



Research paper

Comparing methods for analysing time scale dependent correlations in irregularly sampled time series data

Maria Reschke^{a,b,*}, Torben Kunz^a, Thomas Laepple^{a,1}

^a Alfred Wegener Institute Helmholtz Centre for Polar and Marine Research, Telegrafenberg A45, 14473, Potsdam, Germany

^b Institute of Earth and Environmental Sciences, University of Potsdam, Karl-Liebknecht-Str. 24/25, 14476, Potsdam, Germany



ARTICLE INFO

Keywords:

Irregular sampling
Time scale dependence
Correlation
Resampling methods

ABSTRACT

Time series derived from paleoclimate archives are often irregularly sampled in time and thus not analysable using standard statistical methods such as correlation analyses. Although measures for the similarity between time series have been proposed for irregular time series, they do not account for the time scale dependency of the relationship. Stochastically distributed temporal sampling irregularities act qualitatively as a low-pass filter reducing the influence of fast variations from frequencies higher than about $0.5 (\Delta t_{\max})^{-1}$, where Δt_{\max} is the maximum time interval between observations. This may lead to overestimated correlations if the true correlation increases with time scale. Typically, correlations are underestimated due to a non-simultaneous sampling of time series.

Here, we investigated different techniques to estimate time scale dependent correlations of weakly irregularly sampled time series, with a particular focus on different resampling methods and filters of varying complexity. The methods were tested on ensembles of synthetic time series that mimic the characteristics of Holocene marine sediment temperature proxy records. We found that a linear interpolation of the irregular time series onto a regular grid, followed by a simple Gaussian filter was the best approach to deal with the irregularity and account for the time scale dependence. This approach had both, minimal filter artefacts, particularly on short time scales, and a minimal loss of information due to filter length.

1. Introduction

Variations in past climate and environmental conditions have been recorded in many forms, including sediment and ice cores, from which they can be extracted as paleoclimate time series data. Analysing these data sets provides valuable insights into the drivers affecting both climate and the ecosystem on regional and global scales. One obstacle that complicates the analysis of paleoclimate proxy records using standard statistical methods is the temporal irregularity of the available data.

This patchiness and resulting temporal irregularity may have different reasons. For instance, the physical sampling of the archive might be irregular due to a lack of analysable material. In the case of sedimentary archives, sedimentation or accumulation rates generally vary over time, thus a regular spatial sampling of the core may lead to an irregular temporal sampling of past conditions.

Time series data often contain several processes acting on different time scales (e.g. monthly, decadal, centennial) and one approach to

identify these processes is to compare different data sets and analyse their relationship at different time scales. This can be achieved by estimating the coherency, although robustly determining the phase and the relationship as a function of frequency requires long time series.

A simpler statistic that describes the similarity between two time series is the linear correlation. By definition, the correlation estimate includes information about the shared variability from the Nyquist frequency (half the sampling rate Δt^{-1}) to the reciprocal length of the record.

For many applications, a time scale dependent estimate of the correlation would be useful.

For example, two time series containing the same low frequency signal, but different uncorrelated (white) noise, would exhibit a higher correlation if the correlation estimate was restricted to long time scales (i.e., low frequencies), because at short time scales (i.e., high frequencies) the signal is largely masked by the noise (Fig. 1). Thus, it is necessary to specify for which frequency range the correlation is

* Corresponding author. Alfred Wegener Institute Helmholtz Centre for Polar and Marine Research, Telegrafenberg A45, 14473, Potsdam, Germany.

E-mail address: mreschke@awi.de (M. Reschke).

¹ **Author contributions:** M.R. and T.L. designed the research, M.R. and T.K. and T.L. developed the methodology. M.R. wrote the computer code, performed the analysis and wrote the first draft of the manuscript. All authors contributed to the interpretation and to the preparation of the final manuscript.

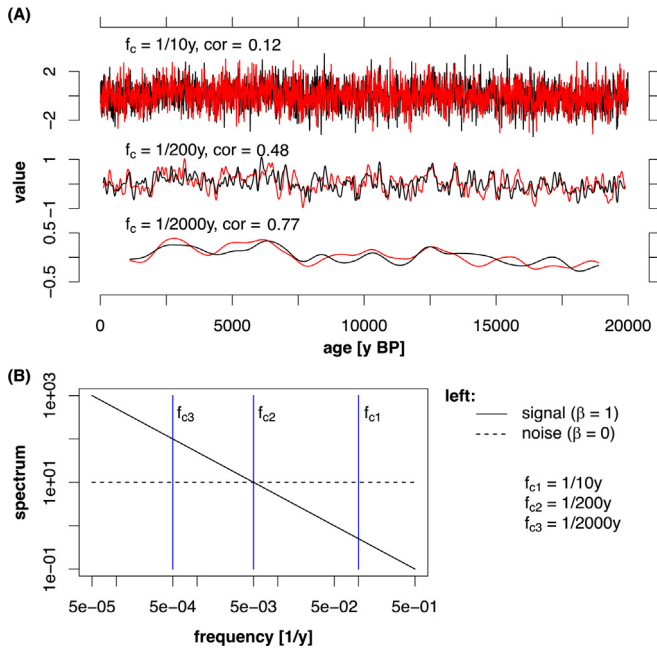


Fig. 1. Time scale dependent correlation of an irregularly sampled time series pair (A) and spectrum of the signal and white noise (B). The time series 1 and 2 (red and black) contain the same signal, generated as red noise with $\beta_s = 1$ and independent white noise components. Both time series are filtered with three cut-off frequencies f_c and the resulting time series are shown (A). The time series containing most frequencies ($f_c = 1/10y$) is basically uncorrelated while slow variations are positively correlated as seen after filtering with $f_c = 1/2000y$.

This can be also understood in the spectral domain (B). The power spectral density of the signal (solid black line) increases towards lower frequencies whereas the power spectral density of the noise (dashed black line) is constant. For high cut-off frequencies such as $f_c = 1/10y$ the variance of the noise is higher than the variance of the signal as given by the integrated spectra of the signal and the noise up to f_c (area between axes, f_c and the power spectral density). For decreased f_c the integrated spectrum of the signal becomes more dominant relative to the noise. The higher amount of the signal compared to the noise component results in higher correlations for time series filtered with lower f_c .

estimated.

Standard routines for estimating the correlation of time series do not usually provide a possibility to make these estimates frequency dependent. Furthermore, it is impossible to apply these routines directly to irregular time series data. Instead, irregular data is typically resampled at a constant rate before applying the standard statistical routines. However, depending on the resampling method this may bias the analysis as spectral estimates may be affected by the resampling or interpolation method whose impact strongly increases with frequency (Broersen et al., 2000). Depending on the resampling frequency, this approach may add power on short time scales due to aliasing, which can only be diminished by using a higher resampling frequency (Broersen, 2009). Nearest neighbour interpolation results in a bias due to shifting the observation times to regular sampling (Broersen, 2009). Linear interpolation causes a variance reduction toward high frequencies of the spectrum (Schulz and Stettgeger, 1997). Thus, interpolation acts as a low-pass filter which causes a reddening of the original signal (Schulz and Stettgeger, 1997; Schulz and Mudelsee, 2002). This bias not only affects spectral analyses, but also time-domain methods (Schulz and Stettgeger, 1997).

There are a few types of techniques (mainly developed for spectral analyses) that can be directly applied to irregularly sampled data. Three approaches are shortly summarized by Broersen et al. (2000): (1) direct transform methods like the Lomb-Scargle method (Lomb, 1976; Scargle,

1989), (2) slotting methods (Edelson and Krolik, 1988; Tummers and Passchier, 1996), and (3) model-based estimators (van Maanen and Oldenziel, 1998; Müller et al., 1998). Rehfeld et al. (2011) tested these techniques for estimating the correlation of irregularly sampled time series including different kernel methods in the time domain such as slotting and non-rectangular kernels. While these techniques enable correlation estimates directly from irregularly sampled data, they do not consider possible time scale dependencies of the correlation.

In this study, we investigate methods to estimate the time scale dependent correlation of weakly irregularly sampled time series. As fast Fourier transform based estimators have been shown to be inferior for these kinds of analyses (Rehfeld et al., 2011), we focus on methods in the time domain and two aspects in particular: (i) the time scale dependency, and (ii) how best to handle the sampling irregularity.

While the first part focuses on finding methods that are sensitive to the time scale dependency in general, the second part evaluates their applicability to irregularly sampled time series. This is accomplished through surrogate time series that show the same characteristics as Holocene marine sediment temperature proxy records.

2. Methods

2.1. Time scale dependency

Since the correlation of time series is generally a function of time scale (i.e. the frequency range), we applied linear response filters to the data in the time domain before estimating the correlation. Mathematically, the filter process is defined as the convolution of an input time series with the impulse response function of the filter. In the frequency domain, a filter is defined by its transfer function given by the Fourier transform of the impulse response function. Commensurate with the filter effect in the time domain, the power spectrum of a filtered time series is given by the product of the power spectrum of the input time series and the squared transfer function.

In the following, the correlation at time scale T_c is defined as the correlation obtained from the low-pass filtered data with cut-off frequency $f_c (=T_c^{-1})$. Specifically, we tested finite response filters of varying complexity: (A) a moving average (boxcar) filter, common in paleoclimate research; (B) a Gaussian filter, as a compromise of a short filter both in the time and frequency domains; and (C) a finite response filter with a sharp cut-off in the frequency domain. The frequency response of the moving average filter depends on the filter length which we defined as half of the length of the time scale for which we analysed, and the statistical filter weights were set to the reciprocal filter lengths. The width parameter σ in the Gaussian filter, commonly known as the standard deviation of the Gaussian function, was set to

$$\sigma = \frac{1}{f_c \pi \sqrt{\frac{2}{\ln(2)}}}, \quad (1)$$

which ensures that the squared transfer function of the filter decreases to 0.25 at the cut-off frequency f_c . We set the filter length to 6σ as filter weights beyond $\pm 3\sigma$ are sufficiently close to zero and thus negligible.

The finite response filter consisted of the low-pass filter from Bloomfield (1976), where the weights depend on both the cut-off frequency and the filter length that defines the filter sharpness. A wider time window of the filter will cause the frequency response of the low-pass to converge toward a perfect cut-off filter in the frequency domain. The filter length was set to five times the time scale for which we analysed. An example of all three filter types is shown in Fig. 2.

2.2. Irregularity

We define an irregularly sampled time series of length N as a sequence of observations $x_i^{(irr)}$ whose observation times $t_i^{(irr)}$ increase monotonically with variable inter-observation time steps

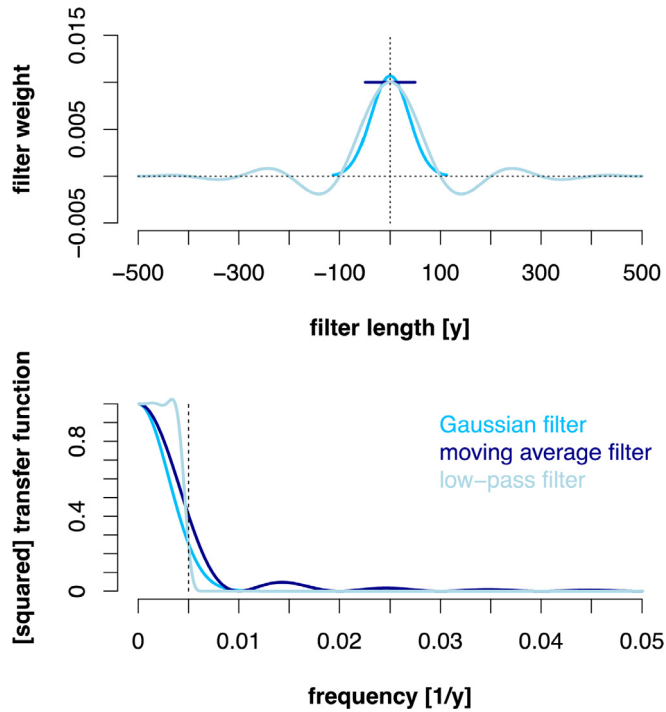


Fig. 2. Characteristics of the filters used in this study. Upper panel: Filter weights in the time domain. Lower panel: Squared transfer function in the frequency domain with marked cut-off frequency (dashed line). The low-pass filter is characterized by a sharp cut-off in the frequency domain, the moving average by its short length in the time domain. The Gaussian filter has the same simple shape in the time and the frequency domain. These examples are for a time scale of 200y.

$\Delta t_i = t_i^{(irr)} - t_{i-1}^{(irr)}$ and a mean sampling interval.

The time domain filters were applied at regular output positions $t_j^{(reg)}$ and different filtering procedures were tested: (1) direct filtering of the irregular time series, (2) an integrand interpolation method that interpolates the integrand of the convolution integral of the filter, and (3) an interpolation method that interpolates the irregular time series onto a regular sampling grid before filtering. The first approach filtered and resampled the irregular time series without interpolation and produced a filtered time series, \tilde{x}_j , at the regular sampling points $t_j^{(reg)}$ as

$$\tilde{x}_j = \frac{\sum_{i=1}^N W(t_j^{(reg)} - t_i^{(irr)}) \cdot x_i^{(irr)}}{\sum_{i=1}^N W(t_j^{(reg)} - t_i^{(irr)})} \quad (2)$$

where $W(t)$ are filter weights of our tested symmetric filters. The integrand interpolation method used the procedure from Eng and Gustafsson (2008), by which the irregular time series was filtered at regular positions $t_j^{(reg)}$ and the integrand of the filter-specific convolution integral $I_i^{(irr)}$ interpolated. We applied a linear interpolation, yielding a filtered time series

$$\tilde{x}_j = \sum_{i=1}^N \frac{I_i^{(irr)} + I_{i-1}^{(irr)}}{2} \Delta t_i \quad (3)$$

with the integrand of the convolution integral being

$$I_i^{(irr)} = W(t_j^{(reg)} - t_i^{(irr)}) x_i^{(irr)} \quad (4)$$

Particularly in the geosciences, irregularly sampled time series data are commonly interpolated using linear or nearest neighbour interpolations to produce a regular time series, $x_i^{(reg)}$ at $t_i^{(reg)}$, which is then filtered at the regular positions $t_j^{(reg)}$ to yield

$$\tilde{x}_j = \frac{\sum_{i=1}^N W(t_j^{(reg)} - t_i^{(reg)}) \cdot x_i^{(reg)}}{\sum_{i=1}^N W(t_j^{(reg)} - t_i^{(reg)})} \quad (5)$$

To identify the best method for estimating time scale dependent correlations between irregularly sampled time series, we tested each method on surrogate time series data (with known correlations) that were modelled on Holocene marine sediment temperature proxy records. We employed a Monte Carlo approach with 10 000 repetitions to generate pairs of irregularly sampled surrogate time series with a pre-defined correlation. They were resampled and filtered at regular positions with a sampling interval well below the mean sampling interval, before calculating the Pearson correlation with standard routines and comparing the result with the known correlations.

2.3. Surrogate data

2.3.1. Construction of surrogate signals

Climate signals are often dominated by low frequencies (Hasselmann, 1976). The variability in climate proxy time series on interannual to multimillennial time scales can be described as a power-law process whose power spectral density (PSD) is given by $PSD = f^{-\beta}$, with frequency f and spectral slope β (Huybers and Curry, 2006). For temperature related paleoclimate signals, typical values for β are in range of red noise between 0.3 and 1.6 (Huybers and Curry, 2006) with $\beta \approx 1$ for the Holocene (Laepple and Huybers, 2014). Due to various non-climate-related factors such as a varying sedimentation rate or different sampling protocols, each proxy record contains an unknown noise component (Laepple and Huybers, 2013), which, to a first degree, can be described as an uncorrelated white noise, equivalent to a power-law process with $\beta = 0$.

Our surrogate data was therefore constructed from pairs of annually resolved time series X and Y of length $N = 10\,000$, each containing a superposition of a climate signal S with a spectral slope β_S and independent random noise components ϵ_x and ϵ_y with a spectral slope β_ϵ :

$$X = S + c\epsilon_x \quad (6)$$

$$Y = S + c\epsilon_y$$

where c is a scaling factor. All components of both time series are uncorrelated, i.e. $cor(S, \epsilon_x) = cor(S, \epsilon_y) = cor(\epsilon_x, \epsilon_y) = 0$, and the associated variances σ_S , σ_{ϵ_x} and σ_{ϵ_y} are equal to 1.

The Pearson correlation of two time series is defined as their standardized covariance given by

$$\rho_{XY} = \frac{cov(X, Y)}{\sqrt{\sigma_X^2} \sqrt{\sigma_Y^2}} \quad (7)$$

As signal and noise are uncorrelated, the covariance of X and Y is equal to the variance of the signal component. In addition, the characteristics of the noise components are identical which yields identical scaling factors. This simplifies the correlation of the time series to

$$\rho_{XY} = \frac{\sigma_S^2}{\sigma_S^2 + \sigma_\epsilon^2 c^2} \quad (8)$$

The scaling factor c is used to scale the amplitudes of the noise components and thus to set the true correlation ρ_{orig} . To get an analytical expression of c , the correlation coefficient from Eq. (8) has to be solved for the scaling factor

$$c = \sqrt{\frac{\sigma_S^2}{\sigma_\epsilon^2} \left(\frac{1}{\rho_{orig}} - 1 \right)} \quad (9)$$

Finally, we normalize the time series by multiplying them with $\sqrt{\rho_{orig}}$ to get a unit variance.

Typical correlations between pairs of Holocene marine proxy records are low because the true climate signal is masked by strong non-climatic noise. Laepple and Huybers (2014) obtained mean correlations

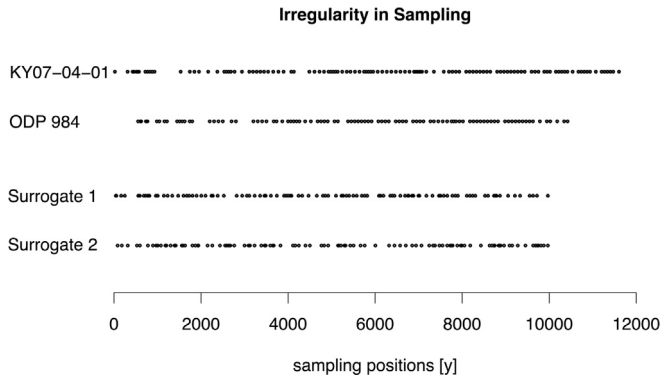


Fig. 3. Examples of sampling irregularity. The sampling times of two typical marine sediment cores KY07-04-01 (Kubota et al., 2010) and ODP 984 (Came et al., 2007) are shown in the upper two rows. The lower two rows show examples of surrogate sampling times generated using gamma-distributed inter-observation time steps as described in the Method section.

of 0.18 (for Mg/Ca proxy records) and 0.08 (for Uk37 proxy records) for marine sediment proxy record time series data collected less than 5000 km apart and for time scales greater than 500y.

2.3.2. Construction of irregular sampling

To obtain a realistic distribution of inter-observation time steps we characterized the properties of 56 marine temperature proxy time series from the Holocene dataset of Marcott et al. (2013). The inter-observation time steps of each time series were normalised by their respective means. The result was found to be gamma distributed and, using the method of moments, with both the shape and rate parameters equal to 2.225. The overall mean inter-observation time step between samples was 108.56y.

To mimic the irregular sampling in our surrogate time series X and Y , we generated gamma distributed individual sequences of inter-observation time steps for each time series, using the above shape and rate parameters, and multiplying the result with a mean sampling interval of 100y (Fig. 3). Each observation time $t_i^{(irr)}$ corresponds to the sum of all previous inter-observation time steps Δt_j : $t_i^{(irr)} = \sum_{j=1}^i \Delta t_j$. The subsampling of the annually sampled time series at a position $t_i^{(irr)}$ was accomplished by block averaging observations between $(t_i^{(irr)} - \Delta t_i/2)$ and $(t_i^{(irr)} + \Delta t_i/2)$. We consciously chose to use averages rather than to interpolate because particularly for marine data, samples often include adjacent depths or the sample distance is smaller than the typical mixing depth in the sediment (Berger and Heath, 1968).

2.4. Evaluation of the estimation methods

Our methods were evaluated by comparing the correlation of our irregularly sampled surrogate time series with the known true correlation using two types of targets: (i) the true time scale dependent correlation (that would be obtained from a perfect filter), and (ii) the filter-specific time scale dependent correlation (to account for a non-perfect frequency response of the actual filter). The true time scale dependent correlation can be determined by considering ρ_{XY} as a ratio of variances [cf., Eq. (8)]. To account for the time scale dependency, the variances of the power-law processes of both the signal and noise components can be calculated by integrating the power spectral density of the filtered power-law processes

$$\sigma^2 = \int_{N^{-1}}^{f_{Ny}} f^{-\beta} G df \quad (10)$$

The lower and upper boundaries of the integral are the lowest detectable frequency, i.e., the inverse time series length N^{-1} , and the Nyquist frequency, respectively. For a perfect cut-off filter, G is a step

function, defined as: $G=1$ ($f \leq f_c$) and $G=0$ ($f > f_c$). In the filter-specific case, G is equal to the squared transfer function which is the squared absolute value of the Fourier transform of the impulse response function of our tested filter (Fig. 2). In this study, we solved the above integral numerically.

The standard deviation and bias of the deviation from the targets was used to evaluate the quality of the estimators. Since the estimated time scale dependent correlation depends on the applied filter, we focussed on the bias relative to the filter-specific true time scale dependent correlation. For practical reasons we limit our analysis on time scales smaller than 1000y (1/10th of the time-series length). On longer time scales, the standard deviation of the correlation estimates is already higher than half of the correlation estimate which is caused by a small number of independent samples used for estimating the correlation.

3. Results

3.1. Correlation of red signal - white noise time series

We first analyse the case consisting of a red signal ($\beta_s = 1$) and white noise ($\beta_e = 0$) with predefined correlation ρ_{orig} so that the time scale dependent correlation (for time scale $> 200y$) is $\rho_{XY} = 0.1$. In this case, the true time scale dependent correlation increases for longer time scales (Fig. 4) as the relative influence of the noise is smaller at long time scales. Since the filters do not have a perfect cut-off in the frequency domain, the filter-specific true time scale dependent correlations slightly differ from the true time scale dependent correlation. While the moving average filter underestimates and the Gaussian filter overestimates the correlation for all analysed time scales, the low-pass filter produces correlations that are almost identical to the true time scale dependent correlation (Fig. 4). All methods exhibit an increase of the correlation estimates for longer time scales similar to the true correlations (Fig. 4), but differences between the estimated and the filter-specific true correlation remain.

On time scales shorter than 200y, all estimators overestimate the correlation, except for the direct filtering method using a low-pass filter which creates a strong bias. At these shorter time scales, the integrand interpolation method produced very similar results that are almost identical with the filter-specific true correlations, and only the moving average overestimated the true correlation (Fig. 4).

On longer time scales, the estimated correlations generally converge to the filter-specific true correlation as indicated by the decreasing absolute bias as the time scales are increased (Fig. 4), with the

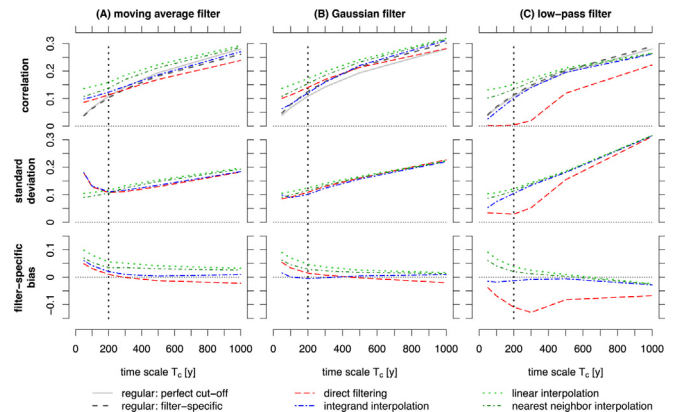


Fig. 4. Performance of the different estimation methods for a red signal ($\beta_s = 1$) with white noise component ($\beta_e = 0$). Top row: True correlation (grey, black) and estimated correlation (colours) as a function of time scale. Middle row: Standard deviation of the estimators. Bottom row: Bias of the estimators relative to the filter-specific true correlation. The dashed vertical line marks the time scale corresponding to the mean sampling interval.

exception of the low-pass filter and the direct filtering approach which both maintain a negative bias.

In general, the standard deviations of the estimators, describing the uncertainty of estimating the correlation from a pair of finite time series, are higher than the filter-specific biases. Due to the higher uncertainties of estimating the correlation on long time scales from short datasets, the standard deviations increase for longer time scales. With the exception of the moving average method on short time scales, all methods exhibit higher standard deviations and biases when interpolating compared to the direct filtering and integrand interpolation approach. The different uncertainties are mainly due to different filter lengths. In general, longer filters have higher standard deviations and lower filter-specific biases (except for parts of the low-pass based procedures) than short filters.

3.2. Correlation of white signal - white noise time series

To further investigate the properties of the estimators, we analyse the superposition of two white noise signals ($\beta_s = \beta_e = 0$) with a predefined correlation $\rho_{\text{orig}} = 0.25$. There is no predominance of one component at any frequency, so that the true correlation is time scale independent, which also holds for the filter-specific true correlations (Fig. 5). For all tested methods, the estimated correlations exhibit a negative bias that appears independent of both the procedure that handles the irregularity and the time domain filters used. The mismatch was high for short time scales and vice versa (Fig. 5).

The interpolation approach, and in particular the linear interpolation vis-à-vis the nearest-neighbour, yielded the best approximation of the true correlation.

On long time scales, the estimated correlations of the interpolation and the integrand interpolation procedure are similar (except for the moving average filter), while on short time scales the integrand interpolation method underestimates the filter-specific true correlation (Fig. 5). This behaviour stands in direct contrast to the results obtained for the red signal and white noise (see above).

With regard to the uncertainties in agreement with the results from the red signal with white noise, the standard deviations typically increased for longer time scales in all tested methods except the moving average filter on short time scales using the direct filtering and integrand interpolation approach.

The standard deviation of an estimator tends to be larger than the filter-specific bias. With longer filters and for long time scales, the bias

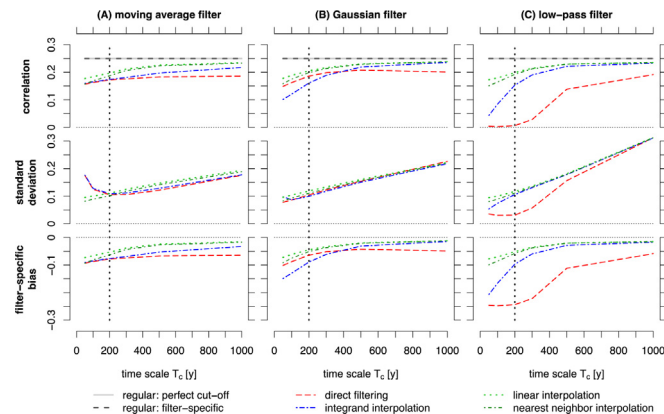


Fig. 5. Performance of the different estimation methods for a white noise signal ($\beta_s = 0$) with white noise component ($\beta_e = 0$). Top row: True correlation (grey, black) and estimated correlation (colours) as a function of time scale. Middle row: Standard deviation of the estimators. Bottom row: Bias of the estimators relative to the filter-specific true correlation. Vertical dashed lines mark the time scale corresponding to the mean sampling interval. As the spectrum of the white noise signal and noise is constant, the true correlation is time scale independent.

was typically lower while those values were nearly equal for short time scales. The observed high absolute bias of the integrand interpolation method for short time scales stands in direct contrast to the low bias observed for the red signal plus white noise scenario.

4. Discussion

We focussed on three particular aspects: (1) how does temporal irregularity in sampling affect the time scale dependency of correlations, (2) what is the best procedure to handle this irregularity, and (3) which is the best filter to account for the time scale dependence of the correlation. We tested three different approaches to find correlations in irregularly sampled time series data, namely: (i) direct filtering of irregular time series, (ii) an integrand interpolation method that interpolates the integrand of the convolution integral of the filter, and (iii) an interpolation that interpolates the irregular time series onto a regular sampling grid before filtering. The approaches were combined with three finite response filters: (A) a moving average (boxcar) filter, (B) a Gaussian filter, and (C) a finite response filter with a sharp cut-off in the frequency domain. While all tested approaches yielded similar results, there were some notable differences.

4.1. Effect of irregularity and non-simultaneousness in sampling

The irregularity of the sampling affects the correlation estimate in two ways: (i) they shift the correlation toward estimates related to longer time scales due to different inter-observation time steps; and (ii) the non-simultaneous sampling of time series introduces a general bias that leads to an underestimation of the true correlation. Fig. 6 illustrates these effects in case of a time scale dependent ($\beta_s = 1, \beta_e = 0$) as well as a time scale independent ($\beta_s = 0, \beta_e = 0$) correlation.

The sampling rate, Δt^{-1} , of a time series acts qualitatively as a first

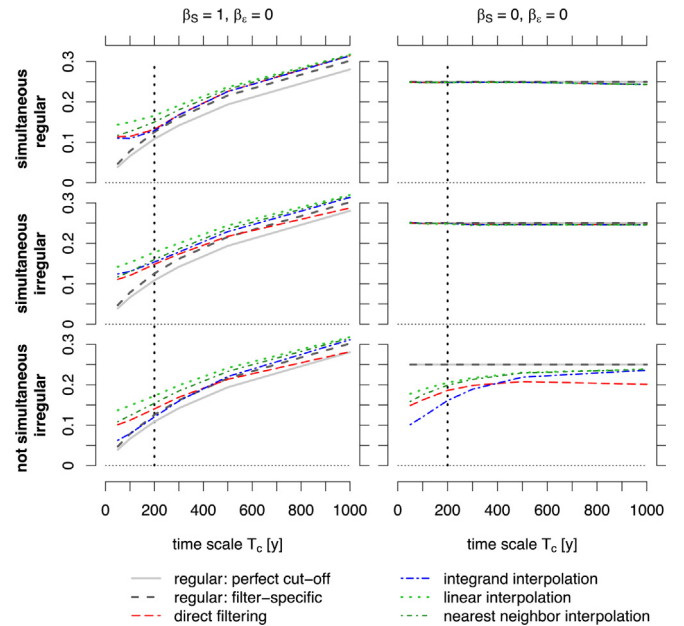


Fig. 6. Effect of irregular and non-simultaneous sampling on the correlation estimates. Correlations are estimated by applying a Gaussian filter on regularly (top), irregularly but simultaneously (mid), and irregularly but non-simultaneously (bottom) sampled time series as a superposition of a red signal and white noise (left), and white noise signal and white noise (right). The vertical dashed lines mark the time scale corresponding to the mean sampling interval. The simultaneous, irregular sampling is affected by a shift of the correlation estimates toward longer time scales, which is not visible in case of the white signal and noise (right). Sampling that is irregular but non-simultaneous will underestimate the correlation.

low-pass filter with a smoothed cut-off. Since it is necessary to have at least two observations per detectable oscillation, the sampling-related cut-off (i.e., Nyquist) frequency is $f_{Ny} = 0.5\Delta t^{-1}$. To analyse the time scale dependency of the correlation, we filtered the time series a second time to remove information about oscillations smaller than the considered time scale. If the first filtering process has a smaller cut-off frequency than the second one, there is an additional loss of information at frequencies $>f_{Ny}$ (Fig. 6, top left). In the presence of a red signal, the time scale dependent correlation of a pair of time series would thus be overestimated while a white noise signal would exhibit the expected time scale independency of the correlation (Fig. 6, top right).

For two time series that were sampled at irregular intervals but simultaneous times, the short time scale correlation estimates are shifted toward those for longer time scales, which is a direct result of the irregular sampling intervals and the sampling-related cut-off becomes roughly $f_{Ny} = 0.5(\Delta t_{max})^{-1}$. Therefore, the shift only affects time scales $<2\Delta t_{max}$. If $f_c > f_{Ny}$, the correlation is overestimated for red signals while remaining unchanged for white noise (Fig. 6, middle row).

Sampling two irregular time series at non-simultaneous times removes information at different time scales and different time windows from each data set. The correlation estimate of such time series is affected by both the shift of short time scale related correlations toward the estimate at a time scale of $2\Delta t_{max}$ and an underestimation that affects all time scales. White data sets are only affected by the non-simultaneity in sampling. The shift of short time scale related correlations toward those observed for longer time scales we did not observe as the correlation of a white noise signal is time scale independent (Fig. 6, bottom right). Both effects are seen for the data set containing the red signal (Fig. 6, bottom left). In this case, the good approximation of the estimated correlation of the integrand interpolation on the filter-specific true correlation is an artefact which is caused by the error-proneness of the method in application on time series with a low resolution and non-simultaneous sampling.

4.2. Choosing the best method

A procedure that should handle irregularity and account for time scale dependencies must above all be robust and reliable.

4.2.1. Handling irregularity

If Δt_{max} exceeds the applied filter length, the filter integral becomes unreliable and the fewer observations we have within the filter range the higher the associated error. Direct filtering is similar to the numerical solution of the filter integral for a high resolution regularly sampled time series which is equal to the sum of the products of the filter weights with their corresponding observations. Due to the irregular sampling combined with a low mean sampling rate, there are only a few observations left which are used for the sum. There is thus a systematic error, caused by the application of this procedure on irregularly sampled time series. This error affects all time scales, but especially shorter ones. This effect was clearly visible with time series pairs whose correlation was time scale independent (Fig. 5).

The filter integrand interpolation procedure commonly reduces the errors caused by direct filtering, because it also works directly on the irregularly sampled time series, but uses additional interpolated values during the summation of the products of the filter weights with their corresponding observations. Again, a higher number of observations per filter window reduces the systematic error of this approach because of less over- or underestimated interpolated values of the filter integral. Hence, the error is lower on longer time scales. On short time scales, there is a high bias despite filter integrand interpolation (Fig. 5).

4.2.2. Accounting for time scale dependency

The applied filters cause artefacts based to their length as well as their cut-off behaviour. For Δt greater than the applied filter length, gaps occurred in the filtered result. According to the filter weights,

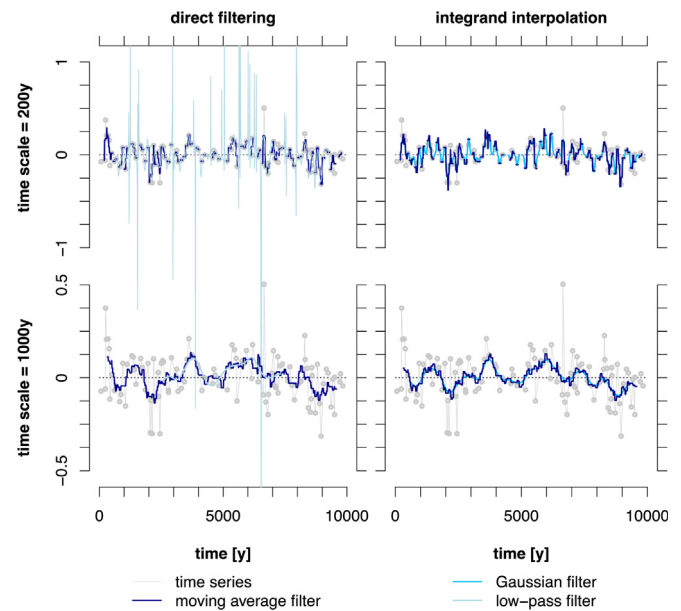


Fig. 7. Filter artefacts related to different time scales caused by filtering an irregularly sampled time series. Short time scales (top) are affected by gaps (moving average filter) for filter lengths shorter than the inter-observation time steps. They also tend to overshoot when there are fewer observations covered by positive and negative filter weights (low-pass). The analysis of longer time scales (bottom) minimizes the occurrence of artefacts due to a higher number of observations covered by the filter windows.

which are related to the cut-off behaviour of the applied filter, there were artefacts like overshooting and steps within the filtering results (Fig. 7).

The moving average filter had the shortest length of all tested filters. There were less observations within the filter window so that the standard deviation of the estimated time scale dependent correlations was lower than for the other filters (Fig. 4). Higher standard deviations on short time scales which are related to the direct filtering and integrand interpolation procedures are the result of gaps within the filtering results. These gaps occur if the filter length is shorter than the inter-observation time steps. Especially on short time scales, this happened frequently (Fig. 7, top). Since the weights of the moving average filter were equal, the filtering result appeared step-like. This is independent of the analysed time scale (Fig. 7).

The low-pass filter was the longest tested filter and contained positive and negative filter weights. Both properties can lead to artefacts. The underestimation of the correlation related to longer time scales might be the consequence of the long filter length as well as the finite length of the time series itself. Especially short time series would suffer from a significant loss of information due to filtering with low cut-off frequencies, so that the filtered time series would be much shorter than the original data (Fig. 7, bottom left).

In addition, this caused a high standard deviation of the correlations for the low-pass filter. On short time scales, the low correlation estimates related to the direct filtering and integrand interpolation procedures resulted from filter artefacts (overshooting) due to less observations as well as positive and negative filter weights. Especially in parts of the time series where the inter-observation time steps were larger than the mean sampling interval, the filtered time series was affected by overshooting (Fig. 7, top left). The overshooting can be reduced by increasing the number of observations per filter window or through a more regular sampling.

The Gaussian filter had nearly the same characteristics as the low-pass filter from Bloomfield (1976) with a filter length of one-time the analysed time scale because both filters only contain positive filter weights. Therefore, here the main artefact of the low-pass filter – the

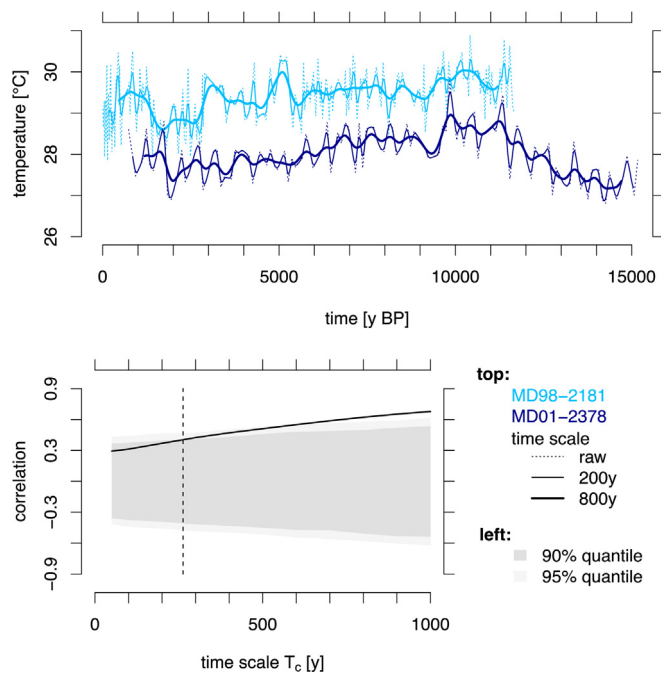


Fig. 8. Filtering results and time scale dependent correlation for two marine sediment temperature proxy records. The upper panel shows the raw and filtered data of both time series using a cut-off frequency of $1/200$ and $1/800 \text{ y}^{-1}$. The lower panel shows the time scale dependent correlation (black line) as well as the 90% (grey polygon) and 95% (light grey polygon) quantiles of the null hypothesis of noise with $\beta = 1$. The vertical dashed line marks the time scale corresponding to the maximum mean sampling interval of both time series.

overshooting which is caused by negative filter weights – was limited or non-existent (Fig. 7, top right). Sometimes, slight overshooting (not shown) was possible on short time scales when Δt was close to the Gaussian filter length.

4.3. Example application to observed proxy records

As a final example, we applied one combination of our tested procedures and filters to data from two marine sediment cores which consisted of sea surface temperature records, based on Mg/Ca measured in planktonic foraminifera (Nürnberg et al., 1996). The Morotai Basin core MD98-2181 (Stott et al., 2007), located in the westernmost Pacific Ocean (6.3°N , 125.8°E), and the Timor Sea core MD01-2378 (Xu et al., 2008), located in the easternmost Indian Ocean (13.1°S , 121.8°E), were collected 2190 km apart, have a high mean resolution (MD98-2181: 50y, MD01-2378: 130y) and overlap in time between 730 and 11 660y BP (Fig. 8).

We use the combination of a linear interpolation followed by Gaussian filtering to estimate the correlation related to different time scales, as this method produced the fewest artefacts. To test the statistical significance of the correlation estimates, we compared the estimates against the null hypothesis of uncorrelated time series with a noise with power-law scaling with $\beta = 1$ (Appendix A).

The correlation of the time series pair increased for longer time scales. Given the relatively high noise level in Mg/Ca records (Laepplé and Huybers, 2013), this suggests that the common (climate) signal is masked on short time scales. Filtering the time series led to a greater reduction in the noise component compared to the common climate component and resulted in statistically significant correlation estimates ($p = 0.05$) for time scales $> 500\text{y}$ (Fig. 8).

5. Conclusion

In this study, we investigated different methods to estimate the time scale dependent correlation of weakly irregularly sampled time series. Each method was restricted to the time domain and overcame the sampling irregularities by interpolation or direct filtering. Each method also had to account for the time scale dependency of the correlation which was achieved through filtering. As data sets we used surrogate time series that mimicked the characteristics of Holocene marine sediment temperature proxy records.

The irregularity in sampling causes problems during the correlation analyses as it qualitatively acts as a low-pass filter with a cut-off frequency of roughly $0.5(\Delta t_{\max})^{-1}$. Thus, estimates for time scales $< 2\Delta t_{\max}$ exhibited a bias in addition to a time scale independent underestimation of the correlation which we attributed to the non-simultaneous sampling of time series.

Both the linear and nearest neighbour interpolations were the best performers in terms of handling the sampling irregularity and neither was prone to artefacts which occur if some Δt are much larger than the mean sampling interval. Although the bias and the standard deviation were both higher for the interpolation method than for the integrand interpolation, we suggest the application of the interpolation method due to more reliable estimates as the integrand interpolation method failed on short time scales due to the existence of artefacts.

The tested low-pass filter with a length five times the analysed time scale provided the best approximation to the analytical solution for a perfect cut-off filter. However, the filtering result was affected by artefacts due to the long filter length as well as negative filter weights. Both aspects limit the applicability of this filter.

The Gaussian filter was a compromise between the low standard deviation of the moving average filter and the cut-off behaviour of the low-pass filter. The cut-off behaviour and filter weights were comparable to the properties of a shorter version of the low-pass filter of the same definition, but with a length of one-time the analysed time scale. By reason of a smoother filtering result without gaps (moving average filter) and overshooting (low-pass filter), we suggest the application of the Gaussian filter.

Irrespective of the applied method and filter, there was the tendency for most artefacts to occur when analysing short time scales if there were less observations due to a low mean sampling rate. Therefore, it is meaningful to restrict the estimation of the time scale dependent correlation to time scales which are obviously larger than the mean inter-observation time step of the analysed time series.

Computer code availability

The code used in this paper is based on R ($\geq 3.2.2$). It can be downloaded as the R package ‘corit’ from the public git repository at <https://github.com/EarthSystemDiagnostics/corit>. The main functions allow a time scale dependent estimation of the Pearson correlation based on three methods resampling irregular data to an equidistant spacing before or during filtering the time series, a significance test of the estimate based on surrogate records as well as an output of the filtered time series. All main functions are provided with a simple example in the R Documentation.

Acknowledgement

We would like to acknowledge Jürgen Kurths, Kira Rehfeld, Andrew Dolman, and Igor Kröner for fruitful discussions. This study was supported by the Initiative and Networking Fund of the Helmholtz Association Grant VG-NH900 as well as the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement no. 716092).

Appendix A. Significance test for time scale dependent correlation estimates

To test the significance of the time scale dependent correlation estimates, we resort to a Monte Carlo procedure. We generate surrogate records on an annual scale, subsample the records at the same resolution as the analysed time series and then apply our estimators. Finally, we provide the 90% and 95% quantiles of the correlations obtained from the surrogate data.

As null-hypothesis we assume uncorrelated time series, thus time series that only consist of noise. The statistical significance will depend on the temporal correlation of the noise, i.e., its power spectrum. Let us assume a power-law spectrum, either with a prescribed slope, or with a slope estimated from the datasets to be tested.

For the latter, we estimate the spectral slopes of our time series by linearly interpolating the time series onto an equidistant grid and estimating the power spectrum from the detrended time series. Finally, we estimate the slope from the power spectrum by a linear regression in the frequency range of the reciprocal half of the length of the overlapping time window of the time series as well as the inverse twofold maximum mean resolution of both time series. Our choice of the frequency range minimized the impact of the detrending as well as the impact of the irregular sampling and the linear interpolation. However, especially for short time series, the estimates of spectral slopes are very uncertain.

For Holocene sea surface temperature paleoclimate time series, we suggest to prescribe a spectral slope of 1 (Laepple and Huybers, 2014) as this estimate is based on analysing a large number of sediment records and thus more robust than noisy estimates based on single time series.

References

- Berger, W.H., Heath, G.R., 1968. Vertical mixing in pelagic sediments. *J. Mar. Res.* 26, 134–143.
- Bloomfield, P., 1976. *Fourier Analysis of Time Series: an Introduction*. John Wiley, New York.
- Broersen, P.M., 2009. Five separate bias contributions in time series models for equidistantly resampled irregular data. *IEEE Trans. Instrum. Meas.* 58 (5), 1370–1379.
- Broersen, P.M.T., de Waele, S., Bos, R., 2000. The accuracy of time series analysis for laser-Doppler velocimetry. In: *Proceedings of the 10th International Symposium on Applications of Laser Techniques to Fluid Dynamics*, Lisbon, Portugal.
- Came, R.E., Oppo, D.W., McManus, J.F., 2007. Amplitude and timing of temperature and salinity variability in the subpolar North Atlantic over the past 10 k.y. *Geology* 35 (4), 315–318. <https://doi.org/10.1130/G23455A.1>.
- Edelson, R.A., Krolik, J.H., 1988. The discrete correlation function - a new method for analyzing unevenly sampled variability data. *Astrophys. J.* 333, 646–659.
- Eng, F., Gustafsson, F., 2008. Downsampling non-uniformly sampled data. *EURASIP J. Appl. Signal Process.* <https://doi.org/10.1155/2008/147407>.
- Hasselmann, K., 1976. Stochastic climate models Part I. Theory, *Tellus* 28 (6), 473–485.
- Huybers, P., Curry, W., 2006. Links between annual, Milankovitch and continuum temperature variability. *Nature* 441 (7091), 329–332. <https://doi.org/10.1038/nature04745>.
- Kubota, Y., Kimoto, K., Tada, R., Oda, H., Yokoyama, Y., and Matsuzaki, H., 2010. Variations of East Asian summer monsoon since the last deglaciation based on Mg/Ca and oxygen isotope of planktic foraminifera in the northern East China Sea. *Paleoceanography*, 25, PA4205. doi:10.1029/2009PA001891.
- Laepple, T., Huybers, P., 2013. Reconciling discrepancies between UK37 and Mg/Ca reconstructions of Holocene marine temperature variability. *Earth Planet Sci. Lett.* 375, 418–429.
- Laepple, T., Huybers, P., 2014. Ocean surface temperature variability: large model–data differences at decadal and longer periods. *Proc. Natl. Acad. Sci. Unit. States Am.* 111 (47), 16682–16687. <https://doi.org/10.1073/pnas.1412077111>.
- Lomb, N.R., 1976. Least-squares frequency analysis of unequally spaced data. *Astrophys. Space Sci.* 39 (2), 447–462.
- Marcott, S.A., Shakun, J.D., Clark, P.U., Mix, A.C., 2013. A reconstruction of regional and global temperature for the past 11,300 years. *Science* 339 (6124), 1198–1201. <https://doi.org/10.1126/science.1228026>.
- Müller, E., Nobach, H., Tropea, C., 1998. Model parameter estimation from non-equidistant sampled data sets at low data rates. *Meas. Sci. Technol.* 9 (3), 435–441.
- Nürnberg, D., Bijma, J., Hemleben, C., 1996. Assessing the reliability of magnesium in foraminiferal calcite as a proxy for water mass temperatures. *Geochim. Cosmochim. Acta* 60 (5), 803–814.
- Rehfeld, K., Marwan, N., Heitzig, J., Kurths, J., 2011. Comparison of correlation analysis techniques for irregularly sampled time series. *Nonlinear Process Geophys.* 18 (3), 389–404. <https://doi.org/10.5194/npg-18-389-2011>.
- Scargle, J.D., 1989. Studies in astronomical time series analysis. III - Fourier transforms, autocorrelation functions, and cross-correlation functions of unevenly spaced data. *Astrophys. J.* 343, 874–887.
- Schulz, M., Mudelsee, M., 2002. REDFIT: estimating red-noise spectra directly from unevenly spaced paleoclimatic time series. *Comput. Geosci.* 28 (3), 421–426.
- Schulz, M., Stattegger, K., 1997. SPECTRUM: spectral analysis of unevenly spaced paleoclimatic time series. *Comput. Geosci.* 23 (9), 929–945.
- Stott, L., Timmermann, A., Thunell, R., 2007. Southern hemisphere and deep-sea warming led deglacial atmospheric CO₂ rise and tropical warming. *Science* 318 (5849), 435–438. <https://doi.org/10.1126/science.1143791>.
- Tummers, M.J., Passchier, D.M., 1996. Spectral estimation using a variable window and the slotting technique with local normalization. *Meas. Sci. Technol.* 7 (11), 1541–1546.
- van Maanen, H.R.E., Oldenziel, A., 1998. Estimation of turbulence power spectra from randomly sampled data by curve-fit to the autocorrelation function applied to laser-Doppler anemometry. *Meas. Sci. Technol.* 9 (3), 458–467.
- Xu, J., Holbourn, A., Kuhnt, W., Jian, Z., Kawamura, H., 2008. Changes in the thermocline structure of the Indonesian outflow during Terminations I and II. *Earth Planet Sci. Lett.* 273 (1), 152–162.