



REDFIT: estimating red-noise spectra directly from unevenly spaced paleoclimatic time series[☆]

Michael Schulz^{a,*}, Manfred Mudelsee^b

^a*Institut für Geowissenschaften, Universität Kiel, Olshausenstr. 40, D-24118 Kiel, Germany*

^b*Institut für Meteorologie, Universität Leipzig, Stephanstr. 3, D-04103 Leipzig, Germany*

Received 20 May 2000; received in revised form 27 March 2001; accepted 2 April 2001

Abstract

Paleoclimatic time series are often unevenly spaced in time, making it difficult to obtain an accurate estimate of their red-noise spectrum. A Fortran 90 program (REDFIT) is presented that overcomes this problem by fitting a first-order autoregressive (AR1) process, being characteristic for many climatic processes, directly to unevenly spaced time series. Hence, interpolation in the time domain and its inevitable bias can be avoided. The program can be used to test if peaks in the spectrum of a time series are significant against the red-noise background from an AR1 process. Generated and paleoclimatic time series are used to demonstrate the capability of the program. © 2002 Elsevier Science Ltd. All rights reserved.

Keywords: Spectral analysis; Irregular sampling intervals; Lomb–Scargle Fourier transform

1. Introduction

Spectral analysis is an important tool in climate research because it allows the variance of a time series to be separated into contributions associated with different time scales. It thus helps to understand better the physical processes, which generate the variability recorded in a time series. Spectra of paleoclimatic time series frequently show a continuous decrease of spectral amplitude with increasing frequency (“red-noise”). Hasselmann (1976) demonstrated that a first-order autoregressive (AR1) process is sufficient to explain this climatic red-noise signature. Accordingly, the AR1 model is often used as null hypothesis to assess whether or not the variability recorded in a time series is consistent with a stochastic origin of this type (Gilman et al., 1963). Such a test involves estimation of an AR1

parameter from the time series under consideration. For evenly sampled time series this is a relatively straightforward procedure (e.g. Percival and Walden, 1993). However, most paleoclimatic time series are unevenly spaced (i.e., intervals between sampling times are not constant), and the application of estimation techniques for evenly spaced time series would require some sort of interpolation. Unfortunately, this procedure results in a significant bias because interpolation in the time domain alters the estimated spectrum of a time series by enhancing the low-frequency components at the expense of high-frequency components. That is, the estimated spectrum of an interpolated time series becomes too “red” compared to the true spectrum (e.g. Schulz and Stattegger, 1997).

We present a computer program which estimates the AR1 parameter directly from unevenly spaced time series, that is, without requiring interpolation. The estimated AR1 model is then transformed from the time domain into the frequency domain. Comparison of the spectrum of the time series with that of the AR1 model allows to test the hypothesis that the time series originates from an AR1 process. Following a brief

[☆]Code available from server at <http://www.iamg.org/CGEditor/index.htm>

*Corresponding author.

E-mail addresses: mschulz@email.uni-kiel.de (M. Schulz), mudelsee@rz.uni-leipzig.de (M. Mudelsee).

description of the numerical procedure and its implementation in a computer program, we apply the program to a synthetic time series and a paleoclimatic record.

2. Method

A discrete AR1 process r for times t_i ($i = 1, 2, \dots, N$) with arbitrary spacing is given by (Robinson, 1977)

$$r(t_i) = \rho_i r(t_{i-1}) + \varepsilon(t_i), \quad (1)$$

$$\rho_i = \exp(- (t_i - t_{i-1})/\tau).$$

The constant τ is the characteristic time scale of the AR1 process (a measure of its “memory”) and ε indicates “white” Gaussian noise with zero mean and variance $\sigma_\varepsilon^2 \equiv 1 - \exp(-2(t_i - t_{i-1})/\tau)$. This value of σ_ε^2 ensures that the AR1 process is stationary and has unit variance. The spectrum $G_{rr}(f_j)$ corresponding to the time-domain process of Eq. (1) is (e.g. Percival and Walden, 1993)

$$G_{rr}(f_j) = G_0 \frac{1 - \rho^2}{1 - 2\rho \cos(\pi f_j / f_{Nyq}) + \rho^2} \quad (2)$$

where f_j denotes discrete frequency up to the Nyquist frequency f_{Nyq} (cf. Schulz and Stettegger, 1997) and G_0 is the average spectral amplitude. The “average autocorrelation coefficient” ρ is calculated from the arithmetic mean of the sampling intervals $\Delta t = (t_N - t_1)/(N - 1)$ as $\rho \equiv \exp(-\Delta t/\tau)$.

The unknown value of τ is estimated from an unevenly spaced time series using the least-squares algorithm devised by Mudelsee (2002). The spectrum of an irregularly spaced time series is determined without the need for interpolation by means of the Lomb–Scargle Fourier transform (Lomb, 1976; Scargle, 1982, 1989). Schulz and Stettegger (1997) presented a computer program for this purpose which makes additional use of the so-called Welch-overlapped-segment-averaging (WOSA) procedure (Welch, 1967). This algorithm splits a time series into n_{50} segments which overlap by 50%, the final spectral estimate is derived from averaging the n_{50} periodograms.

With an estimate for τ as well as an appropriate value for G_0 it should then be possible to overlay the red-noise spectrum after Eq. (2) and the spectrum estimated from the data. Provided that the probability distribution of G_{rr} at each frequency follows a χ^2 distribution (e.g. Percival and Walden, 1993), it is finally possible to test if the data spectrum is consistent with a red-noise model. Unfortunately, this approach is hampered by an inherent aspect of the Lomb–Scargle Fourier transform: in contrast to the classical Fourier transform, the individual Lomb–Scargle Fourier components are not necessarily independent of each other and, as a

consequence, an estimated spectrum based on the Lomb–Scargle transform may be biased (Lomb, 1976; Scargle, 1982). In particular, spectral amplitudes at the high-frequency end of a spectrum are often overestimated. Therefore, a red-noise spectrum Eq. (2) which is based on an unbiased estimate of τ for a given time series will not necessarily coincide with the “Lomb–Scargle spectrum” of the same time series. We therefore seek for a bias correction for the Lomb–Scargle Fourier transform.

3. Numerical procedure

The systematic deviation between a theoretical red-noise spectrum Eq. (2) and one estimated from an unevenly spaced time series by means of the Lomb–Scargle Fourier transform depends on the distribution of the sampling times in the interval $[t_1, t_N]$ (Lomb, 1976; Scargle, 1982). For some arbitrary distribution of sampling times the lack of an analytical solution for the deviation prevents a direct bias correction of a Lomb–Scargle spectrum. To circumvent this problem, we turn to a Monte–Carlo technique. Based on the actual sampling times, an ensemble of N_{sim} AR1 time series is generated after Eq. (1) with fixed τ . The deviation of the average spectrum of the ensemble from the known theoretical spectrum is then employed for the required bias correction. The computational steps to obtain a red-noise spectrum of an unevenly spaced time series $x(t_i)$ which is consistent with the estimated value of τ are as follows:

1. Estimate τ from $x(t_i)$ using the time-domain algorithm of Mudelsee (2002). If more than one WOSA segment is used for spectral analysis ($n_{50} > 1$), an average value for τ is calculated from τ estimates for each individual segment. The individual τ estimates (Mudelsee, in press) are bias corrected, based on the number of data points in each WOSA segment.

2. Estimate spectrum $\hat{G}_{xx}(f_j)$ of $x(t_i)$ in the interval $[0, f_{Nyq}]$ using the Lomb–Scargle Fourier transform as described in Schulz and Stettegger (1997). Determine the area under $\hat{G}_{xx}(f_j)$ which is an estimate for the variance of $x(t_i)$.

3. Monte Carlo simulation loop.
Repeat N_{sim} times

- create AR1 time series according to Eq. (1), using the sampling times of the input data (t_i), the estimated τ , and an independent set of $\varepsilon(t_i)$ for each simulation
- estimate spectrum of the generated AR1 time series, $\hat{G}_{rr}(f_j)$

- scale $\hat{G}_{rr}(f_j)$ such that the area under the spectrum is identical to the area under $\hat{G}_{xx}(f_j)$

Determine arithmetic mean of the N_{sim} independent red-noise spectral estimates $\langle \hat{G}_{rr}(f_j) \rangle$.

4. Calculate theoretical AR1 spectrum $G_{rr}(f_j)$ for the estimated value of τ (Eq. (2)). (Note that $G_{rr}(f_j)$ is not affected by the bias of the Lomb–Scargle Fourier transform, because the critical parameter τ is estimated in the time domain.)

5. Select G_0 (see Eq. (2)) such that the area under $G_{rr}(f_j)$ is identical to the area under $\hat{G}_{xx}(f_j)$. (This step is required since the true noise variance of the process under consideration is unknown.)

6. Calculate a correction factor $c(f_j)$ for the bias adjustment of the Lomb–Scargle spectrum as $c(f_j) = \langle \hat{G}_{rr}(f_j) \rangle / G_{rr}(f_j)$.

7. Using $c(f_j)$, determine a bias-corrected version of the spectrum of the data as $\hat{G}'_{xx}(f_j) = \hat{G}_{xx}(f_j) / c(f_j)$.

8. For assessing the statistical significance of a spectral peak, the upper confidence interval of the AR1 noise is calculated for various significance levels (based on χ^2

distribution; degrees of freedom depend on the actual spectral analysis setting; cf. Schulz and Stattegger, 1997). In addition, significance levels are calculated from percentiles of the Monte Carlo ensemble.

9. Check appropriateness of the AR1 model to describe $x(t_i)$ by testing the equality of $G_{rr}(f_j)$ and $\hat{G}'_{xx}(f_j)$ using a non-parametric runs test (Bendat and Piersol, 1986).

The assumptions underlying this procedure are: (i) The noise background recorded in a time series can indeed be approximated by an AR1 process (tested in step 9), that is, the potential effect of non-AR1 signal components (e.g. harmonic signals) can be neglected. Although it would be possible to identify and subtract harmonic signal components prior to estimating τ (see Mann and Lees, 1996 for evenly spaced time series), this approach may fail if there are quasi-periodic signals (e.g. narrow-band noise), which often occur in climatic time series. For most practical problems such refinement is unwarranted because such signals cover only a small portion of the entire frequency range and have only a marginal effect on the estimated value of τ (Gilman et al., 1963). Situations in which non-AR1 features do affect




Fig. 1. Red-noise spectrum of synthetic AR1 data. Unevenly spaced AR1 time series (A) generated according to Eq. (1) with $\tau = 15$ yr. (B) Theoretical red-noise spectrum $G_{rr}(f_j)$ based on estimated value of τ (thick solid line). Lomb–Scargle spectrum of time series $\hat{G}_{xx}(f_j)$ (thin dashed line; $n_{50} = 1$; rectangular window) and average of $N_{sim} = 1000$ simulated red-noise spectra $\langle \hat{G}_{rr}(f_j) \rangle$ (thick dashed line) deviate from expected shape of $G_{rr}(f_j)$, especially for $f > 0.09$ (1/yr). Correcting for this bias, inherent to spectral estimates of unevenly spaced data, results in spectrum of time series $\hat{G}'_{xx}(f_j)$ (thin solid line) which is consistent with $G_{rr}(f_j)$. Note that spectral amplitudes are plotted on logarithmic decibel [dB] scale.

the estimation of τ can be identified by visual inspection of the resulting red-noise spectrum and the runs test of step 9. (ii) The distribution of data points along the time axis is not too clustered (Horne and Baliunas, 1986).

A computer program (REDFIT) that performs the above steps is freely available via anonymous ftp from ftp.rz.uni-kiel.de (file: /pub/sfb313/mschulz/redfit35.zip) or from the IAMG server. The zip-archive includes Fortran 90 source code, binaries for Windows 95 (or above), program documentation and example files. The program offers the same functionality for univariate spectral analysis as the SPECTRUM program (Schulz and Statterger, 1997) and uses the same format for input files. To cope better with the computational demand of the Monte-Carlo simulation, the program is command-line driven and can therefore be run in batch mode.

4. Example computations

The first test signal is a pure AR1 process after Eq. (1) with $\tau = 15$ yr and $N = 324$ data points (Fig. 1A). The uneven time axis is generated by treating the time interval between subsequent sampling times as a random

variable following a gamma distribution with 3 degrees of freedom (which is a geologically realistic model; Schulz and Statterger, 1997). The estimated value for τ is 15 yr (90% confidence interval: $10 < \tau < 20$ yr). The uncorrected Lomb–Scargle spectrum of the AR1 time series, $\hat{G}_{xx}(f_j)$, does not show the characteristic red-noise shape, instead spectral amplitudes increase slightly for $f > 0.09$ (1/yr) (Fig. 1B). As expected, the same holds true for the mean, $\langle \hat{G}_{rr}(f_j) \rangle$, of the $N_{sim} = 1000$ simulated red-noise spectra (Fig. 1B). Compared to the theoretical spectrum of the generated AR1 process, $G_{rr}(f_j)$, (based on estimated value of τ) the Lomb–Scargle Fourier transform clearly overestimates the spectral amplitudes for a large part of the spectrum (Fig. 1B). Applying the bias correction (steps 6 and 7) results in a spectral estimate $\hat{G}'_{xx}(f_j)$ which is, of course, consistent with $G_{rr}(f_j)$ (Fig. 1B). At the low-frequency end of the spectrum we observe that $\hat{G}'_{xx}(f_j) > \hat{G}_{xx}(f_j)$. This effect is caused by the finite length of the time series, which leads to an underestimation of the spectral amplitudes for periods exceeding the length of the time series (independently of the spectral-analysis technique being used and the spacing of the time axis). Thus, as a side effect, the bias correction accounts also for this problem inherent in all spectral analysis techniques.

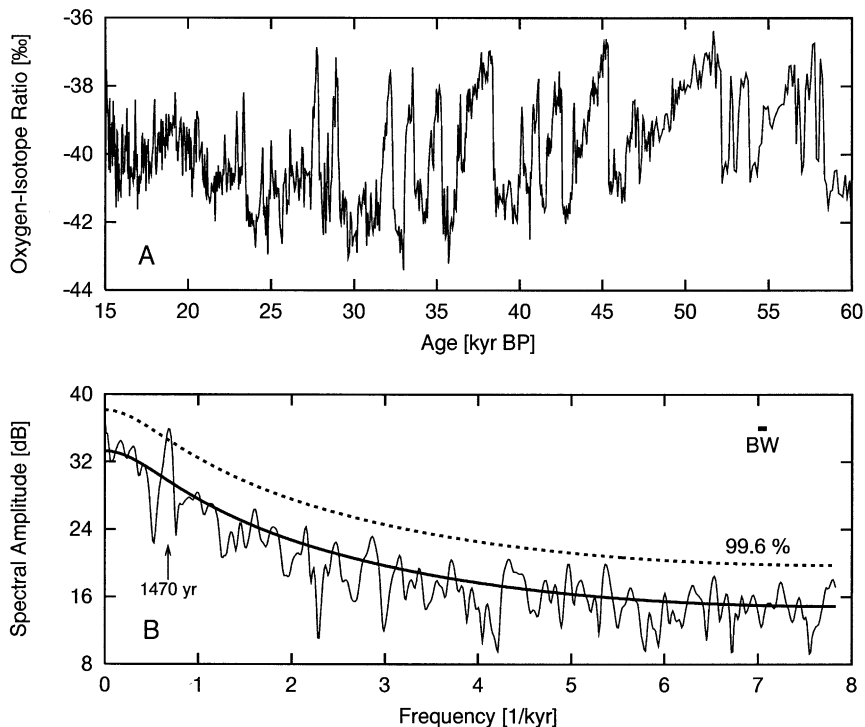


Fig. 2. (A) Oxygen-isotope time series from Greenland GISP2 ice core (Grootes and Stuiver, 1997) between 15–60 thousand years before present (kyr BP). (B) Bias-corrected spectrum of time series in (A) (thin solid line), theoretical red-noise spectrum based on estimated τ (thick solid line) and false-alarm level (99.6%, after Thomson, 1990). Spectral peak at period of 1470 yr (arrow) is inconsistent with AR1 origin. Horizontal bar indicates 6-dB bandwidth (BW).

In the second example, we investigate the glacial part of the oxygen-isotope record from the GISP2 ice core from Greenland (Grootes and Stuiver, 1997; Fig. 2A), which reflects, to a large extent, air temperature above Greenland. In the initial step of the analysis we determine whether or not the spectrum of this time series is consistent with a red-noise model. Based on the periodogram of the time series ($n_{50} = 1$; rectangular window; cf. Schulz and Stattegger, 1997) and $N_{sim} = 1000$ Monte-Carlo simulations, the runs test indicates that the AR1 model is indeed appropriate to characterize this record (5% significance level). The estimated mean value of τ is 310 yr with 90% confidence interval $240 < \tau < 380$ yr. Next we test if any non-AR1 components can be identified in the time series. For this purpose we repeat the analysis, but increase the number of WOSA segments in the spectral analysis in order to obtain a consistent spectral estimate (we refer the reader to Schulz and Stattegger, 1997 for details of the spectral-analysis technique). Setting $n_{50} = 4$ and selecting a Welch spectral window to reduce spectral leakage results in the spectrum depicted in Fig. 2B. We scale the theoretical red-noise spectrum by an appropriate percentile of the χ^2 -probability distribution to obtain a false-alarm level, which marks the maximum spectral amplitude expected if the time series would have been generated by an AR1 process. Accordingly, spectral peaks exceeding the false-alarm level indicate non-AR1 components in a time series, and should be considered significant. We follow Thomson (1990) and select a false-alarm level of $(1 - 1/n) \times 100\%$, where n is the number of data points in each WOSA segment. For the example at hand, a false-alarm level of 99.6% results. At this level the spectrum indicates the presence of a single peak at $f = 1/(1470 \text{ yr})$ which is not consistent with the red-noise model. This spectral peak is associated with the so-called Dansgaard-Oeschger oscillations, the dominant mode of millennial-scale climate fluctuations during the last glacial period (e.g. Grootes and Stuiver, 1997). However, care should be taken when interpreting these results because the assumption of weak stationarity of the time series may be violated.

5. Conclusions

We present a computer program (REDFIT) for testing whether or not the red-noise shape, often observed in paleoclimatic time series, is consistent with the generation by a first-order autoregressive (AR1) process. In contrast to existing approaches, REDFIT allows direct processing of unevenly spaced time series and, hence, the usual prerequisite of data interpolation is not required. Since interpolation of an unevenly spaced time series is equivalent to low-pass filtering, reddening of an estimated spectrum will result and consequently a

biased test result may be the outcome. As an aside, by correcting for the effect of correlation between Lomb-Scargle Fourier components, the program removes the bias of this Fourier transform for unevenly spaced data. A real-world example demonstrates the capability of REDFIT to detect spectral feature not consistent with an AR1 origin. Although REDFIT indicates whether or not the main assumption (i.e., adequacy of the AR1 model) is violated, the program should not be used as black-box tool without checking the structure of a time series prior to its analysis.

Acknowledgements

This work received support from the BMBF 'Klimaforschungsprogramm' (MS) and DFG Research Grant MU 1595/1-1 (MM).

References

- Bendat, J.S., Piersol, A.G., 1986. *Random Data*, 2nd edn. Wiley, New York, 566pp.
- Gilman, D.L., Fuglister, F.J., Mitchel Jr., J.M., 1963. On the power spectrum of red noise. *Journal of the Atmospheric Sciences* 20 (2), 182–184.
- Grootes, P.M., Stuiver, M., 1997. Oxygen 18/16 variability in Greenland snow and ice with 10^{-3} – 10^5 -year time resolution. *Journal of Geophysical Research* 102 (C12), 26455–26470.
- Hasselmann, K., 1976. Stochastic climate models: Part I. Theory. *Tellus* 28 (6), 473–485.
- Horne, J.H., Baliunas, S.L., 1986. A prescription for period analysis of unevenly sampled time series. *The Astrophysical Journal* 302 (2), 757–763.
- Lomb, N.R., 1976. Least-squares frequency analysis of unequally spaced data. *Astrophysics and Space Science* 39, 447–462.
- Mann, M.E., Lees, J.M., 1996. Robust estimation of background noise and signal detection in climatic time series. *Climatic Change* 33 (3), 409–445.
- Mudelsee, M., 2002. TAUEST: a computer program for estimating persistence in unevenly spaced weather/climate time series. *Computers & Geosciences*, 28 (1) 69–72.
- Percival, D.B., Walden, A.T., 1993. *Spectral Analysis for Physical Applications*. Cambridge University Press, Cambridge, 583pp.
- Robinson, P.M., 1977. Estimation of a time series model from unequally spaced data. *Stochastic Processes and their Applications* 6, 9–24.
- Scargle, J.D., 1982. Studies in astronomical time series analysis. II. Statistical aspects of spectral analysis of unevenly spaced data. *The Astrophysical Journal* 263 (2), 835–853.
- Scargle, J.D., 1989. Studies in astronomical time series analysis. III. Fourier transforms, autocorrelation functions, and cross-correlation functions of unevenly spaced data. *The Astrophysical Journal* 343 (2), 874–887.

- Schulz, M., Stattegger, K., 1997. SPECTRUM: Spectral analysis of unevenly spaced paleoclimatic time series. *Computers & Geosciences* 23 (9), 929–945.
- Thomson, D.J., 1990. Time series analysis of Holocene climate data. *Philosophical Transactions of the Royal Society of London. Series A* 330, 601–616.
- Welch, P.D., 1967. The use of fast Fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms. *IEEE Transactions on Audio and Electroacoustics* 15 (2), 70–73.