rows of the Fig. 2b matrices). First, the code creates two target matrices whose number of rows, *N*, and columns, *M*, equal the number of PSV data in the target and candidate sequences, respectively (Fig. 2a). The one-column target PSV vectors (declination and inclination) are then replicated *M* times to fill all the columns of the two target matrices (Fig. 2a). A similar process generates two candidate matrices, though the candidate PSV vectors are transposed to row vectors and replicated to fill all *N* rows of the candidate matrices (Fig. 2b).

The algorithm then determines θ (eqs 2–5) between all possible candidate-target pairings, populating the theta matrix (Fig. 2c), and from θ calculates the resulting CD (eq. 6) to fill the cost matrix (Fig. 2d). For example, the angular distance between the PSV data (declination and inclination) determined from the third row (equivalent to the 3rd stratum) of the target vector time-series (-85° and 82° , respectively; Fig. 2a) and the last column (youngest stratum) of the candidate vector time-series (-110° and 77° , respectively; Fig. 2b) yields a θ of 6.6553° [theta matrix (3,4), Fig. 2c] and a CD of 0.0067 [cost matrix (3,4), Fig. 2d].

The cost matrix is multiplied by the *edge* parameter (Fig. 2e), which serves to discourage (increased value) or encourage (decreased value) the corresponding stratal pairing. The four cost matrix edges are modified in clockwise order, beginning with the first row, and ending with the first column (Fig. 2e; purple ellipsoids). For cost matrix element (3,4) = 0.0067, an (arbitrary) *edge* value of 0.1 results in a lower value of 0.00067 [edge-modified matrix (3,4), Fig. 2e]. This clockwise implementation modifies the matrix corners twice [once per edge; see edge-modified matrix (1,1), Fig. 2e].

The *edge*-modified cost matrix is variably augmented by *g* as the algorithm computes the cumulative cosine distance matrix (Fig. 2f). The purpose of the *g* parameter is to allow for different sediment accumulation rates between the two vector time-series and to allow for the insertion of hiatus surfaces in either the candidate or target vector time-series. The algorithm begins in the upper left-hand cell of the cumulative cosine distance matrix (1,1), which is not altered from the *edge*-modified matrix. Moving horizontally right along the first row of the matrix, each cell is recalculated as the sum of the corresponding *edge*-modified matrix cell value and the values of all preceding *edge*-modified matrix cells in the column, representing an accumulation of cost (horizontal purple arrow, Fig. 2f). For example, cumulative cosine distance matrix (1,3) is calculated as the sum of *edge*-modified matrix (1,3), *edge*-modified matrix (1,2) and *edge*-modified matrix (1,1), equal to the sum of the values 0.0004, 0.00018 and 0.000038, respectively. This calculation yields 0.00062 (Fig. 2f). The same is done vertically for column 1 of the cumulative cosine distance matrix (vertical purple arrow, Fig. 2f).

The algorithm then moves down to element (2,2) and calculates the accumulation of cost for every cell in column 2 and row 2. In contrast to the calculation of the accumulation of cost above, the value of each of these cells (gold arrows, Fig. 2f) is computed as the sum of a cell’s corresponding *edge*-modified matrix value and the minimum value of the 3 preceding cumulative cosine distance matrix cells [(*n*, *m*−1), (*n*−1, *m*), (*n*−1, *m*−1)]. For example, cumulative cosine distance matrix (6,2) is computed as *edge*-modified matrix (6,2), or 0.0014, plus the minimum of the three preceding cumulative cosine distance matrix cells, which in this case is cumulative cosine distance matrix (5,2), or 0.00202, yielding 0.00342 (Fig. 2f). In this example, we adopted an (arbitrary) *g* value of 1.00; however, for values of *g* greater or lesser than 1 the horizontal and vertical preceding cells are augmented by *g* [*g* × (*n*−1, *m*) and *g* × (*n*, *m*−1)] such that the minimum value of the preceding cell can change.

At this point, the algorithm has computed values for all cells in rows 1 and 2, and columns 1 and 2, in the cumulative cosine distance matrix. Values for the remaining 10 empty cells (rows 3–7 and columns 3–4, grey arrows in Fig. 2f) are calculated as the sum of the corresponding *edge*-modified matrix and the minimum of the preceding two rows and columns [(*n*−1, *m*), (*n*, *m*−1), (*n*−1, *m*−1), (*n*−2, *m*−1), (*n*−1, *m*−2), (*n*−2, *m*), (*n*, *m*−2), (*n*−2, *m*−2)] of the cumulative cosine distance matrix. These 8 preceding cells are augmented by *g* and a cell-position-dependent coefficient as follows: elements (*n*−1, *m*) and (*n*, *m*−1) are computed as *g* × *g*; elements (*n*−2, *m*−1) and (*n*−1, *m*−2) are computed as 1.05*g* × *g*; elements (*n*−2, *m*) and (*n*, *m*−2) are computed as 1.1*g* × *g*. Note that diagonal preceding cells [(*n*−1, *m*−1) and (*n*−2, *m*−2)] are not augmented by *g*. For example, cumulative cosine distance matrix (3,4) is computed as *edge*-modified matrix (3,4), or 0.00067, plus the minimum preceding cell value, in this case cumulative cosine distance matrix (1,2), equal to 0.00022. Due to its position (*n*−2, *m*−2), this cell’s value is not augmented, yielding a value of 0.00089 for cumulative cosine distance matrix (3,4) (Fig. 2f). Once all of the remaining empty cells have been filled, the cumulative cosine distance matrix is complete (Fig. 2f).

An alignment between the candidate and target time-series is derived by determining the path through the cumulative cosine distance matrix that accumulates the least cost. The algorithm begins in the lower right-hand corner of the cumulative cosine distance matrix (lighter purple box in Fig. 2g, replicated from Fig. 2f) and selects each forward step from the lowest value of the adjacent eight cells ahead to ensure an optimal path (Fig. 2g). The alignment path begins when it enters the matrix (diverges from a starting edge; purple arrows in Fig. 2g) and ends when it exits the matrix (meets

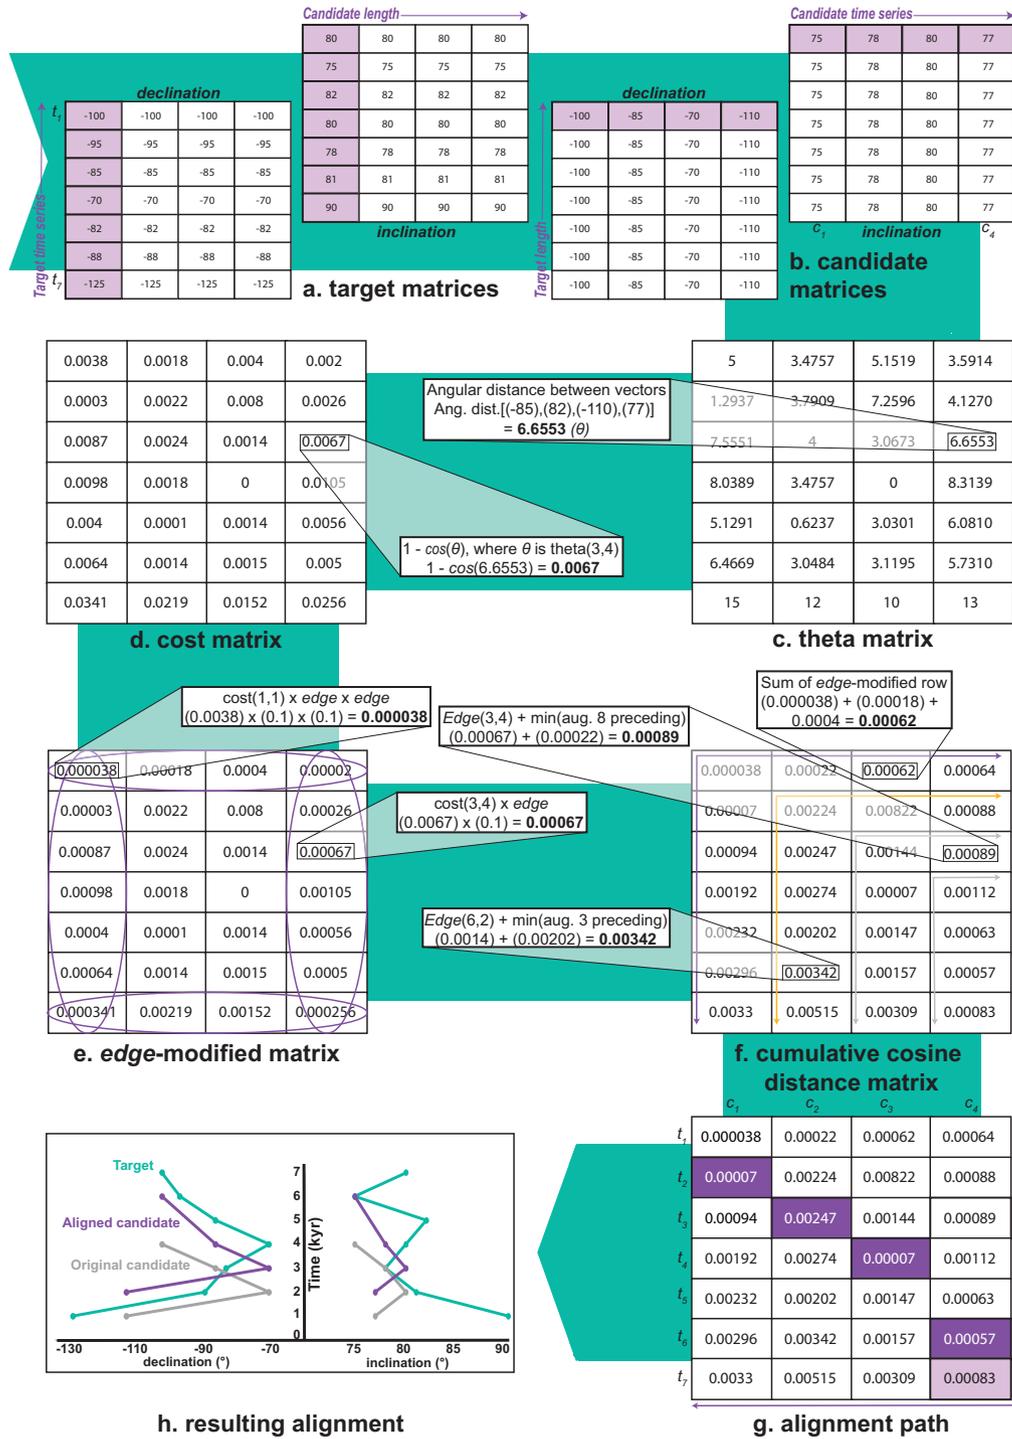


Figure 2. Demonstration of PSV dynamic time warping. We use synthetic data to walk through each calculation involved in this dynamic time warping algorithm to align PSV data. (a) Target data set matrices, with declination on the left and inclination on the right. The synthetic target data sets themselves are shown in the light purple first columns of the matrices. (b) Candidate data set matrices, with declination on the left and inclination on the right. The synthetic candidate data sets themselves are transposed and shown in the light purple first row of the matrices. Note that the matrices in a and b are of equal size, with the number of rows defined by the length of the target data set and the number of columns defined by the length of the candidate data set. (c) Theta matrix, in which each cell is filled with the angular distance between the corresponding target and candidate vectors from a and b. (d) Cost matrix, in which each cell is filled with the cosine distance as calculated with θ from c. (e) Edge-modified matrix, in which the edges of d are modified by *edge* in clockwise order (0.1 in this example; see purple ellipses; note matrix corners are modified twice). (f) Cumulative cosine distance matrix, in which the cells of e are modified in one of three ways: summed along matrix edges (purple arrows); summed with the minimum possible preceding step (gold arrows); summed with the minimum of the two possible preceding steps (grey arrows). Calculations begin in the upper left corner of the matrix (1,1), and thus this cell is unmodified from e. (g) the alignment path through the resulting augmented cost matrix (f), which begins in the lower right corner (7,4; light purple box) and works its way up and left (see purple arrows), optimized by accruing the least possible cost. (h) the resulting target-candidate alignment for the declination (left-hand panel) and inclination (right-hand panel) records. The target data set is shown in teal, the aligned candidate data set is shown in purple, and the original candidate is shown in grey.

an ending edge; left and top edges in Fig. 2g). The resulting path represents the alignment itself, indicating where each successive candidate vector stratum (PSV sample) corresponds best with the target vector stratum (Fig. 2h).

2.2.2 Goodness of fit

Linear goodness of fit parameters, such as R^2 values, are not applicable for comparing two vector time-series. For example, calculating an R^2 value between two declination time-series can be problematic as the true angular distance between two depends on inclination (e.g. large declination differences at steep inclinations can actually be small angular differences). Accordingly, this algorithm uses a goodness of fit cross correlation parameter (XC) that is calculated using the same principle as an R^2 —normalizing the residual from a model by a metric of variance.

The algorithm transforms target and candidate inclinations and declinations into their x , y and z vector components to calculate a linear model, akin to the linear models used in typical calculations of R^2 , for each variable (Fig. S1). The residual cosine distances (RCD) are calculated as the sum of the CDs, where θ is measured between the modelled and original candidate vectors. The total cosine distance (TCD) is calculated as the sum of the CDs, where θ is measured between the candidate vectors and the mean of the candidate vectors. Using RCD and TCD, XC can be calculated as follows:

$$XC = 1 - \frac{RCD}{TCD}. \quad (7)$$

2.2.3 Statistical tests

While the example PSV vector time-series we use in this discussion have radiocarbon chronologies that can be used to assess the accuracy of the resulting alignments, the vector time-series for which this algorithm may be the most useful will likely not have independent chronostratigraphic constraints. This algorithm employs a Monte Carlo method (the Markov Chain Monte Carlo statistical test of Haam & Huybers 2010), adapted for PSV data, to assess the statistical significance of resulting alignments. To do this, this algorithm warps 10 000 synthetic candidate vector time-series to best match the target vector time-series across a range of g values (0.98–1.01) with a fixed *edge* value (*edge* = 0.2 to force overlap). To assess statistical significance, the algorithm compares the average XC for the original candidate alignment across all *edge* values for a given g value to the XC distribution from the 10 000 synthetic candidate vector time-series. The statistical significance test investigates if the resulting alignment between the candidate data set and the target data set produces a better XC than could be achieved through random chance ($p < 0.05$; rejection of null hypothesis).

The synthetic PSV sequences are constructed to approximate the vector mean, variance, and frequency characteristics of the candidate data set. The x , y and z vector components of the candidate data set are calculated from the input inclination and declination. These vector components are then centered to a zero mean and normalized by their standard deviation. Synthetic vector components are generated using Cholesky's decomposition technique (Haam & Huybers 2010; Hay *et al.* 2019) and transformed to match the mean and variance of the input vector components. Inclination and declination are recalculated from these synthetic x , y and z vector components.

The alignment library can be narrowed by eliminating solutions deemed statistically insignificant. However, a failure to reject the null hypothesis (alignment no better than random chance) does not necessarily indicate that the particular alignment is incorrect. Localized basin-scale fluctuations in sedimentation, magnetic lock-in depth, or diagenesis could allow for sequences that would otherwise match well to produce high p -values. Therefore, we only advocate for the elimination of an alignment based on statistical significance if there is another alignment in the library that is found to be significant ($p < 0.05$).

3 RESULTS AND DISCUSSION

3.1 Ideal data: aligning MD99-2322 to MD99-2269

In our first example, we discuss the resulting alignment of MD99–2322 (candidate) and MD99–2269 (target), which we emphasize required no manual tuning. We consider these ideal data for the following reasons:

(i) Geographically, MD99–2322 and MD99–2269 are very close to one another (<500 km), making a strong case for the assumption of a common geomagnetic field history.

(ii) MD99–2269 and MD99–2322 have already been shown to record high fidelity PSV via supervised correlations consistent with the cores' radiocarbon constraints (Stoner *et al.* 2007) and comparison to other records on regional and hemispheric spatial scales (Ólafsdóttir *et al.* 2013, 2019; Stoner *et al.* 2013; Walczak *et al.* 2017; Reilly *et al.* 2019).

(iii) Both vector time-series are high resolution and have similarly high accumulation rates, meaning the alignment may be assisted by the g parameter reward.

(iv) Both MD99–2322 and MD99–2269 span roughly the same period of time (~11.5 ka), meaning the alignment may be assisted by the *edge* parameter reward.

3.1.1 A library of alignments

Rather than find a single, mathematically optimal alignment path, this algorithm generates a library of possible alignments arising from pairings of g and *edge* values. The impact of various g values on alignment of MD99–2322 to MD99–2269 PSV data can be illustrated by plotting the accumulated cost matrices (Fig. 3). Visually, areas of the cost matrix with low CDs (blues in Fig. 3) reflect good alignment between the cores such that warp paths through these zones in the resulting accumulated cost matrices yield low accumulated cost. Travel in various directions through the cost matrix is weighted by g , with higher g values penalizing travel to lateral positions and, accordingly, rewarding travel in a diagonal direction. This is illustrated in Fig. 3, particularly for high g values ($g = 1.01$) where regions of low accumulated cost form parallel to the accumulated cost matrix diagonal. In this example, an intermediate g value (1.00) provides the best visual fit between the vector time-series, finding balance between the g penalty/reward and the distribution of CDs across the cost matrix.

Comparing the warp paths (3978 solutions) from a wide range of g (0.97–1.02) and *edge* (0.01–0.2) values illustrates the tradeoff between g and the cost matrix (Fig. 4). For very low g values (<~1), lateral travel through the accumulated cost matrix is rewarded and warp paths invoke large hiatuses and/or very rapid accumulation, hence the stair-step character of these paths (purples in Fig. 4). Vertical travel through the accumulated cost matrix instead invokes

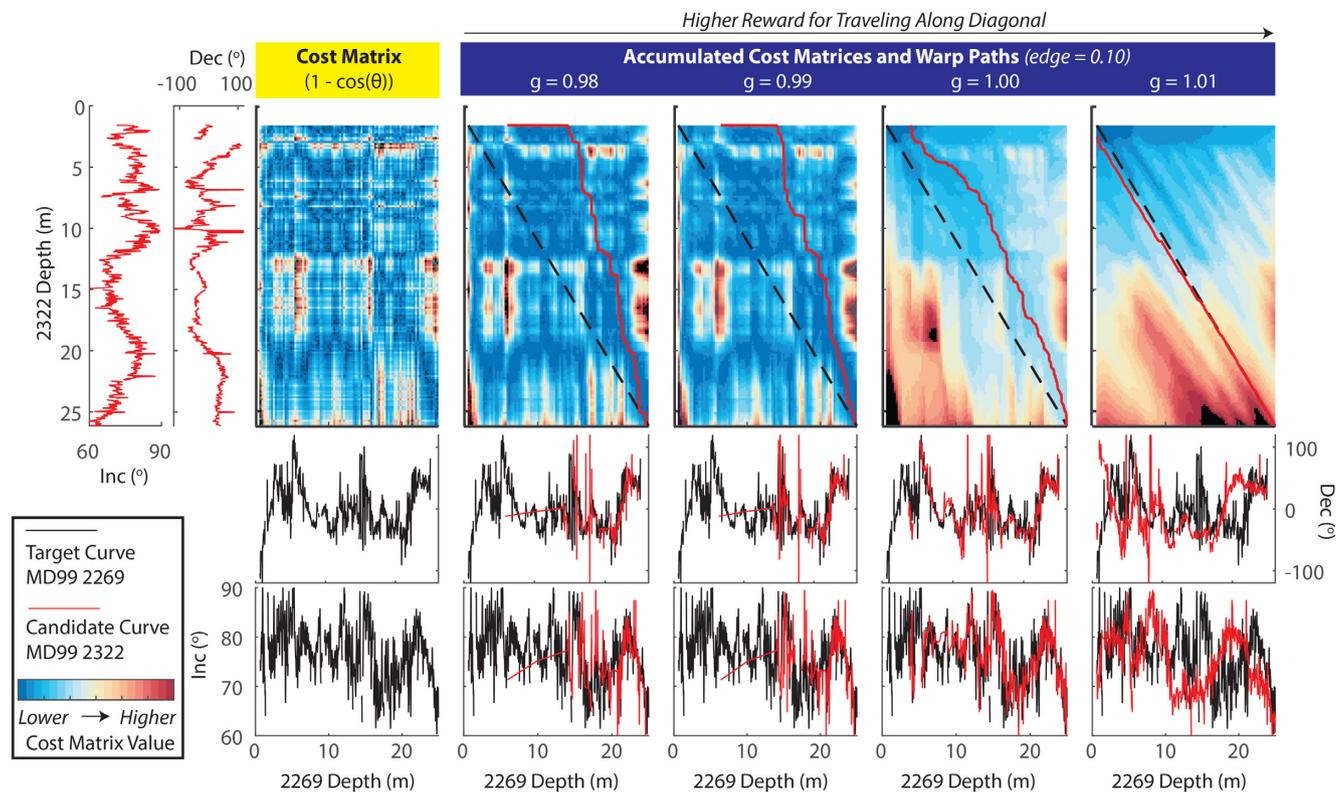


Figure 3. Example cost matrix, accumulated cost matrices, and warp paths. The CD [$1 - \cos(\theta)$] is calculated for each paired data point between the target curve (MD99–2269; black) and candidate curve (MD99–2322; red) to generate the cost matrix. Accumulated cost matrices are calculated from the cost matrix and vary based on the user's choice of g and $edge$. In all cases, blues/whites indicate lower values and reds/blacks indicate higher values. Higher g values reward warp paths (red lines) parallel to the cost matrix diagonal (black dashed line). Higher $edge$ values reward warp paths that begin/end at the edges of the cost matrix. Example warp paths for g ranging from 0.98 to 1.01 and $edge = 0.10$ are included above the resulting alignment of the candidate vector time-series (red) to the target vector time-series (black).

a hiatus in the candidate data set. The low g value warp paths are also very sensitive to the choice of $edge$ values, with low $edge$ values resulting in reduced overlap between the two vector time-series (reds in Fig. 4). At very high g values ($> \sim 1$), the warp paths follow the cost matrix diagonal and the warp paths become less sensitive to the choice of $edge$ parameter (greens and blues in Fig. 4). Overlap between the two vector time-series is generally lowest when g and $edge$ are small, and highest when g and $edge$ are large (Fig. 5b).

3.1.2 Assessing the goodness of fit of the alignment

In situations where the target and candidate sequences include no independent stratigraphic constraints, the user will need to evaluate the library of alignment paths to choose the preferred, or groups of preferred, solution(s) using statistical likelihood, geological context, and expert knowledge. For the alignment of Core 2322 to Core 2269, however, radiocarbon constraints facilitate the development of a depth–depth relationship and associated confidence interval through comparison of their Undatable (Lougheed & Obrochta 2019) age–depth models (Figs 4 and 5g).

The highest XCs between the target and warped candidate PSV vector time-series occurs at low g and $edge$ values (Fig. 5a) where overlap between the alignments is low (Fig. 5b), suggesting the high XCs are an artifact of using a small subset of the available data. Considering only alignment paths with an overlap of at least 500 data points, the highest XCs fall around $g = 1$ in a band that extends the entire range of $edge$ values, suggesting the best

mathematical solutions are on first-order insensitive to the choice of $edge$ parameter (Fig. 5a). This is likely in part because Cores 2269 and 2322 span the same time interval and thus have a good solution that does not require the reward provided by $edge$.

We evaluate the library of alignment paths against the cores' independent radiocarbon stratigraphies by comparing the absolute difference between the cores' median age–depth model, assessing the standard deviation of the difference between the cores' median age–depth model, and determining the per cent of data that fall within the ± 1 and the $\pm 2\sigma$ ranges. In each case, alignments with $g \sim 1$ across a wide range of $edge$ values provide the best alignment (Figs 5c–f). For alignments with $g = 1$ and $edge \geq 0.02$, on average 88 per cent of data fall within the $\pm 1\sigma$ range and 99 per cent of data fall with the $\pm 2\sigma$ range. This suggests that if the PSV alignment were the 'true' depth–depth alignment between these cores, we are overestimating the uncertainty structure of the radiocarbon alignment by using a conservative xfactor of 1 in the Undatable age-model (Lougheed & Obrochta 2019). The highest XC (0.863) was found at $g = 0.999$ and $edge = 0.08$ (Figs 5g and h); however, there are a number of other alignments that give essentially the same result (Fig. 5g).

3.1.3 Statistical test of significance

Each of the 10 000 synthetic time-series generated using the characteristics of the input Core 2322 (candidate) vector time-series (four examples in Fig. 6b) are warped to the target (Core 2269) vector time-series with an $edge = 0.2$ to ensure overlap. The Fisher

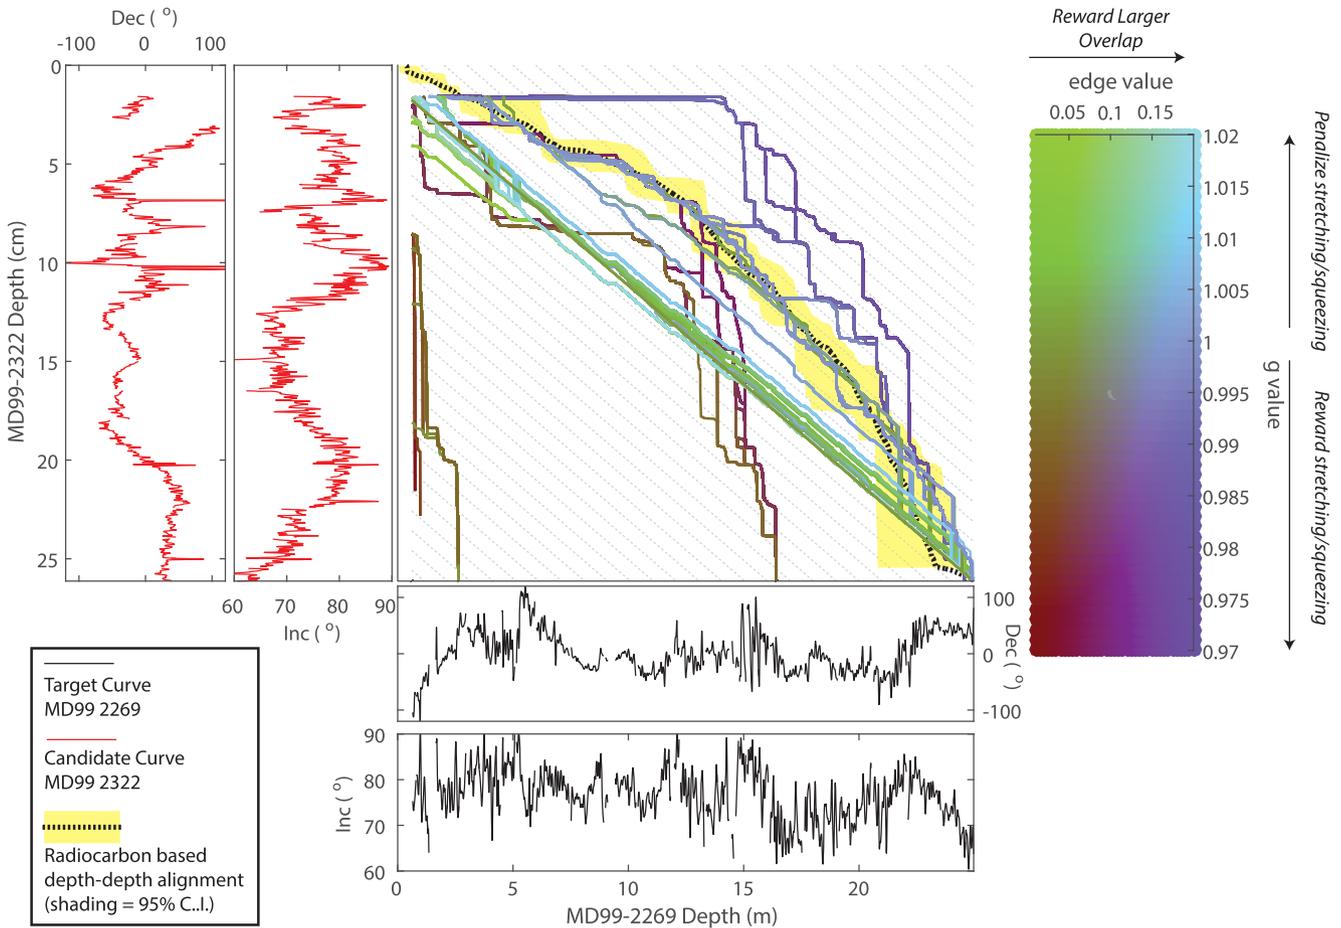


Figure 4. Library of warp paths given a wide range of g and $edge$ values. 3978 alignments of the candidate curve (MD99–2322; red) to the target curve (MD99–2269; black) for a wide range of g and $edge$ values (various color lines correspond to color matrix at right). Some g and $edge$ pairs yield identical results and depth–depth relationships plot on top of each other. A depth–depth alignment using the independent radiocarbon chronologies for the candidate and target vector time-series is included (black dashed line) with its 95 per cent confidence interval (yellow shading). Paths parallel to the diagonal of the cost matrix are indicated (light grey dashed lines). Note, higher g values typically follow paths parallel to the diagonal of the cost matrix while lower g values tend to avoid paths parallel to the diagonal of the cost matrix. Intermediate g values (~ 1) perform best when compared to the radiocarbon-based depth–depth alignment.

mean of these synthetic vector time-series generally are within a few degrees of the Fisher mean of the input vector time-series and have similar circular standard deviations (Fisher 1953). The warped synthetic XC distributions (blue probability distribution functions, Fig. 6c) are compared to the XCs of the Core 2322 warped paths (red histograms, Fig. 6c). Typically, when warping the synthetic data sets, $g \leq 1$ yield reasonable XCs, presumably because the synthetic data sets can be stretched and squeezed to match the target data set. Conversely, when $g > 1$, the amount of stretching and squeezing is limited and XCs are generally low. Using $p \leq 0.05$ as a cut-off for significance, we find values of g between 0.993 and 1 result in statistically significant alignments (not considering solutions with overlap less than 500 data points; Fig. 6d). Alignments within this g range were also the alignments found to be largely within the 2σ age uncertainty envelope derived from the independent radiocarbon age model (red horizontal band in Fig. 5f).

3.2 Non-ideal data: warping IODP Site U1305 to align with MD99-2269

We perform a similar experiment to align PSV data from Site U1305 to MD99–2269, with no input from the palaeomagnetist. While Site U1305 has been demonstrated to have well-resolved PSV, it is

characterized by a few factors that may challenge this algorithm and PSV stratigraphic correlation, including:

- (i) Geographically, Site U1305 is more distal to MD99-2269 (~ 1800 km) than MD99–2322. This distance can partially be accounted for through the relocation of the record via its VGP path to a common location as has been done in other PSV stratigraphy studies (e.g. Korte *et al.* 2018, Reilly *et al.* 2018, Ólafsdóttir *et al.* 2019). The difference between the relocated (to the location of Core MD99–2269) and original data is shown in Fig. 7 (dark red is relocated; light red is original).
- (ii) Accumulation rates in MD99-2269 are a factor of 2–6 greater than those in Site U1305, which means that the g value reward for travel along the cost matrix diagonal may not be helpful in finding the correct solution with 1:1 sampling resolution.
- (iii) Differences in accumulation rates likely result in different magnetic age offsets as a result of depth dependent pDRM processes, with the high-resolution MD99–2269 core having little age offset and moderate resolution Site U1305 having offsets on the order of a few hundred years (see Stoner *et al.*, 2013).
- (iv) The interval of Site U1305 we use here, as studied by Stoner *et al.* (2013), is shorter than the interval covered by MD99–2269 (~ 8.5 ka versus ~ 11.5 ka), meaning that increasing the $edge$ parameter may not be helpful in finding the correct solution.

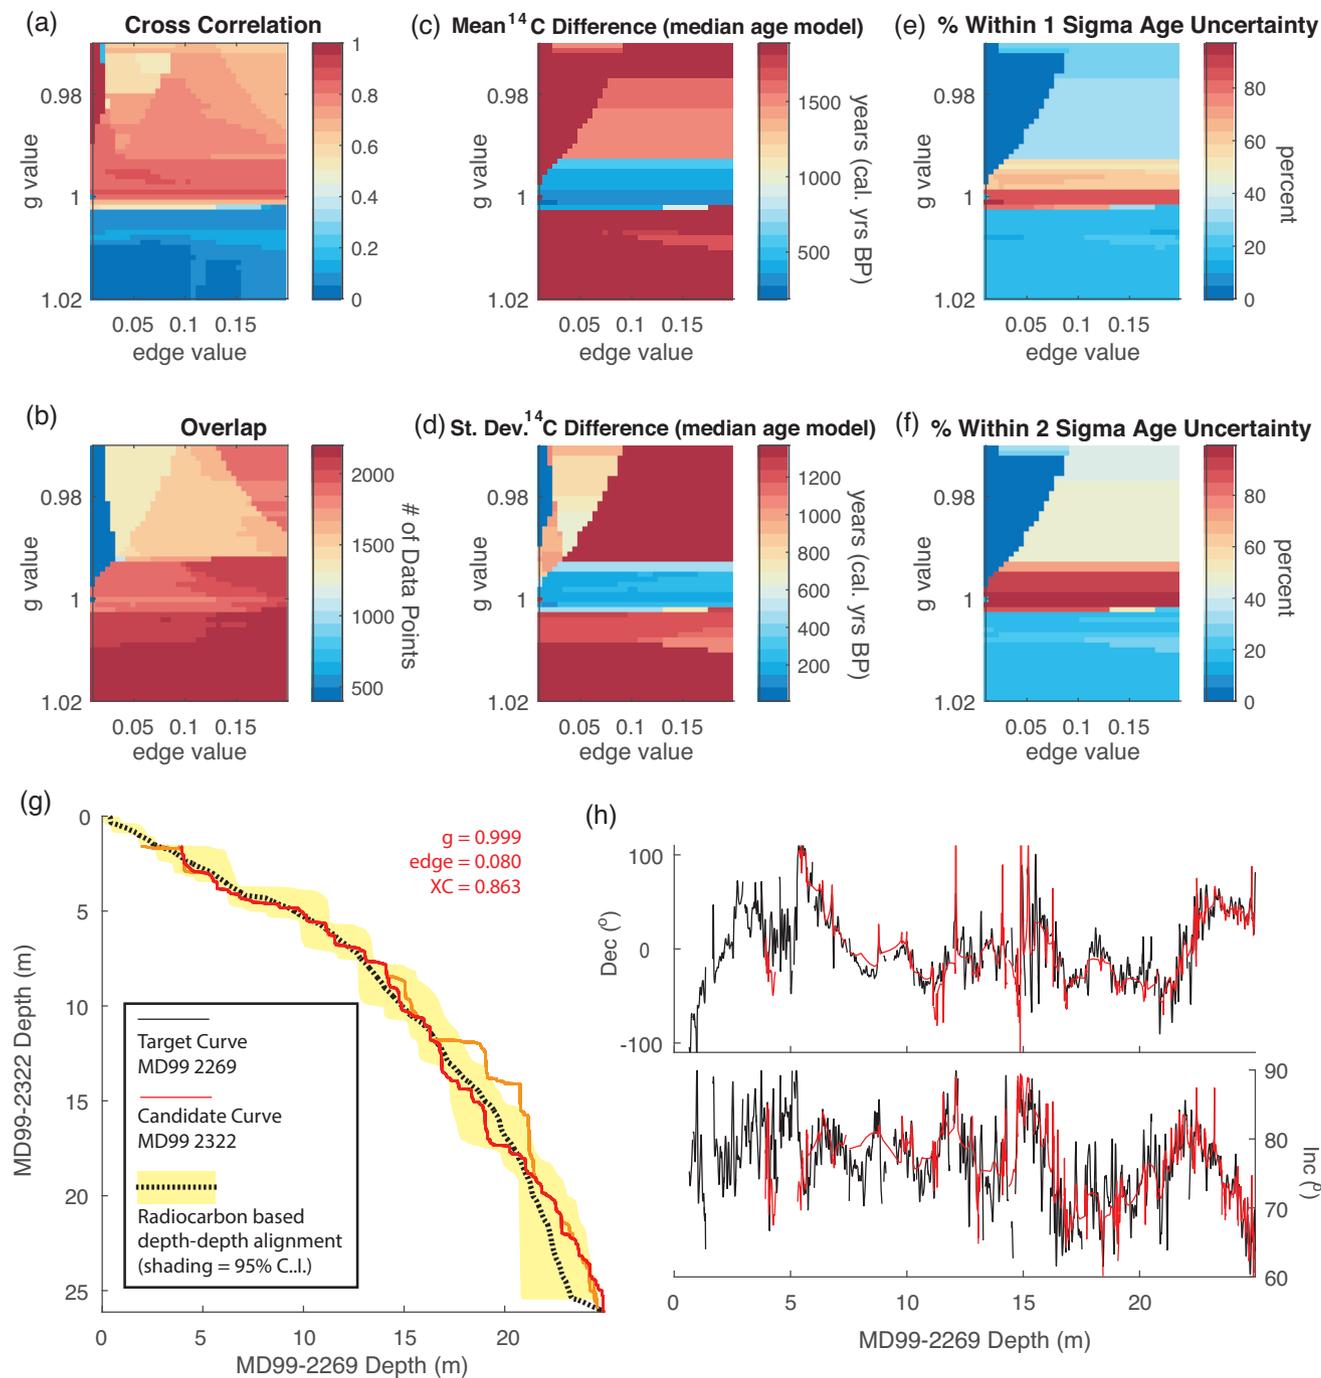


Figure 5. Evaluating the performance of the library of alignment paths given a wide range of g and $edge$ values. In panels a–f we investigate $g = 0.97–1.02$ and $edge = 0.01–0.2$ (a) XC of the candidate (MD99–2322) and target (MD99–2269) vector time-series. (b) Number of overlapping data points in each alignment path. Note, that regions with little overlap (low $edge$) may have anomalously high XC in a. (c) Mean difference between the two cores’ independent radiocarbon age models at correlated points using the median modeled age. (d) Standard deviation of the difference between the two cores’ independent radiocarbon age models at correlated points using the median modeled age. (e) Percent of correlated points that are within the $\pm 1\sigma$ age uncertainty of the two cores’ independent radiocarbon age models. (f) As in e, but for the $\pm 2\sigma$ age uncertainty. (g) Plot of the highest performing alignment paths (XCs > 0.85 and overlaps > 500 data points) in orange with best performing (highest XC) alignment path in red. The radiocarbon-based depth–depth alignment (black dashed line) and 95 per cent confidence interval (yellow shading) are also indicated. (h) PSV comparison of the best performing alignment path (target vector time-series in black; warped candidate vector time-series in red).

3.2.1 Sensitivity of g to sampling resolution

Unlike the previous example, the thickness of Holocene aged sediments are considerably different between Site U1305 and Core 2269. Accordingly, the length of the data sets, at their original 1 cm

sampling resolution are different, result in cost matrix dimensions that are more elongate than the previous example. High values of g reward accumulation rates at Site U1305 on par with those observed at Core 2269, which are much greater than would be expected if our

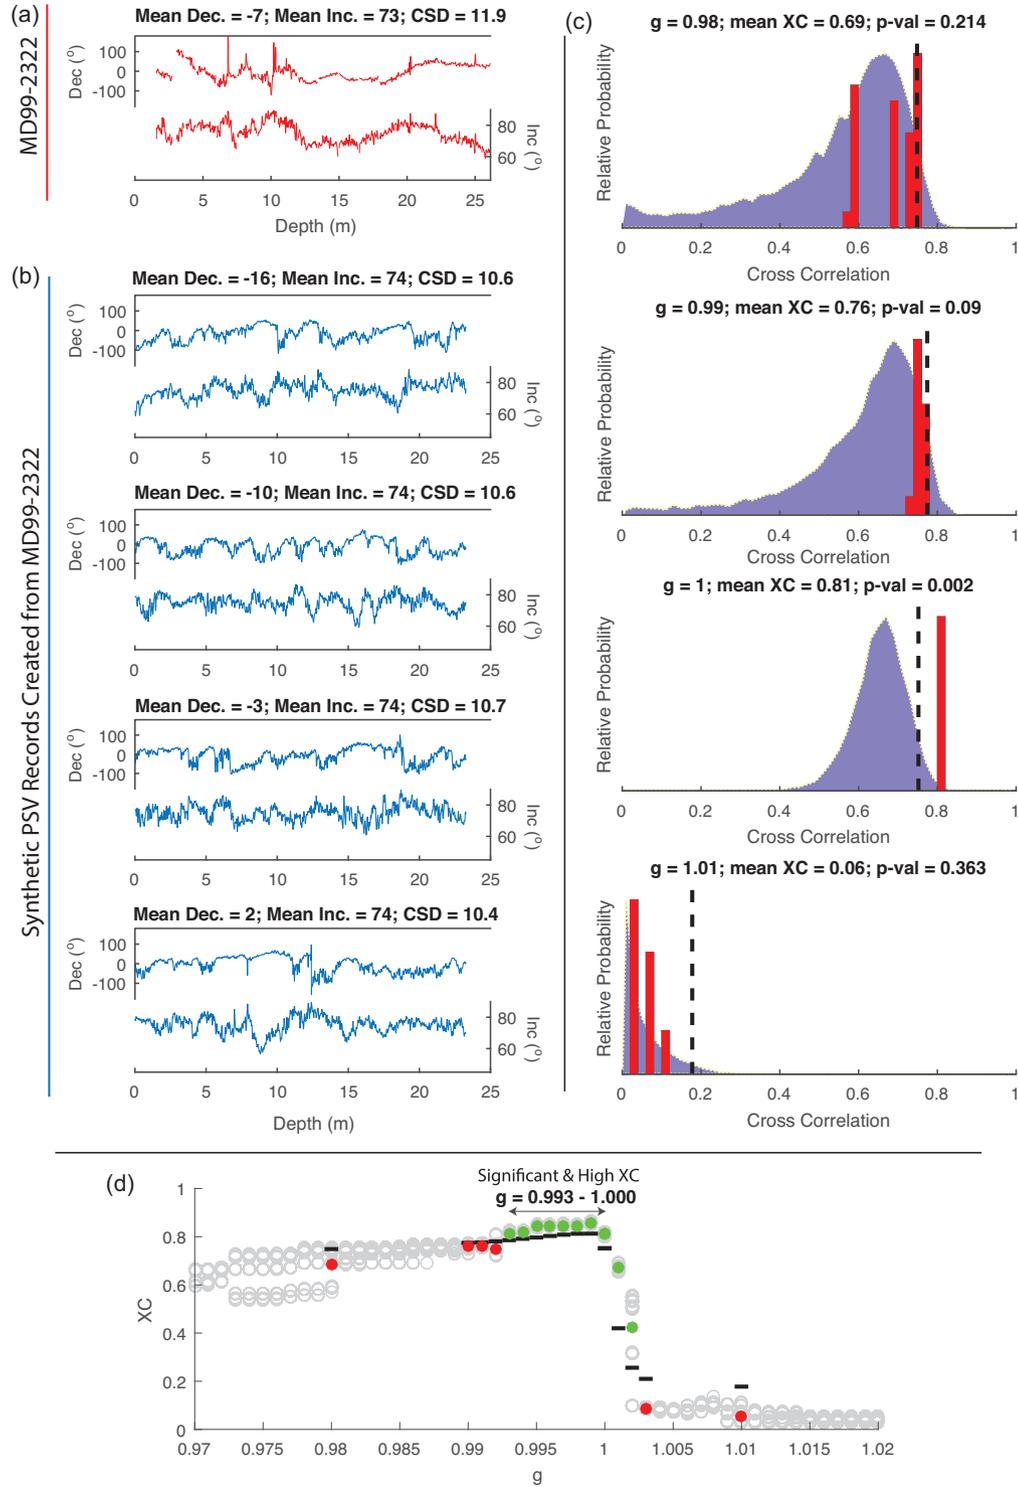


Figure 6. Example synthetic PSV vector time-series and statistical significance tests. (a) MD99–2269 PSV vector time-series, with vector mean declination/inclination and the circular standard deviation (CSD) indicated. (b) Four random example synthetic PSV vector time-series that have frequency characteristics, mean, and variance in their x, y and z vector components based on the MD99–2322 PSV vector time-series. (c) Summary of Markov Chain Monte Carlo significance tests, plotting the XC distributions (blue shading) from the alignment of 10 000 candidate synthetic PSV vector time-series (based on MD99–2322) to the target MD99–2269 PSV vector time-series for four g values with the 95 per cent level indicated (vertical black dashed line). Actual mean XC values (for a given g) between alignments of MD99–2322 and MD99–2269 (with overlap >500 data points and at various *edge* values) are also indicated (red histogram). Note, $g = 1$ consistently generates highly significant alignments where the data alignments are a better match than can be achieved by random chance 99.8 per cent of the time. (d) Significance tests (as in c) highlighting the range of statistically significant alignments at the 95 per cent level that also achieve high XC. Grey open circles are the XC of the MD99–2322 and MD99–2269 alignments (with overlap >500 data points), green filled circles are the mean of those values if significant at the 95 per cent level, red filled circles are the mean of those values if not significant at the 95 per cent level, black horizontal line indicate the 95 per cent level of the synthetic alignments.

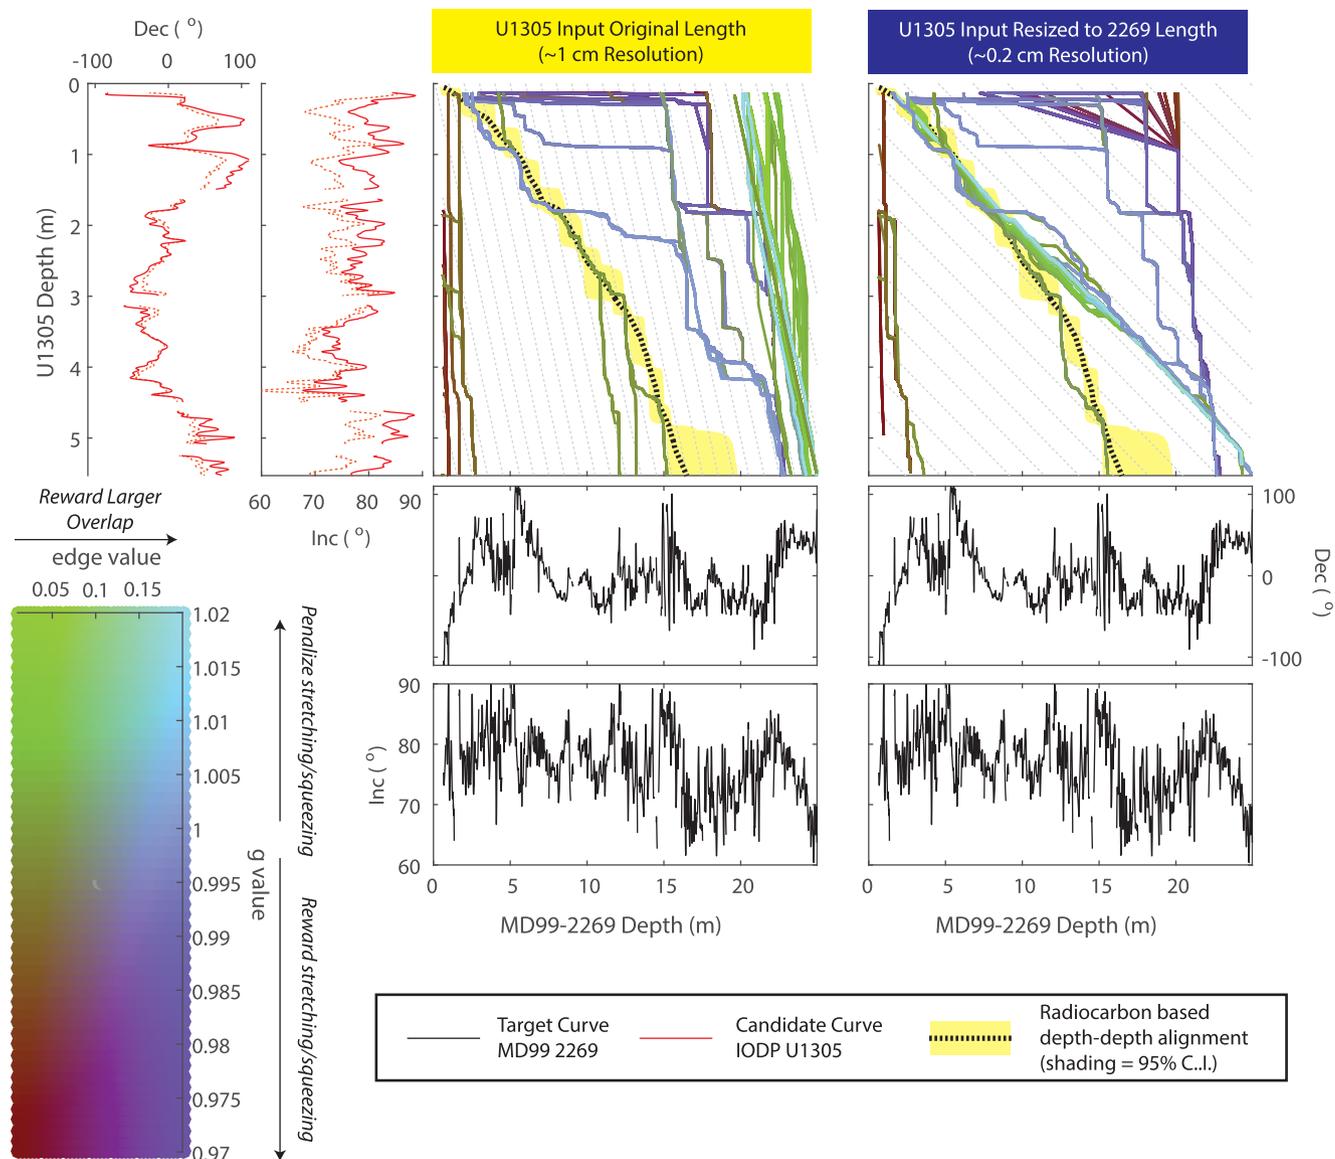


Figure 7. Site U1305 warp paths and sensitivity to cost matrix diagonal. Alignment of the relocated candidate curve (Site U1305; relocated data in red solid line—primary data in light red stippled line) to the target curve (MD99–2269; black) for a wide range of g and $edge$ values. A depth–depth alignment using the independent radiocarbon chronologies for the candidate and target vector time-series is included (black dashed line) with its 95 per cent confidence interval (yellow shading). Alignment was done using Site U1305 at its original data measurement resolution (~ 1 cm; plots on left under yellow box) and after interpolation to match the length of the much higher resolution MD99–2269 input (~ 0.2 cm; plots on right under blue box). In both scenarios, the resolution of the MD99–2269 input is ~ 1 cm. Paths parallel to the diagonal of the respective cost matrix are indicated (light grey dashed lines). Note, higher g values in first scenario reward alignment paths that equate to similar accumulation rates, while higher g values in the second scenario reward alignment paths that equate to the two vector time-series spanning similar duration.

assumption was that each core spans the majority of the Holocene epoch. Thus, we can change our assumption of the two cores having roughly similar accumulation rates to having the proportional sedimentation rates implied by the records spanning roughly the same interval of time by rescaling the lower accumulation rate candidate vector time-series (Site U1305) to the length of the higher accumulation rate vector time-series (MD99–2269), or vice versa (cost matrix diagonal lines are grey dashed lines in Fig. 7). The result of this change is clear in the alignment paths (at high g values) when the input resolution of the Site U1305 data set is adjusted (Fig. 7). While this has a limited impact on the alignment paths for the best performing solutions in the upper 3 m of Site U1305, the assumption of different sedimentation rates (versus equal accumulation rates)

performs better below when compared to the independent radiocarbon depth–depth relationship (Fig. 7). This likely indicates that the correlation is more robust in the upper part of the stratigraphy (less sensitive to cost matrix size) and less robust below 3 m, requiring proper assumptions to arrive at a plausible solution. As in the previous example, lower g values typically result in alignments that imply intervals with longer hiatuses and/or rapid accumulation (Fig. 7).

3.2.2 Assessing the performance of the time warping

Like the last example, high XCs for both the original and interpolated (resized) lengths are found at low g and $edge$ values, where there is less overlap between the two data sets (see Figs 8a, b, g

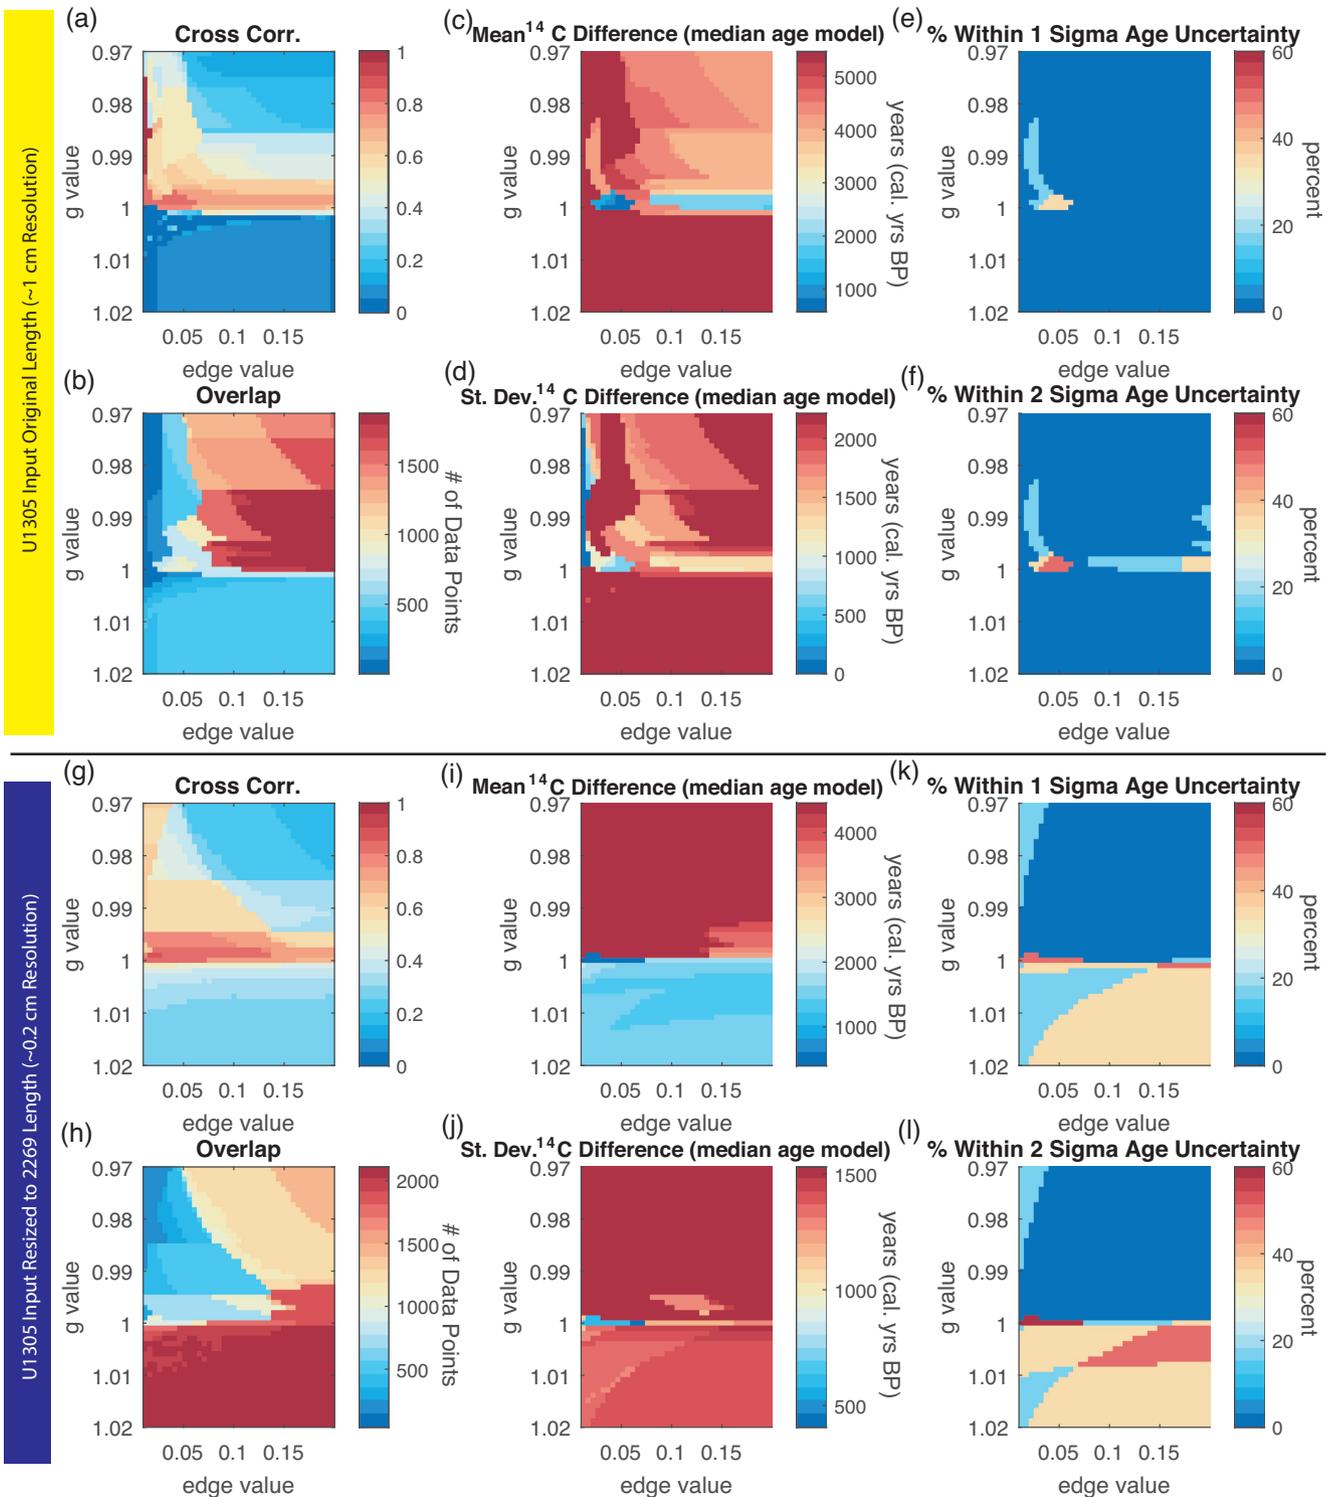


Figure 8. Evaluating the performance of the library of Site U1305 warp paths given a wide range of g and $edge$ values and length of input data. (a–f) Results using the Site U1305 original measurement resolution (~ 1 cm). (g–l) Results after resizing Site U1305 to match the length of the much higher resolution MD99–2269 vector time-series (~ 0.2 cm resolution). (a, g) XC of the candidate (Site U1305) and target (MD99–2269) vector time-series. (b, h) Number of data points that overlap in each alignment path. Note, that regions with little overlap at low $edge$ may have anomalously high XC in a. (c, i) Mean difference between the two cores’ independent radiocarbon age models at correlated points using the median modeled age. (d, j) Standard deviation of the difference between the two cores’ independent radiocarbon age models at correlated points using the median modeled age. (e, k) Percent of correlated points that are within the $\pm 1\sigma$ age uncertainty of the two cores’ independent radiocarbon age models. (f, l) As in e, but for the $\pm 2\sigma$ age uncertainty. Note for e, f, k and l, a large number of solutions lie almost entirely outside the $\pm 1\sigma$ and $\pm 2\sigma$ range (dark blue).

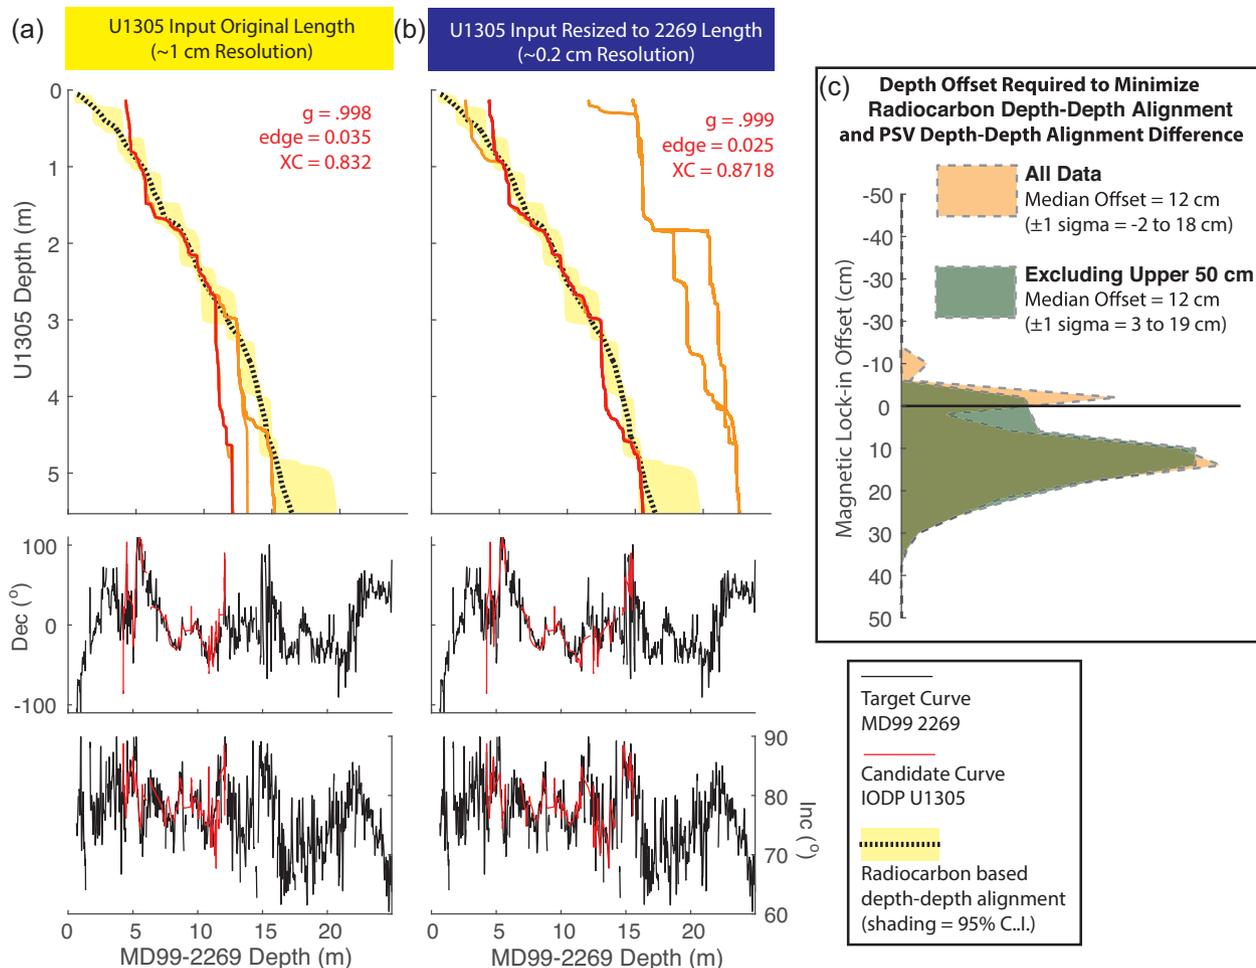


Figure 9. Comparing the highest performing results ($XC > 0.75$) for Site U1305 and evaluation of offset between the radiocarbon and PSV alignments. (a–b) Comparison of the highest XCs (> 0.75) for alignment paths with sufficient overlap (> 500 data points) for the Site U1305 input at the primary resolution (a) and resized resolution (b) (orange lines). The highest XC alignment path is red and plotted below the depth–depth plot. Other alignment paths are plotted in grey. The radiocarbon alignment is included (black dash line) with 95 per cent confidence (yellow shading). (c) Analysis of the offset between the radiocarbon alignment and the highest XC warped PSV alignment in (b). Probability distribution functions indicate the depth offsets required to minimize the difference between the PSV warped path and 10 000 random radiocarbon alignments based on output from each core’s Undateable age–depth model. This exercise was done with all the data (orange shading), and excluding the upper 50 cm (green shading).

and h). The highest XC values with > 500 aligned data points generally arise from $g \sim 1$. In this case study, the highest XCs with sufficient overlap (> 500 data points) arise from $edge$ values $< \sim 0.1$ as is expected considering the contrasting duration of these vector time-series. The structure of XCs at various g and $edge$ values varies between the two scenarios for original versus resized input length; however, regions of higher XCs and sufficient overlap generally coincide with better performance against the independent radiocarbon stratigraphy (Figs 8c–f, i–l).

Of the warp paths arising from taking the highest XC values with sufficient overlap (> 500 data points), we find little difference between the warped paths in the upper 300 cm, but a range of warped paths below (Fig. 9). For the scenario where we use equivalent sampling resolution (i.e. ~ 1 cm for both the target and the candidate), most of the paths follow the diagonal of the cost matrix. This could be because the higher amplitude PSV variability in the upper 300 cm creates a unique cost matrix structure such that accumulated cost is less sensitive to g . In comparison, lower amplitude PSV variability below 300 cm creates a less-defined cost matrix

structure that is more sensitive to g (Fig. 9a). When the input candidate data are rescaled proportionally to the approximate difference in accumulation rate, more warped paths converge on the independent radiocarbon stratigraphy (Fig. 9b).

3.2.3 Evaluating the offset between PSV and radiocarbon stratigraphies

While the preferred Site U1305 and MD99–2269 alignments do a good job agreeing with the independent radiocarbon depth–depth alignment, the magnetic correlation is often biased towards slightly deeper depths in Site U1305 or shallower depths in MD99–2269 (Fig. 9b). This relationship was previously observed when the PSV vector time-series were correlated by Stoner *et al.* (2013), who found that their PSV tie points had depth offsets ranging from 13 to 25 cm. The age offsets associated with these depth offsets were calculated to range from 170 to 401 yr. We find, when comparing the highest XC alignment solution (red line in Fig. 9b) to depth–depth relationships created from random pairs of iteration of the radiocarbon-based

age–depth Undatable results, constant depth offsets of ~ 12 cm, with $\pm 1\sigma$ range of -2 to 18 cm (or 3 – 19 cm, if the upper 50 cm are excluded, Fig. 9c).

This offset is consistent with previous depth offsets for sediment magnetizations in deep sea sediments (Channell & Guyodo 2004, Suganuma *et al.* 2010, Simon *et al.* 2018), which have been attributed to magnetic acquisition in a lock-in zone, following deposition, as a pDRM. It is possible that some of this offset could also be attributed to older reservoir ages of the mixed planktonic foraminifera that were dated in Site U1305 than the mollusk, gastropod and benthic foraminifera material dated in MD99–2269; however, we consider this unlikely as it would require consistent reservoir ages on the order of a few hundred years older at Site U1305 for the Holocene than what would be typically considered (Bard 1988, Reimer *et al.* 2002).

4 PERSPECTIVE AND CONCLUSION

This algorithm can successfully align Holocene PSV data sets and elucidate multiple possible alignments within their radiocarbon uncertainty envelopes without any manual tweaking. We selected the case studies of the alignment of MD99–2269 (target) with candidate sections MD99–2232 and IODP Site U1305 to demonstrate the algorithm's skill against a tightly constrained independent technique (radiocarbon chronology). These alignments are reproducible and objective, eliminating the need for tedious visual alignments, circumventing potential researcher bias and allowing workers to more easily compare solutions with statistical significance. In our first experiment we demonstrate how this algorithm can be used to objectively align high resolution PSV vector time-series (MD99–2322 and MD99–2269) without utilizing the independent age data available. We then test the algorithm with less ideal PSV data (IODP Site U1305) and find that the algorithm still finds alignments within the independent radiocarbon stratigraphy, especially when resampled to match the data set length of the target vector time-series (MD99–2269). These alignments are biased to slightly deeper depths in Site U1305 or shallower depths in MD99–2269 by ~ 12 cm, which is consistent with previously observed depth offsets for sediment magnetization that result from pDRM processes. Target and candidate PSV vector time-series can carry considerable random directional noise, and the identification of signal from noise can be difficult. Initial filtering of PSV data may serve to minimize this noise and therefore produce a more realistic alignment. Future applications should consider whether filtering is appropriate for their records.

As a final perspective, we note that this algorithm in no way replaces the palaeomagnetist. The results of this algorithm are dependent on the quality of the input vector time-series, which require the same careful evaluation needed for any palaeomagnetic study. Proper evaluation of the library of mathematical solutions produced through this algorithm requires expert knowledge and geological context. However, use of the library of reproducible alignments, tests of statistical significance, and the palaeomagnetist's expertise provides a means for more objective PSV stratigraphy, particularly when independent chronostratigraphic methods, such as radiocarbon or $\delta^{18}\text{O}$, are not available.

ACKNOWLEDGEMENTS

This work was supported by the NSF award EAR-1645411 to Stoner. Hagen acknowledges support from the NSF GRFP and the ARCS Foundation. Data sets are available at the Oregon

State University Palaeo- and Environmental Magnetic Lab website (<http://paleomag.ceoas.oregonstate.edu>). We thank Tilo von Dobe-neck, editor Frederik Simons and an anonymous reviewer, whose comments greatly improved the paper.

Author Contributions: C.J.H., B.T.R., J.S.S. and J.R.C. conceived the study and contributed to analysis. C.J.H. developed the algorithm with help from B.T.R. B.T.R. and C.J.H. wrote the manuscript with input from J.S.S. and J.R.C.

REFERENCES

- Balbas, A.M., Koppers, A.A.P., Clark, P.U., Coe, R.S., Reilly, B.T., Stoner, J.S. & Konrad, K., 2018. Millennial-Scale instability in the geomagnetic field prior to the Matuyama-Brunhes reversal, *Geochem. Geophys. Geosyst.*, **19**, 952–967.
- Bard, E., 1988. Correction of accelerator mass spectrometry ^{14}C ages measured in planktonic foraminifera: Paleoceanographic implications, *Paleoceanography*, **3**, 635–645.
- Caron, M., St-Onge, G., Montero-Serrano, J.-C., Rochon, A., Georgiadis, E., Giraudeau, J. & Massé, G., 2019. Holocene chronostratigraphy of northeastern Baffin Bay based on radiocarbon and palaeomagnetic data, *Boreas*, **48**, 147–165.
- Channell, J.E.T. & Guyodo, Y., 2004. The matuyama chronozone at ODP site 982 (Rockall Bank): evidence for decimeter-scale magnetization lock-in depths, in *Timescales of the Paleomagnetic Field*, pp. 205–219, eds Channell, J.E.T., Kent, D.V., Lowrie, W. & Meert, J.G., American Geophysical Union.
- Channell, J.E.T., Kanamatsu, T., Sato, T., Stein, R., Alvarez Zarikian, C.A. & Malone, M.J. the Expedition 303/306 Scientists (Eds.), 2006. *Proceedings of the IODP*, Vol. 303/306, Integrated Ocean Drilling Program. Retrieved from http://iodp.tamu.edu/publications/exp303_306/30306toc.htm.
- Channell, J.E.T., Xuan, C. & Hodell, D.A., 2009. Stacking paleointensity and oxygen isotope data for the last 1.5 Myr (PISO-1500), *Earth planet. Sci. Lett.*, **283**, 14–23.
- Clark, R.M., 1985. A FORTRAN program for constrained sequence-slotting based on minimum combined path length, *Comput. Geosci.*, **11**, 605–617.
- deMenocal, P.B., Ruddiman, W.F. & Kent, D.V., 1990. Depth of post-depositional remanence acquisition in deep-sea sediments: a case study of the Brunhes-Matuyama reversal and oxygen isotopic Stage 19.1, *Earth planet. Sci. Lett.*, **99**, 1–13.
- Dunhill, G., Andrews, J.T. & Kristjánssdóttir, G., 2004. *Radiocarbon Date List X: Baffin Bay, Baffin Island, Iceland, Labrador Sea, and the Northern North Atlantic (No. 56)*. Occasional Paper, 1–77, University of Colorado, Boulder: Institute of Arctic and Alpine Research.
- Egli, R. & Zhao, X., 2015. Natural remanent magnetization acquisition in bioturbated sediment: General theory and implications for relative paleointensity reconstructions, *Geochem. Geophys. Geosyst.*, **16**, 995–1016.
- Eiriksson, J. *et al.*, 2011. Coupling of palaeoceanographic shifts and changes in marine reservoir ages off North Iceland through the last millennium, *Palaeogeog. Palaeoclimat. Palaeoecol.*, **302**, 95–108.
- Fisher, R., 1953. Dispersion on a Sphere, *Proc. Natl. Acad. Sci. U. S. A.*, **217**, 295–305.
- Haam, K.-W.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**, 1–14.
- Hay, C.C., Creveling, J.R., Hagen, C.J., Maloof, A.C. & Huybers, P., 2019. A library of early Cambrian chemostratigraphic correlation from a reproducible algorithm, *Geology*, **47**, 457–460.
- Irving, E. & Major, A., 1964. Post-depositional detrital remanent magnetization in a synthetic sediment, *Sedimentology*, **3**, 135–143.
- Jennings, A., Thordarson, T., Zalzal, K., Stoner, J., Hayward, C., Geirsdóttir, Á. & Miller, G., 2014. Holocene tephra from Iceland and Alaska in SE greenland shelf sediments, *Geol. Soc., Lond., Spec. Publ.*, **398**, SP398.6, doi:10.1144/SP398.6.

- Johnson, C.L. & McFadden, P., 2007. Time-Averaged field and paleosecular variation, in *Treatise on Geophysics*, pp. 417–453, ed. Schubert, G., Elsevier.
- Korte, M., Brown, M.C., Gunnarson, S.R., Nilsson, A., Panovska, S., Wardinski, I. & Constable, C.G., 2018. Refining Holocene geochronologies using palaeomagnetic records, *Quat. Geochronol.*, **50**, 47–74.
- Kristjánssdóttir, G.B., Stoner, J.S., Jennings, A.E., Andrews, J.T. & Grönvold, K., 2007. Geochemistry of Holocene cryptotephra from the North Iceland Shelf (MD99-2269): intercalibration with radiocarbon and palaeomagnetic chronostratigraphies, *Holocene*, **17**, 155–176.
- Lisiecki, L.E. & Lisiecki, P.A., 2002. Application of dynamic programming to the correlation of paleoclimate records, *Paleoceanography*, **17**, 1–12.
- Lisiecki, L.E. & Raymo, M.E., 2005. A Pliocene-Pleistocene stack of 57 globally distributed benthic $\delta^{18}\text{O}$ records, *Paleoceanography*, **20**, 1–17.
- Lougheed, B. & Obrochta, S., 2016. MatCal: open source Bayesian ^{14}C age calibration in MATLAB, *J. Open Res. Software*, **4**, e42, doi:10.5334/jors.130.
- Lougheed, B. & Obrochta, S., 2019. A rapid, deterministic age-depth modeling routine for geological sequences with inherent depth uncertainty, *Paleoceanogr. Paleoclimatol.*, **34**, 122–133.
- Lund, S.P., 2007. Paleomagnetic secular variation, in *Encyclopedia of Geomagnetism and Paleomagnetism*, pp. 766–776, Springer.
- Lund, S.P. & Keigwin, L., 1994. Measurement of the degree of smoothing in sediment paleomagnetic secular variation records: an example from late Quaternary deep-sea sediments of the Bermuda Rise, western North Atlantic Ocean, *Earth planet. Sci. Lett.*, **122**, 317–330.
- Mazaud, A., 2006. A first-order correction to minimize environmental influence in sedimentary records of relative paleointensity of the geomagnetic field, *Geochem. Geophys. Geosyst.*, **7**, Q07002. doi:10.1029/2006GC001257.
- Mazaud, A., Channell, J.E.T. & Stoner, J.S., 2012. Relative paleointensity and environmental magnetism since 1.2Ma at IODP site U1305 (Eirik Drift, NW Atlantic), *Earth planet. Sci. Lett.*, **357–358**, 137–144.
- Nilsson, A., Muscheler, R. & Snowball, I., 2011. Millennial scale cyclicity in the geodynamo inferred from a dipole tilt reconstruction, *Earth planet. Sci. Lett.*, **311**, 299–305.
- Nilsson, A., Snowball, I., Muscheler, R. & Uvo, C.B., 2010. Holocene geocentric dipole tilt model constrained by sedimentary paleomagnetic data, *Geochem. Geophys. Geosyst.*, **11**. doi:10.1029/2010GC003118.
- Reilly, B.T. et al., 2018. Regionally consistent Western North America paleomagnetic directions from 15 to 35 ka: assessing chronology and uncertainty with paleosecular variation (PSV) stratigraphy, *Quat. Sci. Rev.*, **201**, 186–205.
- Reilly, B.T. et al., 2019. Holocene Break-up and reestablishment of the peterrmann ice tongue, Northwest Greenland, *Quat. Sci. Rev.*, **128**, 322–342.
- Reimer, P.J., McCormac, F.G., Moore, J., McCormick, F. & Murray, E.V., 2002. Marine radiocarbon reservoir corrections for the mid-to-late Holocene in the eastern subpolar North Atlantic, *Holocene*, **12**, 129–135.
- Reimer, P.J. et al., 2013. IntCal13 and marine13 radiocarbon age calibration curves 0–50,000 years cal BP, *Radiocarbon*, **55**, 1869–1887.
- Roberts, A.P. & Winklhofer, M., 2004. Why are geomagnetic excursions not always recorded in sediments? Constraints from post-depositional remanent magnetization lock-in modelling, *Earth planet. Sci. Lett.*, **227**, 345–359.
- Sakoe, H. & Chiba, S., 1978. Dynamic programming algorithm optimization for spoken word recognition, *IEEE Trans. Acoust., Speech Signal Process.*, **26**, 43–49.
- Schwartz, M., Lund, S. & Johnson, T., 1996. Environmental factors as complicating influences in the recovery of quantitative geomagnetic-field paleointensity estimates from sediments, *Geophys. Res. Lett.*, **23**, 2693–2696.
- Simon, Q., Bourlès, D.L., Thouveny, N., Horng, C.-S., Valet, J.-P., Bassinot, F. & Choy, S., 2018. Cosmogenic signature of geomagnetic reversals and excursions from the Réunion event to the Matuyama–Brunhes transition (0.7–2.14 Ma interval), *Earth planet. Sci. Lett.*, **482**, 510–524.
- Stoner, J.S., Channell, J.E.T., Mazaud, A., Strano, S.E. & Xuan, C., 2013. The influence of high-latitude flux lobes on the Holocene paleomagnetic record of IODP Site U1305 and the northern North Atlantic, *Geochem. Geophys. Geosyst.*, **14**, 4623–4646.
- Stoner, J.S., Jennings, A., Kristjánssdóttir, G.B., Dunhill, G., Andrews, J.T. & Hardardóttir, J., 2007. A paleomagnetic approach toward refining Holocene radiocarbon-based chronologies: Paleooceanographic records from the north Iceland (MD99-2269) and east Greenland (MD99-2322) margins, *Paleoceanography*, **22**, 1–23.
- Stoner, J.S. & St-Onge, G., 2007. Chapter three magnetic stratigraphy in paleoceanography: Reversals, excursions, paleointensity, and secular variation, in *Developments in Marine Geology*, Vol. 1, pp. 99–138, Elsevier.
- Strunk, A., Larsen, N.K., Nilsson, A., Seidenkrantz, M.-S., Levy, L.B., Olsen, J. & Lauridsen, T.L., 2018. Relative sea-level changes and ice sheet history in Finderup Land, North Greenland, *Front. Earth Sci.*, **6**, 129.
- Suganuma, Y., Yokoyama, Y., Yamazaki, T., Kawamura, K., Horng, C.-S. & Matsuzaki, H., 2010. ^{10}Be evidence for delayed acquisition of remanent magnetization in marine sediments: Implication for a new age for the Matuyama–Brunhes boundary, *Earth planet. Sci. Lett.*, **296**, 443–450.
- Tauxe, L., 1993. Sedimentary records of relative paleointensity of the geomagnetic field: theory and practice, *Rev. Geophys.*, **31**, 319–354.
- Thompson, R., 1973. Palaeolimnology and Palaeomagnetism, *Nature*, **242**, 182–184.
- Thompson, R. & Clark, R.M., 1989. Sequence slotting for stratigraphic correlation between cores: theory and practice, *J. Paleolimnol.*, **2**, 173–184.
- Valet, J.-P. & Fournier, A., 2016. Deciphering records of geomagnetic reversals, *Rev. Geophys.*, **54**, 410–446.
- Valet, J.-P., Herrero-Bervera, E., LeMouél, J.-L. & Plenier, G., 2008. Secular variation of the geomagnetic dipole during the past 2000 years, *Geochem. Geophys. Geosyst.*, **9**, doi:10.1029/2007GC001728.
- Verosub, K.L., 1977. Depositional and postdepositional processes in the magnetization of sediments, *Rev. Geophys.*, **15**, 129–143.
- Walczak, M.H. et al., 2017. A 17,000 yr paleomagnetic secular variation record from the southeast Alaskan margin: regional and global correlations, *Earth planet. Sci. Lett.*, **473**, 177–189.
- Wanamaker, A.D., Butler, P.G., Scourse, J.D., Heinemeier, J., Eiriksson, J., Knudsen, K.L. & Richardson, C.A., 2012. Surface changes in the North Atlantic meridional overturning circulation during the last millennium, *Nat. Commun.*, **3**, 899.
- Weeks, R. et al., 1993. Improvements in long-core measurement techniques: applications in palaeomagnetism and palaeoceanography, *Geophys. J. Int.*, **114**, 651–662.
- Xuan, C., Channell, J.E.T. & Hodell, D.A., 2016. Quaternary magnetic and oxygen isotope stratigraphy in diatom-rich sediments of the southern Gardar Drift (IODP Site U1304, North Atlantic), *Quat. Sci. Rev.*, **142**, 74–89.
- Ólafsdóttir, S., Geirsdóttir, A., Miller, G.H., Stoner, J.S. & Channell, J.E.T., 2013. Synchronizing Holocene lacustrine and marine sediment records using paleomagnetic secular variation, *Geology*, **41**, 535–538.
- Ólafsdóttir, S., Reilly, B.T., Bakke, J., Stoner, J.S., Gjerde, M. & van der Bilt, W.G.M., 2019. Holocene paleomagnetic secular variation (PSV) near 80° N, Northwest Spitsbergen, Svalbard: implications for evaluating High Arctic sediment chronologies, *Quat. Sci. Rev.*, **210**, 90–102.

SUPPORTING INFORMATION

Supplementary data are available at *GJI* online.

Fig. S1. Visualizing the meaning of the RCD and TCD parameters. (a) Inclination and (b) declination for synthetic target (black) and candidate (red) records that have an XC of 0.50. (c) Projection of the candidate vectors (red circles) with a model generated using linear models of the x , y and z components of the candidate record as a function of the x , y and z components of the target record (black). Blue stippled lines are used to illustrate the difference between

the data and the model. (d) Projection of the candidate vectors (red circles) with blue stippled lines used to illustrate the difference between the candidate vectors and the mean of the candidate vectors. (e–h) As in a–d, but for a synthetic target and candidate record with an XC of 0.27.

Please note: Oxford University Press is not responsible for the content or functionality of any supporting materials supplied by the

authors. Any queries (other than missing material) should be directed to the corresponding author for the paper.

APPENDIX: ARCHIVED CODE

The algorithm used in this work is available to download on Github (<https://github.com/CedricHagen/PSV-dynamic-time-warping>).