

Detecting paleoclimate transitions with Laplacian Eigenmaps of Recurrence Matrices (LERM)

Alexander James¹, Julien Emile-Geay¹, Nishant Malik², Deborah Khider³

¹University of Southern California, Department of Earth Sciences

²School of Mathematics and Statistics, Rochester Institute of Technology, Rochester, NY-14623

³University of Southern California, Information Sciences Institute

Key Points:

- Laplacian eigenmaps of recurrence matrices (LERM) is a novel tool for paleoclimate time series analysis.
- LERM can robustly detect the gradual Mid-Pleistocene Transition in relatively low signal-to-noise ratio scenarios.
- LERM can also be applied to detect abrupt climate transitions like the 8.2ka event, though less robustly.

Corresponding author: Alexander James, akjames@usc.edu

Abstract

Paleoclimate records can be considered low-dimensional projections of the climate system that generated them. Understanding what these projections tell us about past climates, and changes in their dynamics, is a main goal of time series analysis on such records. Laplacian Eigenmaps of Recurrence Matrices (LERM) is a novel technique using univariate paleoclimate time series data to indicate when notable shifts in dynamics have occurred. LERM leverages time delay embedding to construct a manifold that is mappable to the attractor of the climate system; this manifold can then be analyzed for significant dynamical transitions. Through numerical experiments with observed and synthetic data, LERM is applied to detect both gradual and abrupt regime transitions. Our paragon for gradual transitions is the Mid-Pleistocene Transition (MPT). We show that LERM can robustly detect gradual MPT-like transitions for sufficiently high signal-to-noise ratios, though with a time lag related to the embedding process. Our paragon of abrupt transitions is the “8.2ka” event; we find that LERM is generally robust at detecting 8.2ka-like transitions for sufficiently high signal-to-noise ratios, though edge effects become more influential. We conclude that LERM can usefully detect dynamical transitions in paleogeoscientific time series, with the caveat that false positive rates are high when dynamical transitions are not present, suggesting the importance of using multiple records to confirm the robustness of transitions. We share an open source Python package to facilitate the use of LERM in paleoclimatology and paleoceanography.

1 Introduction

Much of the current discussion on our changing climate centers around the concept of tipping points (Alley et al., 2003; Lenton et al., 2008; Steffen et al., 2018). Climate tipping points occur when a change in the climate system becomes self-perpetuating (McKay et al., 2022). They describe moments in the evolution of a climate system during which the behavior of the climate changes in a fundamental way. In other words, they bridge the gap between separate dynamical regimes. In reference to global warming, climate tipping points are typically used to describe moments in which positive feedback loops are created, resulting in runaway warming. More generally, within the context of nonlinear dynamical systems theory, tipping points are the critical thresholds; when crossed, they lead to abrupt and irreversible changes to the dynamics of the underlying system, i.e., these are points in the parameter space of the system, where, due to influences such as noise, perturbations or parameter drift, the shape of the system’s typical trajectory, or attractor, changes significantly (Kaszás et al., 2019).

Within the context of paleoclimate, tipping points are interesting because they can inform us about conditions under which the climate has undergone fundamental changes in the past in response to forcings and might do so again. Given the increasingly unstable nature of our current climate system, understanding when and where tipping points have occurred in the past is deeply valuable for policy makers, scientists, and citizens alike. Additionally, assuming we are able to observe synchronous tipping points at different locations or between different archive types and proxy records, it can inform our understanding of the history of climate teleconnections as well as how changes in climate regimes are reflected in various paleoclimate records.

Developing analytical tools to detect significant changes in system dynamics is an ongoing field of study (Kantz & Schreiber, 2003; Marwan, Carmen Romano, et al., 2007). In this paper we will explore the application of a novel four-step method we colloquially refer to as *Laplacian Eigenmaps of Recurrence Matrices (LERM)*, originally developed and published by Malik (2020). In their paper, Malik (2020) provided evidence that LERM was able to robustly detect changes in the dynamics of an idealized experiment in the presence of noise and missing values before applying it to a Holocene speleothem record to probe questions regarding the climate’s influence on the collapse of the Harappan civ-

ilization. We seek to expand upon their findings, providing further validation that this technique is effective in identifying significant climate regime changes when applied to paleoclimate records, as well as exploring some potential shortcomings. We also present an open source Python package meant to simplify carrying out this workflow, alongside Jupyter notebooks to support the reproducibility of our results (James, 2023b).

2 LERM: a basic algorithm

In this section, we briefly describe the method and its core principles. A much more thorough discussion can be found in the original publication (Malik, 2020). LERM is a recurrence plot-based time series analysis technique. Recurrence plots/matrices are a popular non-linear time series analysis method that transforms a time series into a binary matrix, in which elements with value one correspond to time points close in phase space (Eckmann et al., 1987; Marwan, Carmen Romano, et al., 2007; Zou et al., 2019). Analysis of spatial patterns in a recurrence plot using dynamical systems theory can provide deep insights into the nonlinear and stochastic dynamics of the system underlying the data (Eckmann et al., 1987; Marwan, Carmen Romano, et al., 2007; Zou et al., 2019; Bradley & Kantz, 2015). The LERM method consists of four main steps.

2.1 Step 1: Phase Space Reconstruction

Nonlinear time series analysis relies on phase space reconstruction, which projects the time series on a time-delay coordinate system. This time-delay embedding of a time series is a consequence of the classical theorem by Floris Takens, colloquially known as Takens' theorem, which states conditions under which a topologically equivalent attractor can be constructed from single scalar observations (Takens, 1981; Packard et al., 1980). Time-delay embedding constructs phase space vectors from time-shifted snippets of a time series $x(t)$ of length N . For example, for time delay τ and embedding dimension m , a vector in time-delay embedding would be $\mathbf{x}(t) = [x(t), x(t - \tau), x(t - 2\tau), \dots, x(t - m\tau)]$. The parameter m determines the length of the phase space vectors. The standard technique for choosing m is the method of false nearest neighbor (Abarnabel, 1997; Kantz & Schreiber, 2003). However, heuristics show that in the presence of noise, the principle of over-embedding (Hegger et al., 2000; Malik et al., 2014) is more appropriate. This principle suggests taking $m > 2(D + P)$ where D is the dimensionality of the system and P is the number of time-dependent parameters. Our numerical experimentation indicates that m between 10 and 15 leads to robust results for our application. The parameter τ (time-delay) can be chosen as the time point corresponding to the first minimum of lagged mutual information or the first zero of the autocorrelation function; for details, see Abarnabel (1997). Further discussion of these choices can be found in Abarnabel (1997); Kantz and Schreiber (2003); Malik (2020); Malik et al. (2014).

Although the method of time delay embedding has been known to introduce spurious correlations into phase space trajectories and spurious structures into the recurrence plot (Thiel et al., 2006; Wendi et al., 2017), certain metrics are less dependent on embedding parameter choices. For example, Thiel et al. (2006) observed that second-order Renyi entropy and correlation dimension can be calculated using arbitrary embedding parameter choices. Similarly, as we will show, LERM appears to be robust with respect to embedding parameter choices. Additionally, Wendi et al. (2017) demonstrated that over-embedding leads to more reliable measurement of the determinism metric.

For any given time series, phase space vectors are created for all points along the time axis for which it is possible. Note that, due to indexing constraints, phase space vectors cannot be constructed for the last $m \cdot \tau$ points. The above method of time-delay embedding satisfies the condition of Takens' theorem: the phase space reconstructed using suitable time-delay embedding of time series data is topologically equivalent to the original phase space of the system (Takens, 1981). In practice, uneven spacing of data,

noisy sensors, and imperfect selection criteria for the embedding dimension and delay parameters prevent perfect topological equivalency. However, if proper care is taken in the data selection and embedding steps, the reconstructed phase space can still provide deep insights into the system’s dynamics.

2.2 Step 2: Recurrence Plot

The next step is to analyze recurrence relationships within the reconstructed phase space. Both recurrence quantification analysis (RQA) and recurrence network analysis (RNA) focus on characterizing recurrence plots. Recurrence plots (RP) are graphical representations of the recurrence matrix of a time series, which is a binary square matrix of size N defined as $R_{ij} = \Theta(\epsilon - \|\mathbf{x}_i - \mathbf{x}_j\|)$. \mathbf{x}_i and \mathbf{x}_j are time embedded vectors at time points i and j . Θ is the Heaviside step function, i.e., $\Theta(y) = 1$ if $y > 0$ and otherwise $\Theta(y) = 0$ and $\|\mathbf{x}_i - \mathbf{x}_j\|$ is the distance between embedded vectors \mathbf{x}_i and \mathbf{x}_j (in this work we use the Euclidean norm). The threshold ϵ is interpretable as a radius defining the largest distance that can separate two points in phase space if they are considered in the same neighborhood. If the distance between two points is greater than ϵ , the value inside the Heaviside function will be negative and the recurrence matrix will record a zero at that intersection. If the distance between two points in phase space is less than ϵ , then the recurrence matrix entry is unity at that intersection, indicating that the system is visiting a similar state at those indices. ϵ is typically chosen so that the recurrence density (number of ones in the recurrence matrix divided by the total number of entries) is around 5%, a heuristic which is supported by other studies on the topic (Kraemer et al., 2018; Malik et al., 2014; Malik, 2020).

2.3 Step 3: Laplacian eigenmaps

Laplacian eigenmaps is a manifold learning (nonlinear dimensionality reduction) technique, where the eigenvectors of the Laplacian corresponding to a proximity graph constructed from a point cloud of the data are used to project the data onto lower dimensional space (Belkin & Niyogi, 2003). Laplacian eigenmaps are closely related to spectral clustering techniques and, similarly, preserve the local distance between points. Laplacian eigenmaps are used here to extract low-dimensional structures from an RP, as these low-dimensional structures are the basis of RP-based metrics and analysis. For example, diagonal lines in an RP are related to the determinism of the underlying system (Marwan, Romano, et al., 2007). We expect that, as the system moves between different dynamical regimes, the manifolds extracted through our technique should also evolve and change, and hence these low-dimensional manifolds will track transitions in dynamical regimes.

To calculate the Laplacian, we first define the elements of the weighted adjacency matrix \mathbf{W} of the graph as $W_{ij} = R_{ij} + 1$ and then the corresponding graph Laplacian is $\mathbf{L} = \mathbf{D} - \mathbf{W}$, where \mathbf{D} is a diagonal matrix with $D_{jj} = \sum_j W_{ij}$. To construct \mathbf{W} , the 1 is added to the each element of the recurrence matrix to avoid numerical complications when solving for the eigenvalue problem (see below). The graph Laplacian or the Laplacian matrix \mathbf{L} can be considered the discrete analog to the continuous version of the Laplacian operator, and it is used to model diffusion on graphs (Merris, 1994). To obtain the manifolds, we solve the eigenvalue problem $\mathbf{L}\phi = \lambda\mathbf{D}\phi$. Let $\phi_0 \cdots \phi_{N-1}$ be the solution of this eigenvalue problem with $0 = \lambda_0 \leq \lambda_1 \leq \cdots \leq \lambda_{N-1}$ being the corresponding eigenvalues. The first eigenvector ϕ_0 is dropped as it corresponds to the eigenvalue 0, and all elements in it are ones. The manifolds are obtained by projecting each point in the reconstructed phase space \mathbf{x}_i to the m -dimensional Euclidean space: $[\phi_1(i), \cdots, \phi_m(i)]$.

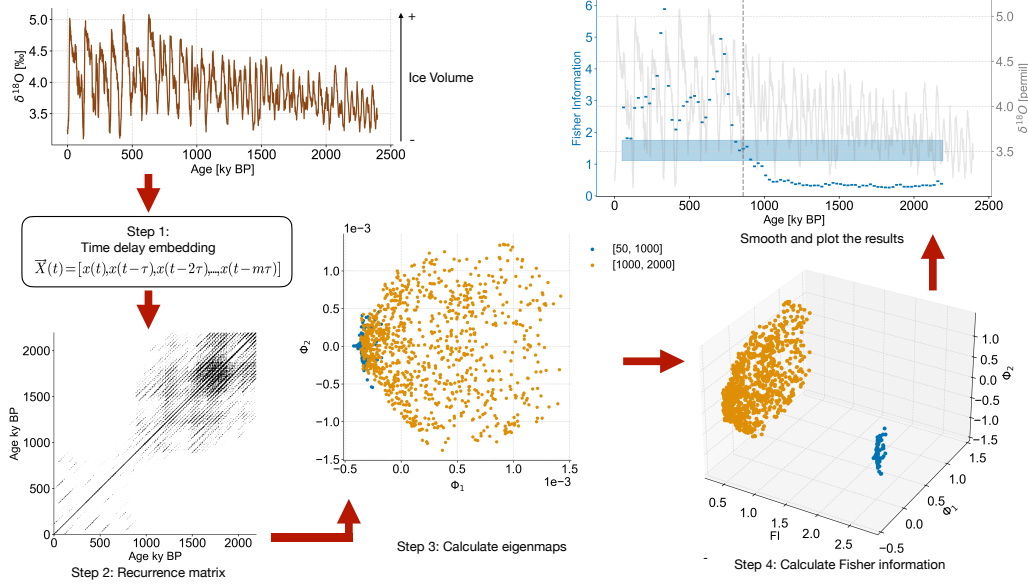


Figure 1: Description of the full workflow presented in this section. In this figure we use the LR04 benthic stack from Lisiecki and Raymo (2005) as an example, and examine the dynamical transition that occurred around 1000 kyr BP known as the Mid-Pleistocene Transition. Step 1, time delay embedding, is described further in Section 2.1. Step 2, generating the recurrence plot, is described in Section 2.2. Step 3, the calculation of the graph laplacian and its eigenvectors (creating the eigenmaps) is described in Section 2.3. The last step, step 4, describing the calculation of the Fisher information statistic is described in Section 2.4. We also show the final result, plotting the evolution of the Fisher information statistic over time. Here we smooth the statistic using a block size of 5 to isolate the dominant statistic behavior, and calculate a confidence interval to detect significant transitions in the statistic. The calculation of this confidence interval is described in Section 2.5. The dashed gray line in the final plot shows the detected transition.

2.4 Step 4: Fisher Information

Laplacian eigenmaps result in p -dimensional projections of the original data, and our numerical experimentation indicates that $p = 4$ produces the most robust results; higher values only add redundant information to the analysis, whereas lower values do not always lead to stable results. From this low-dimensional subspace we then seek to create a univariate statistic that reflects changes in the complexity and dynamics represented by the multidimensional eigenmaps in order to ease interpretability. To do so, Malik (2020) proposed a modified version of the Fisher information statistic (FI). As defined in Ahmad et al. (2016), the FI is an invariant over the manifolds resulting from Laplacian eigenmaps, i.e., as the dynamics of underlying system change regimes, the extracted manifolds change. Consequently, FI captures this regime change as a single numerical value (Malik, 2020). In general, FI is a practical and robust way of discovering shifts in multivariate data's behavior and information content (Ahmad et al., 2016). FI can also be thought of as a way to measure the complexity of the underlying dynamics, as it can capture the complexity of the geometric object that represents a dynamic process, for instance, an attractor. The segments of the time series where the values of FI are higher (lower) are also the sections of the time series where the underlying dynamics are of higher (lower) complexity. Our numerical experimentations indicate that FI behaves like an invariant metric or a constant of motion; values remain the same over the same dynamical regime (when control parameters are kept fixed). This means that when the parameter changes significantly, a change in dynamics has occurred. We discuss how we assess the significance of changes in the next section.

The calculation of FI requires the specification of two key parameters: window size and window increment. These specify the size and step of the sliding window that will be used to calculate FI. The choice of these parameters is arbitrary, and depends on the record and phenomenon being studied. Larger window sizes and increments will result in a smoothing effect, improving the robustness of results while reducing time resolution and smoothing over smaller transitions. Smaller window sizes and increments will tend to introduce more spurious behavior, but will also improve time resolution of the FI and allow for the detection of subtler shifts in time series character. Given this, two competing factors drive the choice of window size. The first is that one wants the window size to be long enough that FI values converge towards a stable value. That is, the window size should not be so small that FI is not convergent or robust. The second is that the time scale on which one would like to resolve the transitions must not be so large that multiple transition points get fused into one. This is especially important when dealing with abrupt events. Further explanation and justification of this choice for this specific problem, and the specific variant of FI we are using, can be found in Malik (2020).

The endpoint of the FI window is used to determine the time index of the FI value for a given window. Typically we then take a block average over another window of several consecutive FI values. This average is then plotted for all the points within that window. That is, all the points within this window are assigned the same average FI. This minimizes the possible artifacts of the start/center/endpoint choice for a block average over a window. In this paper we occasionally do not do this in order to show the actual variability of the FI, which in some cases can be erratic. When this is the case, the FI is plotted with a fill (e.g. Figure 4). When we have smoothed the FI we plot it as a scatter, with the width of scatter points indicating the width of the smoothing window (e.g. Figure 2) A pictorial overview of this workflow is shown in Figure 1.

2.5 Significance of Transitions

To determine whether changes in our FI statistic are significant we employ the same strategy as Malik (2020). Our null hypothesis for this test is that no transition has occurred. This would be indicated by the FI statistic not exhibiting a significant change

in values. Significance in this case is ascertained via the usage of a confidence interval. To do this, we sample with replacement from the FI series, creating an ensemble of M samples with w points each. Typically we set M to 10,000, and $w = 50$. This choice was largely arbitrary, as the analysis did not show strong sensitivity to these parameters. We then take the mean of each of these samples and calculate a confidence interval from the distribution of means. The bounds for this confidence interval are typically taken to be 5% and 95%. When the FI statistic crosses this confidence interval, moving either from above the 95% boundary to below the 5% boundary or vice versa, we claim that this is a significant change, thereby marking a transition in the dynamical regime of the system. The midpoint of this transition is taken as the transition timing. This is only one approach to establishing significance, and others may be possible. We note that this significance test often produces many false positives when applied to time series without dynamical transitions, and as such all results should be verified across multiple records (see below).

3 Detecting gradual transitions

In this section we demonstrate how LERM performs when applied to records that are known to contain a gradual shift in dynamics. For our gradual transition we chose the Mid-Pleistocene Transition (MPT). The MPT was a transition from the “41kyr world” to the “100 kyr world” (Paillard, 2001). That is, the dominant periodicity of the glacial-interglacial cycles switched from 41 thousand years to 100 thousand years. The transition occurred over several hundred thousand years, from around 1200ka to 800ka (Clark et al., 2006; Chalk et al., 2017). There are many theories as to why this transition occurred, which are not germane to our purpose as they all indicate the presence of a dynamical change. We are primarily interested in the ability of the LERM technique to detect the MPT in real paleoclimate archives. In order to study the outcome of applying LERM to data influenced by the MPT we applied the method to five benthic oxygen isotope records drawn from Lisiecki and Raymo (2005) as well as data from a conceptual glacial/interglacial cycle model (Leloup & Paillard, 2022). Oxygen isotopes were chosen as during the Pleistocene they are typically interpreted as representing changes in ocean temperature and global ice volume (Waelbroeck et al., 2002), with ice volume being the dominant signal. If there is a significant change in the dynamics that control global ice volume such as the MPT, we should be able to observe it by applying LERM to benthic foraminiferal oxygen isotope observations.

3.1 Tests with observational data

We apply the technique to oxygen isotope records from marine sediment cores taken at five Ocean Drilling Project (ODP) sites 925 (Bickert et al., 1997; Billups et al., 1998; Franz & Tiedemann, 2002), 927 (Bickert et al., 1997; Franz & Tiedemann, 2002), 929 (Bickert et al., 1997; Billups et al., 1998; Franz & Tiedemann, 2002) and 846 (Mix et al., 1995; Shackleton et al., 1995), and 849 (Mix et al., 1995). Their oxygen isotope records were drawn from the compilation of Lisiecki and Raymo (2005), and the age models for each are those aligned to the age model of the LR04 stack. Core locations are shown in Figure 2 and their traces are shown in Figure 2. These records were chosen due to their length and general lack of hiatuses. Each record was linearly interpolated using their mean time increment (2.67, 3.92, 3.43, 2.46, 3.10 kilo-years respectively) in order to produce a uniform time axis for each record. The records can be roughly subdivided into two geographical groups, those in the East Pacific and those in the West Atlantic. In this case, record locations starting with the number eight lie in the East Pacific and those starting with nine lie in the West Atlantic. The geographic division of these records means that if we observe any local effects, they will likely be apparent in the results. ϵ values were selected by finding the value that produced a density of 5% in the recurrence matrix, in accordance with the recommendation of Kraemer et al. (2018). m was chosen ac-

cording the principle of over-embedding as described by Malik (2020) and set to 13 indices. τ was set by calculating the first minimum of lagged mutual information in accordance with the recommendation of Abarnabel (1997). In this example these values range between 4 and 8 indices. Window size and window increment were set to 50 (roughly 100 - 150 kyr) and 5 indices respectively in pursuance with the recommendation of Section 2.4.

The results of this analysis are shown in Figure 2. There is strong agreement between these records as to the timing of a climate regime transition. The mean value and standard deviation of the transition is 908 ± 66 kyr BP (1σ). This agrees with what we would expect to see, assuming the MPT was the dominant climate regime transition in this set of records. We then place all of the records onto a shared, evenly spaced time axis. The timestep for this shared axis is the mean of each of the records, and the bounds are the maximum of the minimum and the minimum of the maximum of the collection of record time axes. That is, the most conservative endpoints are chosen such that all records cover the full shared time axis. Each record is then linearly interpolated over this shared axis. No changes to the underlying age models are made during this process. By doing so, we find that our mean transition occurs at 911 ± 51 kyr BP (1σ), reducing the uncertainty in this estimate. This reduction in uncertainty, while small in this case, illustrates the importance of time axis considerations when conducting this kind of analysis. We will further explore such considerations in section 5. We also note that the standard deviation of a transition timing across records is not necessarily the best measure of uncertainty. We recommend the employment of ensemble based approaches for more robust uncertainty quantification.

3.2 Tests with synthetic data

3.2.1 A conceptual model for glacial cycles

We applied LERM to the conceptual model presented by Leloup and Paillard (2022). This model generates a unitless variable v which represents normalized ice volume. The equation that controls the evolution of this variable depends on whether the model is in the slow glaciation regime (g) or the fast deglaciation regime (d). The equations that define how each of these states govern the change in v over time are shown in Equation 1.

$$(g) \quad \frac{dv}{dt} = -\frac{I}{\tau_i} + \frac{1}{\tau_g} \quad (1a)$$

$$(d) \quad \frac{dv}{dt} = -\frac{I}{\tau_i} - \frac{v}{\tau_g} \quad (1b)$$

Equation 2 describes when the model is to switch from (g) to (d) and vice versa.

$$(d) \text{ to } (g) : I < I_0 \quad (2a)$$

$$(g) \text{ to } (d) : I + v > V_0 \quad (2b)$$

τ_i , τ_d , and τ_g are time constants, and I is normalized summer insolation forcing at 65°N . By varying the deglaciation parameter V_0 , we can emulate a dynamical change in the evolution of ice volume similar to the one observed during the MPT. We generated a time series of v with a length of 2500 time units and placed a transition from a V_0 value of 3.4 to 5.2 at time step 1000. This was in accordance with Leloup and Paillard (2022), who evaluated which values of V_0 most accurately reflected the pre- and post-MPT ice volume dynamics. We used summer solstice insolation at 65°N as our insolation scenario as this produced the most accurate results for the last 1500 Ma (Leloup

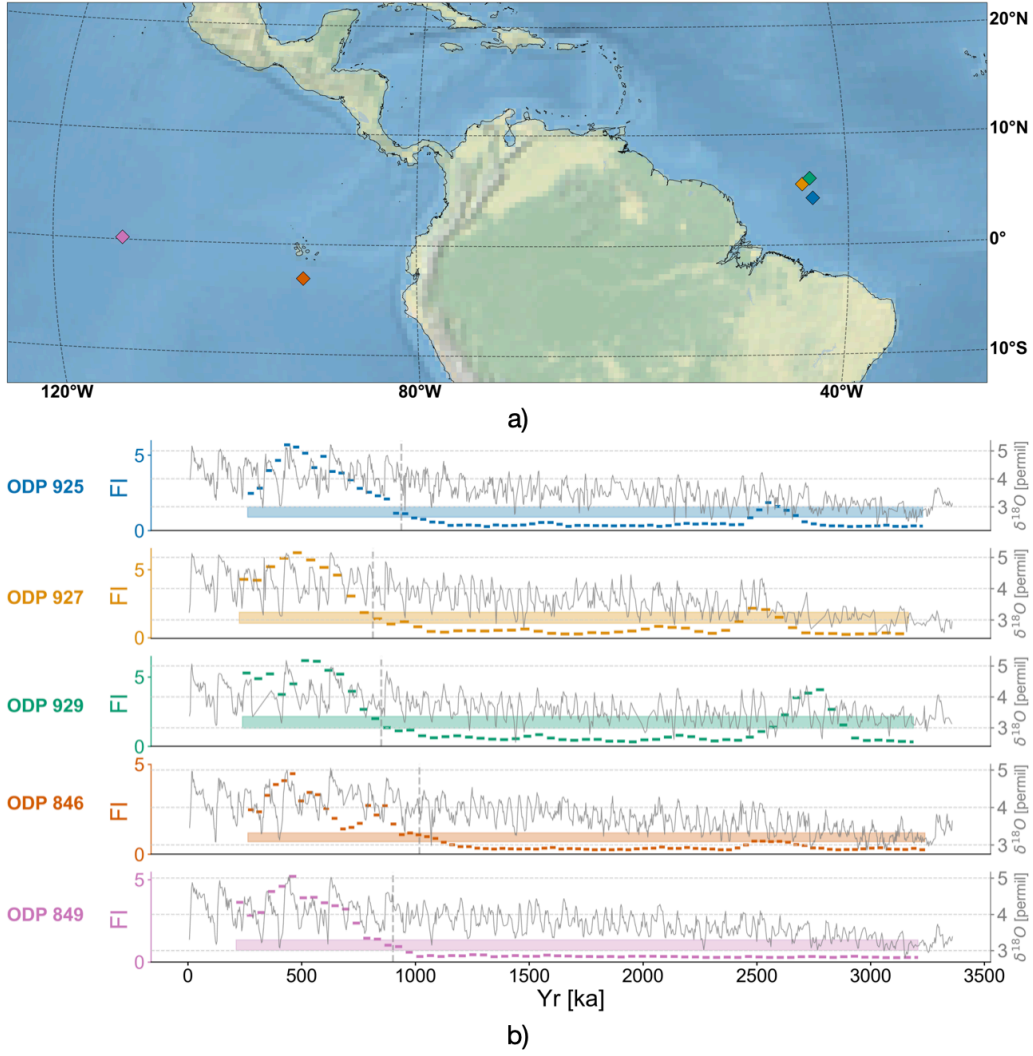


Figure 2: Overview of LERM analysis applied to five ODP records. (a) show a map of the considered records, (b) shows the oxygen isotope time series trace in black, and the Fisher information statistic in color.

& Paillard, 2022). We then bin the series to the time axes of Cores 925, 927, 929, 846, and 849 from the Ocean Drilling Project (ODP). This is done by placing bin edges between time points for each time axis and averaging the conceptual model data over each bin, assigning each time point its associated average. This was done in order to compare the effect of differing time axes on our analysis. The LERM workflow is then applied to each of these series. The same parameters are used here as those in the previous section, though τ now varies between 3 and 4 indices. The results of this analysis are shown in Figure 3.

The method performs reliably when applied to this simple test, locating the point of the transition with relative accuracy, regardless of which time axis was used. The mean transition timing for each of the binned series is 849 ± 44 kyr BP (1σ). We conducted several other tests to see how the method responded to the addition of noise, missing values, and how it behaved when no transition was present at all. These are presented in Figure 4. Each of these tests uses the version of the conceptual model time series binned onto Core 925.

3.2.2 Sensitivity Analysis

Here we present the results of the LERM method when applied to the conceptual model MPT time series depicted in the top panel of Figure 3 when varying levels of noise are present. We used an AR(1) model to create a noise time series and added it to our conceptual model MPT time series "signal". We define the signal to noise ratio (S/N) as being the standard deviation of the MPT signal divided by the standard deviation of the noise. We tested four S/N ratios. An example of one of the tests using a S/N ratio approximately equal to two is shown in Figure 4. The actual signal to noise ratio for a given test approximates the targeted S/N ratio as our noise generation process is imprecise and cannot create a series with exactly the standard deviation necessary to create the targeted S/N ratio.

We repeated the process of creating and analyzing noisy time series 1000 times for each S/N ratio. We then produced a Kernel Density Estimate (KDE) of the distribution of detected transitions and normalized each KDE to have a maximum amplitude of 1. These KDEs are interpreted as representing the "probability" (the normalization process means the y-axis cannot be literally interpreted as such) of a transition occurring at different points in time given a specific S/N ratio. The results of this analysis are shown in the top panel of Figure 7.

For higher S/N ratios, the method is quite consistent in its detection of the primary transition point. However, as the S/N ratio decreases, while the detection of the actual transition point remains consistent, the method begins to return a large number of false positives. This suggests that this tool is best used within a comparative context in order to bolster the confidence of the results. That is, if multiple records agree on the timing of a transition, this is good evidence that the transition is real. If not, we may just be observing spurious system behavior resulting from our requirement that our recurrence matrix be 5% populated. This effect is especially apparent in panels c) and d) of Figure 4. When no transitions are present in a series, we observe random fluctuations in the FI metric. This is likely due to the minimum recurrence density requirement used when choosing a recurrence threshold and from our confidence interval based approach to significance testing. Alternate approaches to selection of the recurrence threshold parameter ϵ and the definition of significant transitions could improve these results. In the meantime, this is further evidence of the necessity of verifying potential dynamical transitions across multiple records.

Additionally, when the stable time series is artificially coarsened and then interpolated, as shown in panels e) and f) of Figure 4, the LERM technique suggests that the coarsened section experienced a change in dynamics. This effect only became noticeable

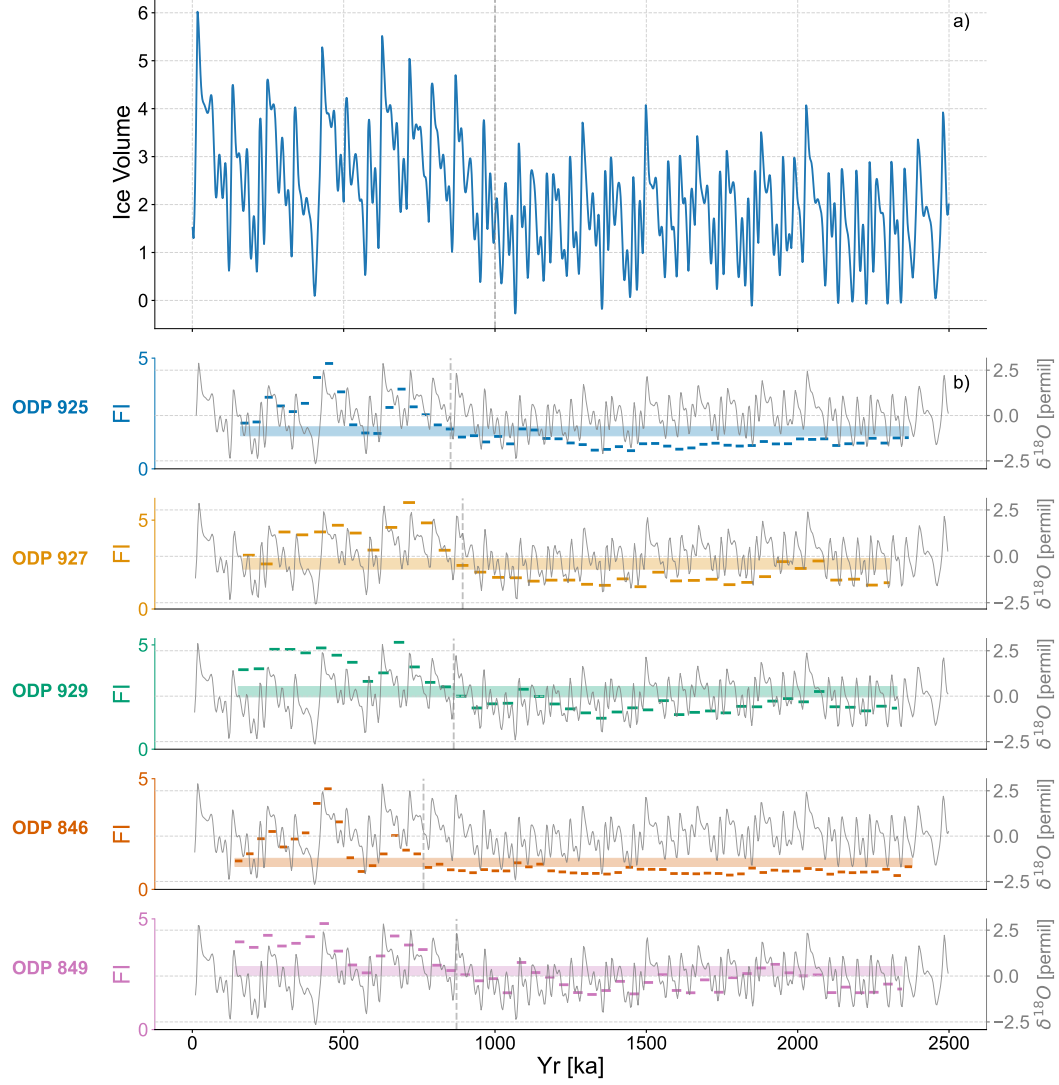


Figure 3: Overview of the application of the LERM pipeline to conceptual model data. The top panel shows data generated by the Leloup Paillard glacial/interglacial cycle conceptual model using with transition of V_0 threshold parameter from 3.4 to 5.2 at time step 1000 (vertical dashed line). The lower panels show the Fisher information statistic of conceptual model data from the top panel after being binned onto the time axis of each ODP core (shown in black). The binning process is described in Section 3.2.1.

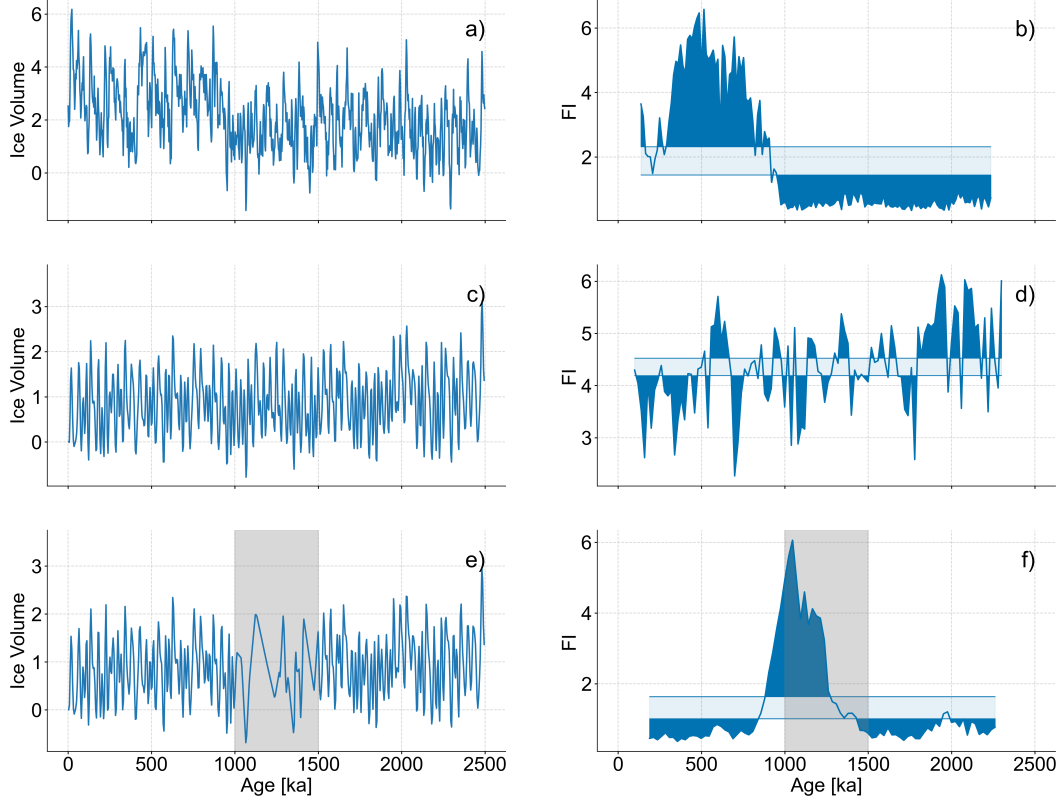


Figure 4: LERM sensitivity analysis for gradual transitions. Conceptual model data is binned onto ODP core 925 oxygen isotope time series time axis, to which we add AR(1) noise. Signal to noise (S/N) ratio is defined as the standard deviation of the conceptual model data divided by the standard deviation of the noise. Shown in the top panel is an example of this test applied to a noisy series with an S/N ratio of 2. In the middle panel we show the result of applying the technique to a series without a transition. In the bottom panel, we apply the method to a stable series that has a coarsened section over which we have interpolated to show the propensity of the technique to return false positives for imputed sections of a record when no transitions are present.

in our experiments when over 60% of the points in a section were removed, though this likely depends heavily on the series being considered and the length of the coarse subsection. As such, caution is advised when applying this method to evenly spaced versions of unevenly sampled time series, which are common within the field of paleoclimatology. In such case it is important that transitions observed near changes in resolution be viewed with skepticism. On the other hand, in paleoclimate research a change in resolution can suggest a change in system dynamics. For example, a decrease in the resolution of speleothems can be interpreted as reflecting a period of aridification in the region (Henselowsky et al., 2021). Teasing apart these effects is difficult and domain-specific. This further emphasizes the importance of comparing multiple records when using this technique. If detected dynamical transitions cannot be reproduced in nearby records, they are unlikely to reflect robust changes in climate dynamics.

4 Detecting abrupt transitions

In this section we demonstrate how LERM perform when applied to a record that contains an abrupt, short-lived, transition from one climate regime to another. To do this, we apply the technique to four Greenland ice core oxygen isotope records and an Antarctic ice core oxygen isotope record, which have been interpreted as being proxies for temperature (Jouzel et al., 1997; Johnsen et al., 2001). In doing so we explore the potential climate regime shifts that occurred around the 8.2ka event, a period of intense, abrupt cold (Alley & Ágústsson, 2005) that has been observed primarily in Greenland ice cores (Thomas et al., 2007), though it has appeared in other archives from the Northern Hemisphere (Cheng et al., 2009) as well as some from the Southern Hemisphere (Chase et al., 2015). We also apply the technique to synthetic data designed to mimic the Greenland ice core oxygen isotope records.

4.1 Tests with observational data

The four Greenland cores we analyzed are NGRIP (Andersen et al., 2004), Renland (Johnsen et al., 1997), GRIP (Johnsen et al., 1992), and GISP2 (Grootes & Stuiver, 1997) ice core records. The time axes of the GRIP and NGRIP records have been aligned to Greenland Ice Core Chronology 2005 (GICC05) (Rasmussen et al., 2006; Vinther et al., 2006). We also analyze oxygen isotope data from a high resolution section of ice cores from EPICA Dome C (Stenni et al., 2010). and the locations and traces of all the ice core records are shown in Figure 5. Each of these records was interpolated to its mean time step (20, 10, 5, 20, 18 years respectively). Among them, we observe strong agreement regarding the effect of the 8.2ka event on climate dynamics in the region.

The agreement between the Greenland records as to the timing of the onset of the change in dynamics is somewhat unsurprising given the evident anomalous nature of the 8.2ka event in the time series. However, the results do illustrate the edge effects that are inherent with time delay embedding techniques, as can be seen in the appearance of a climate regime shift occurring prior to the 8.2ka event. This is caused by the way time delay embedded vectors are constructed. Each vector that is associated with a given time point contains time information that extends $m \cdot \tau$ beyond that time point. In our case we used $m = 12$ and $\tau = 4$, continuing our practice of choosing m via over-embedding. Note that we choose tau by hand here, as the first minimum of mutual information heuristic fails when applied to some of these records, resulting in excessively large values for tau. This means that data from the following 48 time points are included in a given time point of our embedded data. Then, if the resolution of our time series is 20 years, the information from the subsequent 960 years is included in any given year. This can result in a smearing effect, where changes in dynamics that happen at one point in the time axis can affect the result of our analysis at a different point in time. This smearing effect is uni-directional, occurring only in the direction of the time delay embedding, which

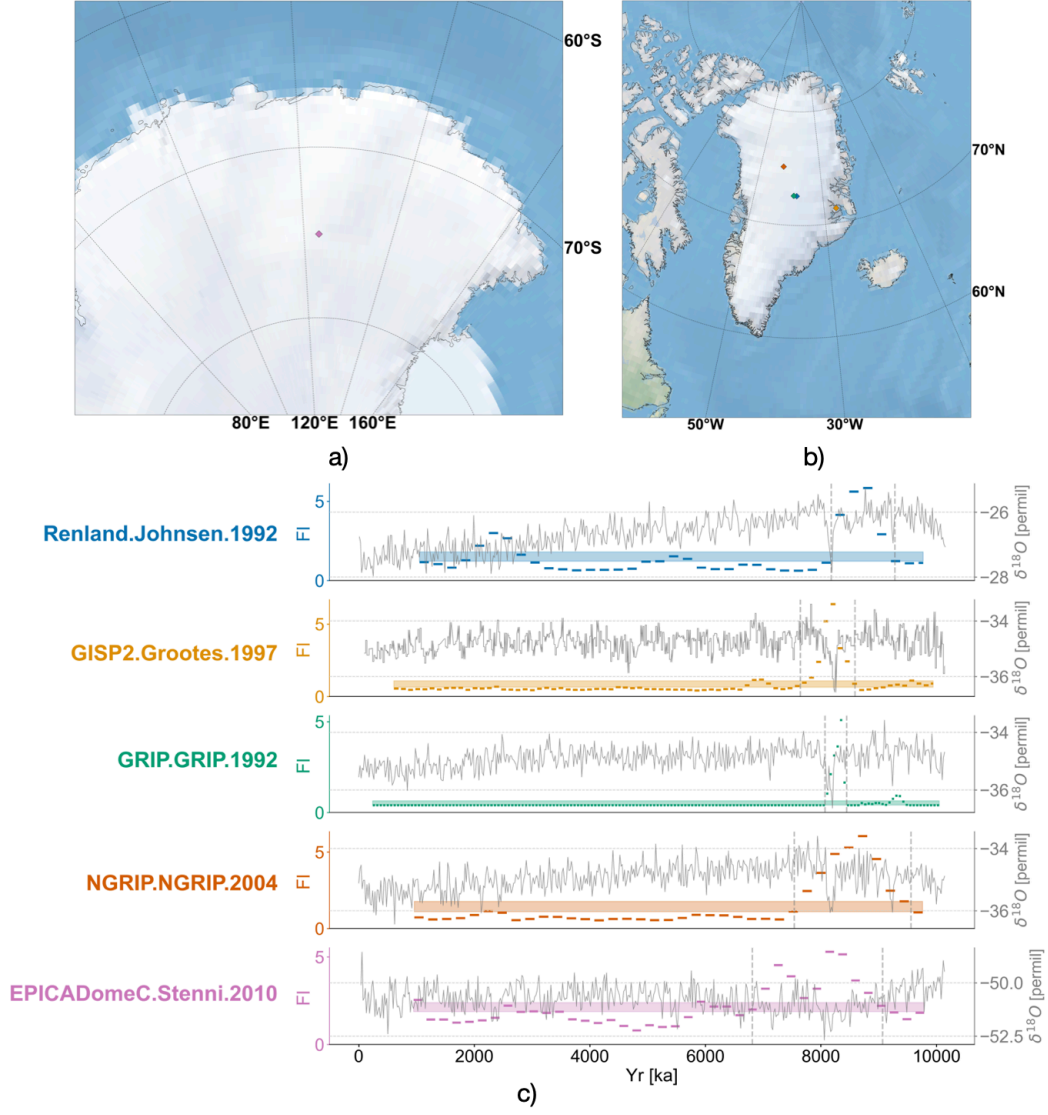


Figure 5: Overview of LERM analysis applied to Greenland and Antarctica records. (a) shows a map of the Greenland records, (b) shows a map of the Antarctica record, (c) shows the oxygen isotope time series trace in black, and the Fisher information statistic in color.

is best taken with the flow of time in order to preserve the temporal structure of the record in the constructed phase space.

We also observe that the choice of window size and window increment can exacerbate this smearing effect. When a large window size is chosen, the detection timing of abrupt transitions tends to "widen", or move outward from the actual edges of the transition. In order to minimize this effect we used a smaller window size of 20 indices (around 100 to 400 years) and window increment of 4 indices for these tests. Window increment tends to have less of an effect (see the accompanying Holocene Ice Window Increment and Holocene Ice Window Size notebooks for examples of this effect (James, 2023b)).

What is somewhat surprising is the agreement between these records and the high resolution oxygen isotope record from EPICA Dome C, suggesting hemispheric synchrony between Greenland and Eastern Antarctica during the 8.2ka event. However, this result should be viewed with some caution, as other records we tested from different regions in Antarctica do not show a synchronous climate regime transition at this point (see the Holocene Ice Analysis notebook from James (2023b)).

4.2 Tests with synthetic data

To investigate LERM's behavior in a controlled setting, we once again resort to synthetic data. Our signal is defined as a ramp with a peak amplitude of -1 . The onset of the spike occurred at 8400 kyr BP, terminating at 7800 kyr BP in order to produce a signal that was not as easily nullified by noise and more consistently represented an 8.2k event-like signal, as shorter events tended to be entirely obscured when noise was added. This signal is shown in Figure 6. To test detection, we added perturbations to this signal using a simple AR(1) process with an autocorrelation coefficient of 1. In this case the S/N ratio is the amplitude of the perturbation divided by the standard deviation of the AR(1) series. We repeated our sensitivity analysis as in Section 3.2.2, this time using relative large S/N ratios of 1, 2, 3, and 4, as the method proved to be less robust for brief, abrupt transitions than gradual ones. We constructed our synthetic series using the same time axis as the NGRIP oxygen isotope record for these experiments. Parameters used were $m = 13$, $\tau = 5$, $w_{size} = 20$ (window size), $w_{inc} = 4$ (window increment). The KDEs from these experiments are shown in the bottom panel of 7. The asymmetrical offset of detection times from the edges of the transition are driven by the edge effects associated with using the FI statistic in the case of the termination side of the event, and a combination of the FI window effect with the uni-directional time-embedding effect in the case of the beginning side of the event (resulting in a greater offset than the termination side).

The results are similar to those of the gradual transition synthetic tests, though the S/N ratios required to achieve reliable detection are much higher for the abrupt transition. This result suggests that unless signal to noise ratios are high, this method will return a large number of false positives when regime transitions present in a record are shorter and less durable. However, it seems that the detection of the shift is consistent even at lower S/N ratios, again suggesting the benefit of applying this technique to an ensemble of records and looking for shared transitions as a way of filtering out false positives.

5 Time Axis Considerations

Throughout this paper we have mentioned the importance of time axis properties at various points. In this section we demonstrate how the detection of dynamical regimes changes depending on the resolution of the time axis. For this we will use the marine sediment oxygen isotope data from ODP Site U1308 (Hodell et al., 2008). This is a high

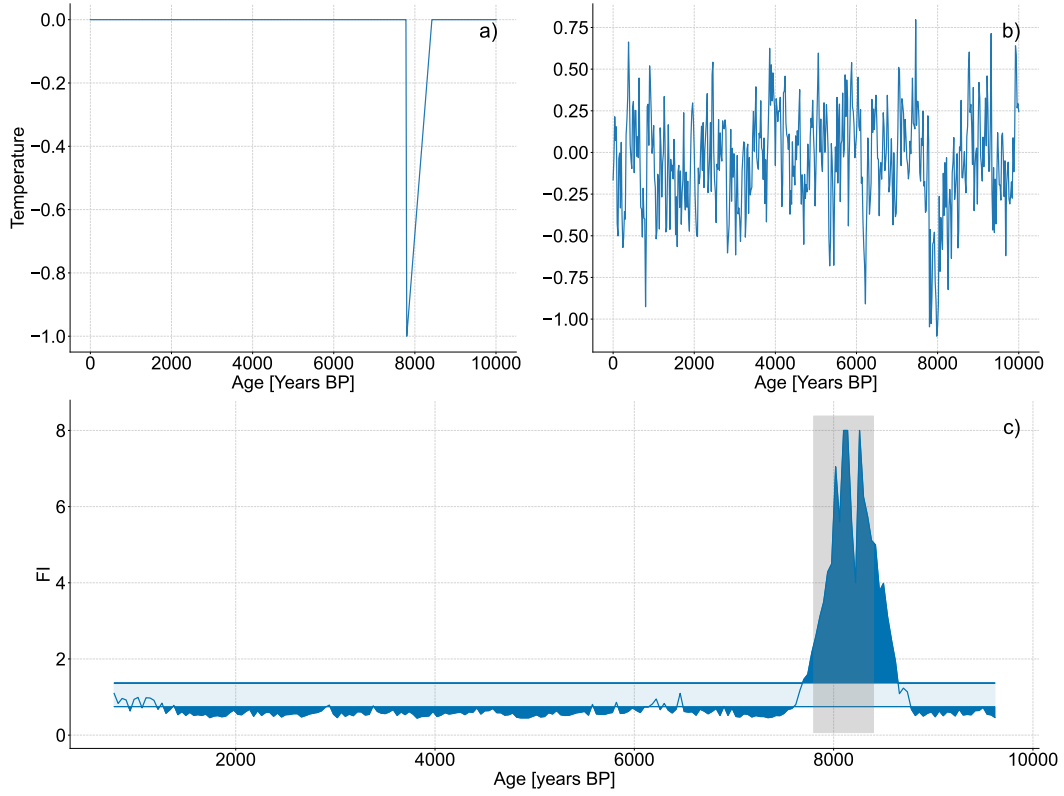


Figure 6: Overview of LERM sensitivity analysis applied to synthetic data designed to mimic the 8.2ka event. The top left panel shows the signal used in these tests. In the top right panel we show the noisy series to be analyzed. In the bottom panel we see the result of the analysis. The grey region indicates the spike interval.




Figure 7: Kernel density estimates (KDE) of detected transitions for MPT-like scenarios (top) and 8.2ka event-like scenarios (bottom). KDEs were normalized against their maximum value.

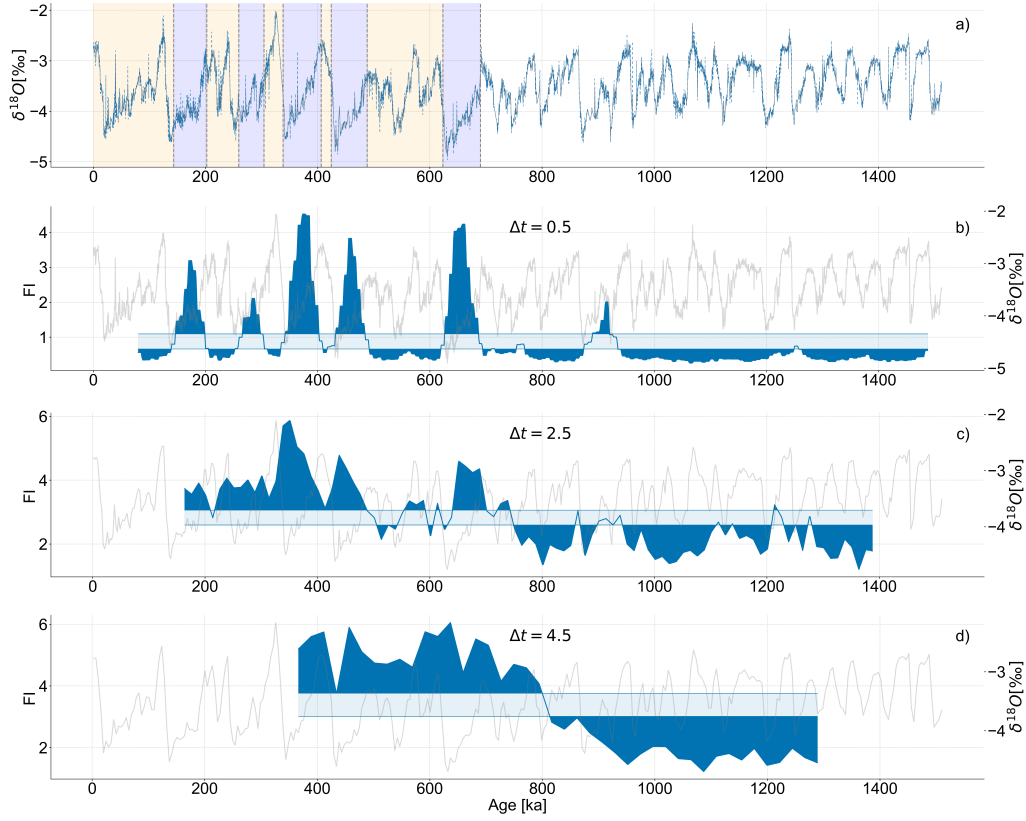


Figure 8: Results of LERM analysis applied to oxygen isotopes from U1308 with different time resolutions. Panel a) shows the original, unaltered oxygen isotope record from U1308 in blue, the version of the record with values averaged over bins of 500 years in grey, and the transitions after the MPT detected in the Fisher information for the binned version of the series highlighted in blue and orange. Panels b), c), and d) show the analysis in blue applied with a time step of .5, 2.5, and 4.5 kiloyears respectively. Oxygen isotopes from U1308 with values averaged over their associated time step are shown in grey.

resolution deep-sea core with a median time step in the published age model of approximately 270 years. By averaging oxygen isotope values across bins of varying sizes, we can coarsen the series to various time steps, and examine how this changes the results of our analysis. First we interpolate the time series using the mean time step, which is approximately 302 years. This is done in order to prevent gaps in the binned version of the time series, as there is one short section with low resolution. We prioritize creating a continuous binned version of the time series in order to emulate our workflow from the previous sections. We use bin sizes of .5, 2.5, and 4.5 kiloyears. We use an embedding dimension of 13 indices and τ values ranging from 5 indices for the maximum bin size and 12 indices for the minimum bin size. Window size and window increment are again set to 50 and 5 indices respectively. The results of this analysis are shown in Figure 8.

In Figure 8, when the time step is relatively large as in panel d), we detect only the Mid-Pleistocene transition. However, as we move to finer time steps, we begin to observe the detection of other regime shifts. These appear to be glacial-interglacial shifts, as dynamical transitions are observed every 100 kiloyears during shifts from glacial to interglacial periods and vice versa. This effect illustrates the importance of temporal resolution when using this technique. Higher resolution allow for the detection of shorter

regime shifts. With coarser time series we primarily detect the gradual regime shifts that occur over longer time scales. If one is primarily interested in a relatively gradual transition like the MPT, it can be useful to coarsen our record in order to minimize the detection of these shorter transitions, like glacial-interglacial cycles.

Another detail worth noting is that we do not observe the detection of glacial-interglacial cycles during the "41ky world" (the interval during which these cycles have a periodicity of 41ky). This is likely due to the minimum resolution we use. Were we to employ an even higher resolution of this time series, we might be able to get at these higher frequency phenomena. However, in this case, only the detection of the slower 100 kiloyear cycles is available to us with this choice of τ and m . The precise relationship between the time scale of the phenomena being detected and the resolution of the time series is not yet known, though it likely depends both on resolution and the parameters chosen. Constraining this relationship further is a subject for future work.

Another effect worth noting is the later detection of the MPT observed in core U1308 compared to the other sediment cores we've examined. The transition detected in this the coarse version of this record shown in Figure 8 panel d) occurs a little before 800 kiloyears ago, which is somewhat at odds with the timing observed in Figure 2. This is caused by the shorter length of the oxygen isotope time series from U1308. If we shorten the time series used in Figure 2, we observe a similarly delayed transition timing (not shown). This is likely due to the way in which we select our recurrence threshold parameter. Because we require a recurrence matrix density of 5%, the recurrence threshold we pick will depend both on the dynamics which are present in a given time series, and the prevalence of those dynamics in a particular record. The more stable the dynamics of a given time series are, the lower and more selective the recurrence threshold will be, and vice versa. Different recurrence thresholds mean that different sections of the time series will be considered recurrent and the detected transition timings will change. In the case of U1308, the record is shorter and less of the pre-MPT interval is present. This results in a lower recurrence threshold, which results in different recurrence patterns, in turn leading to the detected transition timing being pushed towards the end of MPT window. This is the point in time which the 100 kyr cycles have begun in earnest and as such, where the change in dynamics is more evident. There are other approaches to the choice of recurrence threshold that may minimize this effect, though we do not explore them here. The effect described above is demonstrated further in the MPT Core Length Comparison notebook found in the supplement (James, 2023b).

6 Discussion

These examples show that the LERM technique shows promise in the detection of gradual and abrupt regime transitions in paleoclimate records. It robustly detected the Mid-Pleistocene Transition in a set of marine sediment oxygen isotope records, and was resistant against noise and missing values when applied to synthetic data. One caveat is that the method is prone to false positives when a time series does not contain a regime transition. It also shows strong sensitivity to the resolution of the time axis.

As with any recurrence-based technique, there are a few key considerations that must be taken into account when determining what time axis to use for a given record when applying this technique. The time axis must be evenly spaced; this is a strict requirement of methods that rely on uniform time delay embedding. The time axis should also minimize the generation of new data (upsampling). Recurrence analysis based techniques are often very sensitive to changes in time series structure. Interpolating over coarse sections of a time series using too fine a time step can produce false positives. It is best when using these techniques to be as conservative with one's time imputation scheme as possible. When comparing multiple records it can be valuable to align the time axes of each record via a technique such as linear interpolation of the time series onto a shared

time axis. This will minimize the possibility of time axis dependent effects influencing the results of one's analysis. There is a trade-off between consistency of time axes and maintenance of the original time axis. In certain cases this trade-off can be mitigated by usage of skillfully aligned records, such as in the case of the GICC05 time scale used to align Greenland ice cores. However, this requires the independent construction of aligned time axes for a specific set of records, which can either be expensive or impossible depending on the records under consideration. The question of the appropriate time axis to use in these studies is best handled on a case by case basis. Ideally, the results produced by this type of analysis should be reasonably robust across time axes, and variation in the precise timing of transitions due to different time axes and parameter choices should be included in the uncertainty quantification. Another topic related to the choice of time-axis that is worthy of scrutiny is the usage of age model ensembles. None of the records we considered in our analysis had age model ensembles, so we leave the determination of how to handle this kind of uncertainty for future work.

When applied to a set of four oxygen isotopes records from Greenland ice cores and one from an Antarctica ice core, this technique suggested a shift in climate dynamics around the 8.2ka event. In Greenland this was unsurprising given the obvious shift in the character of the time series during the event, but the Antarctica result is intriguing. The result should be viewed with caution, as other records from different regions of Antarctica did not experience the same shift. However, this could be the product of an interdecadal bipolar seesaw mediated by the Atlantic ocean (Chylek et al., 2010; Wang et al., 2015). In this case, due to the abrupt nature of the 8.2ka event, its effects as mediated by this teleconnection could have manifested more as a dynamical disturbance than an opposing temperature response.

7 Conclusion

With appropriate parameter selection and precautions regarding time sampling, the LERM technique shows promise in application to paleoclimate time series data. The technique can reveal, in a holistic way, when changes in the behavior of univariate time series occur. Such changes can be gradual or abrupt, subtle or obvious. However, when noise levels are high or data are unevenly spaced and require the usage of imputation methods, the method can produce a high rate of false positives. Additionally, because the creation of the recurrence matrix relies on a minimum density, it is inherently relative. This means that if there are no dominant changes in time series character present in a series, this method will still report transitions between dynamical regimes. It follows that this method is best applied to sets of records so that synchronicity between records can act to establish robustness to noise and sampling issues. It may also be useful in modern applications when used for tipping point analysis as a dynamically-motivated changepoint detection algorithm. We leave this application for future work.

8 Open Research

v0.0.8 of the Ammonyte Python package (James, 2023a) was used to generate all the examples in this study and the supporting Jupyter Notebooks. Ammonyte is available via a GPL-3.0 license and developed openly at <https://github.com/alexkjames/Ammonyte>. v0.4.0 of the accompanying Jupyter Notebooks (James, 2023b) that provide examples of how each of the figures in this study were produced and additional tests referred to in the text is available via a MIT license and developed openly at https://github.com/alexkjames/Detecting_Paleoclimate_Transitions_with_LERM

References

Abarnabel, H. D. I. (1997). *Analysis of observed chaotic data*. Springer.

- Ahmad, N., Derrible, S., Eason, T., & Cabezas, H. (2016). Using Fisher information to track stability in multivariate systems. *Royal Society Open Science*, 3(11), 160582. doi: 10.1098/rsos.160582
- Alley, R. B., & Ágústssdóttir, A. M. (2005). The 8k event: cause and consequences of a major Holocene abrupt climate change. *Quaternary Science Reviews*, 24(10–11), 1123–1149. doi: 10.1016/j.quascirev.2004.12.004
- Alley, R. B., Marotzke, J., Nordhaus, W. D., Overpeck, J. T., Peteet, D. M., Jr., R. A. P., ... Wallace, J. M. (2003). Abrupt climate change. *Science*, 299(5615), 2005–2010. doi: 10.1126/science.1081056
- Andersen, K. K., Azuma, N., Barnola, J.-M., Bigler, M., Biscaye, P., Caillon, N., ... members, N. G. I. C. P. (2004). High-resolution record of northern hemisphere climate extending into the last interglacial period. *Nature*, 431(7005), 147–151. doi: 10.1038/nature02805
- Belkin, M., & Niyogi, P. (2003). Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15, 1373–1396.
- Bickert, T., Curry, W. B., & Wefer, G. (1997, Jan). *Late pliocene to holocene (2.6–0 ma) western equatorial atlantic deep-water circulation : Inferences from benthic stable isotopes*. Retrieved from <https://pascal-francis.inist.fr/vibad/index.php?action=getRecordDetail&idt=2134271>
- Billups, K., Ravelo, A. C., & Zachos, J. C. (1998). Early pliocene deep water circulation in the western equatorial atlantic: Implications for high-latitude climate change. *Paleoceanography*, 13(1), 84–95. doi: 10.1029/97pa02995
- Bradley, E., & Kantz, H. (2015). Nonlinear time-series analysis revisited. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 25(9), 097610. doi: 10.1063/1.4917289
- Chalk, T. B., Hain, M. P., Foster, G. L., Rohling, E. J., Sexton, P. F., Badger, M. P. S., ... Wilson, P. A. (2017). Causes of ice age intensification across the Mid-Pleistocene Transition. *Proceedings of the National Academy of Sciences*, 114(50), 13114–13119. doi: 10.1073/pnas.1702143114
- Chase, B. M., Boom, A., Carr, A. S., Carré, M., Chevalier, M., Meadows, M. E., ... Reimer, P. J. (2015). Evolving southwest african response to abrupt deglacial north atlantic climate change events. *Quaternary Science Reviews*, 121, 132–136. doi: 10.1016/j.quascirev.2015.05.023
- Cheng, H., Fleitmann, D., Edwards, R. L., Wang, X., Cruz, F. W., Auler, A. S., ... Matter, A. (2009). Timing and structure of the 8.2 kyr b.p. event inferred from $\delta^{18}O$ records of stalagmites from china, oman, and brazil. *Geology*, 37(11), 1007–1010. doi: 10.1130/g30126a.1
- Chylek, P., Folland, C. K., Lesins, G., & Dubey, M. K. (2010). Twentieth century bipolar seesaw of the arctic and antarctic surface air temperatures. *Geophysical Research Letters*, 37(8). doi: 10.1029/2010gl042793
- Clark, P. U., Archer, D., Pollard, D., Blum, J. D., Rial, J. A., Brovkin, V., ... Roy, M. (2006). The middle pleistocene transition: characteristics, mechanisms, and implications for long-term changes in atmospheric pCO_2 . *Quaternary Science Reviews*, 25(23–24), 3150–3184. doi: 10.1016/j.quascirev.2006.07.008
- Eckmann, J.-P., Kamphorst, S. O., & Ruelle, D. (1987, nov). Recurrence plots of dynamical systems. *Europhysics Letters*, 4(9), 973. Retrieved from <https://dx.doi.org/10.1209/0295-5075/4/9/004> doi: 10.1209/0295-5075/4/9/004
- Franz, S.-O., & Tiedemann, R. (2002). *Stable isotope data of odp hole 154-925b (table a2)*. PANGAEA. Retrieved from <https://doi.pangaea.de/10.1594/PANGAEA.67525> doi: 10.1594/PANGAEA.67525
- Groote, P. M., & Stuiver, M. (1997). Oxygen 18/16 variability in Greenland snow and ice with 10^{-3} to 10^5 year time resolution. *Journal of Geophysical Research: Oceans*, 102(C12), 26455–26470. doi: 10.1029/97jc00880
- Hegger, R., Kantz, H., Matassini, L., & Schreiber, T. (2000). Coping with nonstationarity by overembedding. *Physical Review Letters*, 84(18), 4092–4095. doi:

- 10.1103/physrevlett.84.4092
- Henselowsky, F., Eichstädter, R., Schröder-Ritzrau, A., Herwartz, D., Almoazamy, A., Frank, N., ... Bubenzer, O. (2021). Speleothem growth phases in the central eastern desert of Egypt reveal enhanced humidity throughout MIS 5. *Quaternary International*. doi: 10.1016/j.quaint.2021.05.006
- Hodell, D. A., Channell, J. E. T., Curtis, J. H., Romero, O. E., & Röhl, U. (2008). Onset of “Hudson Strait” Heinrich events in the eastern North Atlantic at the end of the middle Pleistocene transition (640ka)? *Paleoceanography*, 23(4), n/a-n/a. doi: 10.1029/2008pa001591
- James, A. (2023a, April). *Ammonyte: A Python package designed for conducting non-linear time series analysis of paleoclimate data*. [Software]. Zenodo. Retrieved from <https://doi.org/10.5281/zenodo.7843881> doi: 10.5281/zenodo.7843881
- James, A. (2023b, November). *Detecting Paleoclimate Transitions with LERM - Jupyter Notebooks* [Software]. Zenodo. Retrieved from <https://doi.org/10.5281/zenodo.10205104> doi: 10.5281/zenodo.10205104
- Johnsen, S. J., Clausen, H. B., Dansgaard, W., Fuhrer, K., Gundestrup, N., Hammer, C. U., ... Steffensen, J. P. (1992). Irregular glacial interstadials recorded in a new Greenland ice core. *Nature*, 359(6393), 311–313. doi: 10.1038/359311a0
- Johnsen, S. J., Clausen, H. B., Dansgaard, W., Gundestrup, N. S., Hammer, C. U., Andersen, U., ... Fisher, D. (1997). The D18O record along the Greenland ice core project deep ice core and the problem of possible Eemian climatic instability. *Journal of Geophysical Research: Oceans*, 102(C12), 26397–26410. doi: 10.1029/97jc00167
- Johnsen, S. J., Dahl-Jensen, D., Gundestrup, N., Steffensen, J. P., Clausen, H. B., Miller, H., ... White, J. (2001). Oxygen isotope and palaeotemperature records from six Greenland ice-core stations: Camp Century, Dye-3, GRIP, GISP2, Renland and NorthGRIP. *Journal of Quaternary Science*, 16(4), 299–307. doi: 10.1002/jqs.622
- Jouzel, J., Alley, R. B., Cuffey, K. M., Dansgaard, W., Grootes, P., Hoffmann, G., ... White, J. (1997). Validity of the temperature reconstruction from water isotopes in ice cores. *Journal of Geophysical Research: Oceans*, 102(C12), 26471–26487. doi: 10.1029/97jc01283
- Kantz, H., & Schreiber, T. (2003). *Nonlinear time series analysis* (2nd ed.). Cambridge University Press. doi: 10.1017/CBO9780511755798
- Kaszás, B., Feudel, U., & Tél, T. (2019). Tipping phenomena in typical dynamical systems subjected to parameter drift. *Scientific Reports*, 9(1), 8654. doi: 10.1038/s41598-019-44863-3
- Kraemer, K. H., Donner, R. V., Heitzig, J., & Marwan, N. (2018). Recurrence threshold selection for obtaining robust recurrence characteristics in different embedding dimensions. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 28(8), 085720. doi: 10.1063/1.5024914
- Leloup, G., & Paillard, D. (2022). Influence of the choice of insolation forcing on the results of a conceptual glacial cycle model. *Climate of the Past*, 18(3), 547–558. doi: 10.5194/cp-18-547-2022
- Lenton, T. M., Held, H., Kriegler, E., Hall, J. W., Lucht, W., Rahmstorf, S., & Schellnhuber, H. J. (2008). Tipping elements in the Earth’s climate system. *Proceedings of the National Academy of Sciences*, 105(6), 1786–1793. doi: 10.1073/pnas.0705414105
- Lisiecki, L. E., & Raymo, M. E. (2005). A Pliocene-Pleistocene stack of 57 globally distributed benthic $\delta^{18}\text{O}$ records. *Paleoceanography*, 20(1), n/a-n/a. doi: 10.1029/2004pa001071
- Malik, N. (2020). Uncovering transitions in paleoclimate time series and the climate driven demise of an ancient civilization. *Chaos: An Interdisciplinary Journal of*

- Nonlinear Science*, 30(8), 083108. doi: 10.1063/5.0012059
- Malik, N., Marwan, N., Zou, Y., Mucha, P. J., & Kurths, J. (2014, Jun). Fluctuation of similarity to detect transitions between distinct dynamical regimes in short time series. *Phys. Rev. E*, 89, 062908. Retrieved from <https://link.aps.org/doi/10.1103/PhysRevE.89.062908> doi: 10.1103/PhysRevE.89.062908
- Marwan, N., Carmen Romano, M., Thiel, M., & Kurths, J. (2007). Recurrence plots for the analysis of complex systems. *Physics Reports*, 438(5), 237–329. Retrieved from <https://www.sciencedirect.com/science/article/pii/S0370157306004066> doi: <https://doi.org/10.1016/j.physrep.2006.11.001>
- Marwan, N., Romano, M. C., Thiel, M., & Kurths, J. (2007). Recurrence plots for the analysis of complex systems. *Physics Reports*, 438(5–6), 237–329. doi: 10.1016/j.physrep.2006.11.001
- McKay, D. I. A., Staal, A., Abrams, J. F., Winkelmann, R., Sakschewski, B., Loriani, S., ... Lenton, T. M. (2022). Exceeding 1.5 °C global warming could trigger multiple climate tipping points. *Science*, 377(6611), eabn7950. doi: 10.1126/science.abn7950
- Merris, R. (1994). Laplacian matrices of graphs: a survey. *Linear Algebra and its Applications*, 197, 143–176. doi: 10.1016/0024-3795(94)90486-3
- Mix, A., Le, J., & Shackleton, N. (1995, August). Benthic foraminiferal stable isotope stratigraphy of site 846: 0–1.8 ma. In *Proceedings of the ocean drilling program, 138 scientific results*. Ocean Drilling Program. Retrieved from <https://doi.org/10.2973/odp.proc.sr.138.160.1995> doi: 10.2973/odp.proc.sr.138.160.1995
- Packard, N. H., Crutchfield, J. P., Farmer, J. D., & Shaw, R. S. (1980). Geometry from a time series. *Physical Review Letters*, 45(9), 712–716. doi: 10.1103/physrevlett.45.712
- Paillard, D. (2001). Glacial cycles: Toward a new paradigm. *Rev. Geophys.*, 39, 325–346. doi: 10.1029/2000RG000091
- Rasmussen, S. O., Andersen, K. K., Svensson, A. M., Steffensen, J. P., Vinther, B. M., Clausen, H. B., ... Ruth, U. (2006). A new greenland ice core chronology for the last glacial termination. *Journal of Geophysical Research: Atmospheres*, 111(D6). doi: 10.1029/2005jd006079
- Shackleton, N., Hall, M., & Pate, D. (1995). 15. pliocene stable isotope stratigraphy of site 846. In *Proc. ocean drill. program sci. results* (Vol. 138, pp. 337–355).
- Steffen, W., Rockström, J., Richardson, K., Lenton, T. M., Folke, C., Liverman, D., ... Schellnhuber, H. J. (2018). Trajectories of the earth system in the anthropocene. *Proceedings of the National Academy of Sciences*, 115(33), 8252–8259. doi: 10.1073/pnas.1810141115
- Stenni, B., Masson-Delmotte, V., Selmo, E., Oerter, H., Meyer, H., Röthlisberger, R., ... Udisti, R. (2010). The deuterium excess records of epica dome c and dronning maud land ice cores (east antarctica). *Quaternary Science Reviews*, 29(1–2), 146–159. doi: 10.1016/j.quascirev.2009.10.009
- Takens, F. (1981). Detecting strange attractors in turbulence. In D. Rand & L.-S. Young (Eds.), *Dynamical systems and turbulence, warwick 1980* (pp. 366–381). Berlin, Heidelberg: Springer Berlin Heidelberg.
- Thiel, M., Romano, M. C., & Kurths, J. (2006). Spurious structures in recurrence plots induced by embedding. *Nonlinear Dynamics*, 44(1–4), 299–305. doi: 10.1007/s11071-006-2010-9
- Thomas, E. R., Wolff, E. W., Mulvaney, R., Steffensen, J. P., Johnsen, S. J., Arrowsmith, C., ... Popp, T. (2007). The 8.2ka event from greenland ice cores. *Quaternary Science Reviews*, 26(1–2), 70–81. doi: 10.1016/j.quascirev.2006.07.017
- Vinther, B. M., Clausen, H. B., Johnsen, S. J., Rasmussen, S. O., Andersen, K. K., Buchardt, S. L., ... Heinemeier, J. (2006). A synchronized dating of three

- 739 Greenland ice cores throughout the Holocene. *Journal of Geophysical Re-*
740 *search: Atmospheres*, 111(D13). doi: 10.1029/2005jd006921
- 741 Waelbroeck, C., Labeyrie, L., Michel, E., Duplessy, J., McManus, J., Lambeck, K.,
742 ... Labracherie, M. (2002). Sea-level and deep water temperature changes de-
743 rived from benthic foraminifera isotopic records. *Quaternary Science Reviews*,
744 21(1-3), 295–305. doi: 10.1016/s0277-3791(01)00101-9
- 745 Wang, Z., Zhang, X., Guan, Z., Sun, B., Yang, X., & Liu, C. (2015). An atmospheric
746 origin of the multi-decadal bipolar seesaw. *Scientific Reports*, 5(1), 8909. doi:
747 10.1038/srep08909
- 748 Wendi, D., Marwan, N., & Merz, B. (2017). In search of determinism-sensitive re-
749 gion to avoid artefacts in recurrence plots. *International Journal of Bifurcation*
750 *and Chaos*, 28(01), 1850007. doi: 10.1142/s0218127418500074
- 751 Zou, Y., Donner, R. V., Marwan, N., Donges, J. F., & Kurths, J. (2019). Complex
752 network approaches to nonlinear time series analysis. *Physics Reports*, 787, 1–
753 97. Retrieved from [https://www.sciencedirect.com/science/article/pii/](https://www.sciencedirect.com/science/article/pii/S037015731830276X)
754 [S037015731830276X](https://www.sciencedirect.com/science/article/pii/S037015731830276X) doi: <https://doi.org/10.1016/j.physrep.2018.10.005>