

A propagation-separation approach to estimate the autocorrelation in a time-series

D. V. Divine^{1,*}, J. Polzehl², and F. Godtlielsen³

¹Department of Mathematics and Statistics, Faculty of Science, University of Tromsø, 9037, Norway

²Weierstrass Institute for Applied Analysis and Stochastics, Mohrenstr. 39, 10117 Berlin, Germany

³Department of Mathematics and Statistics, Faculty of Science, University of Tromsø, 9037, Norway

* also at: Norwegian Polar Institute, Polar Environmental Centre, 9296 Tromsø, Norway

Received: 22 October 2007 – Revised: 12 June 2008 – Accepted: 19 June 2008 – Published: 23 July 2008

Abstract. The paper presents an approach to estimate parameters of a local stationary AR(1) time series model by maximization of a local likelihood function. The method is based on a propagation-separation procedure that leads to data dependent weights defining the local model. Using free propagation of weights under homogeneity, the method is capable of separating the time series into intervals of approximate local stationarity. Parameters in different regions will be significantly different. Therefore the method also serves as a test for a stationary AR(1) model. The performance of the method is illustrated by applications to both synthetic data and real time-series of reconstructed NAO and ENSO indices and GRIP stable isotopes.

1 Introduction

A frequent use of the first order autoregressive model, AR(1), in climate applications is conditioned by its simplicity and efficiency in capturing the inertial nature of climatic phenomena (Hasselmann, 1976). Whereas a choice of a global parametric structure for the fitted model is common, it is often worthwhile to know whether such an approximation provides an adequate description of the analyzed series. This can become an issue, for example, in null hypothesis testing, when the fitted model is used to assess whether or not the variability recorded in a time series is consistent with a stochastic origin of this type.

In this paper we propose a method which, using the idea of structural adaptation, is capable of isolating the periods in the analyzed data where the parameters of the fitted AR(1)

model differ significantly. The method is largely based on ideas implemented in the adaptive weight smoothing (AWS) procedure for local constant modeling, introduced earlier in (Polzehl and Spokoiny, 2000) in the context of image denoising. The AWS technique was later successfully generalized to the case of an arbitrary local linear parametric structure (Chen et al., 2008) and extended to a broad class of nonparametric models, including e.g. the regression, density, Poisson and binary response model. The same idea was applied to estimation of the tail index parameter, classification, density and volatility estimation (Polzehl and Spokoiny, 2002). Theoretical results for exponential families were achieved in (Polzehl and Spokoiny, 2006).

The paper is presented as follows. In Sect. 2 we recall a global likelihood approach to the estimation of the sought parameters in the AR(1) model. Section 3 introduces the concept of local likelihood. A definition of weights is given in Sect. 4. Sections 5 and 6 present a numerical implementation of the method and explain the choice of parameters used, respectively. Examples demonstrating the performance of the method are shown in Sect. 7 followed by brief conclusions in Sect. 8.

2 Global likelihood estimation

The model of the discrete AR(1) process is formulated as

$$y_i = \phi y_{i-1} + \epsilon_i, i = 1, \dots, n \quad (1)$$

where ϕ is the autoregressive parameter, y_i is the i th observed realization of the AR(1) process and ϵ_i are independent, Gaussian distributed innovations with zero mean and variance σ^2 . We assume that the process is weakly-stationary, that is, $|\phi| < 1$.



Correspondence to: D. V. Divine
(dmitry.divine@npolar.no)

Given parameters ϕ and σ^2 , the likelihood function of the vector $\mathbf{Y} = (y_1, \dots, y_n)$ has the form (Brockwell and Davis, 1998) (see 8.7.2-4):

$$L(\mathbf{Y}|\phi, \sigma^2) = (2\pi\sigma^2)^{-\frac{n}{2}} \times \exp \left\{ -\frac{1}{2\sigma^2} (a_1\phi^2 + (1 + \phi^2)a_2 + a_3 - 2\phi a_4) \right\} \quad (2)$$

where

$$\begin{aligned} a_1 &= y_1^2 & a_2 &= \sum_{j=2}^{n-1} y_j^2 \\ a_3 &= y_n^2 & a_4 &= \sum_{j=2}^n y_j y_{j-1} \end{aligned} \quad (3)$$

The global maximum likelihood estimates for ϕ and σ^2 are defined by maximization of the log-likelihood function $\log(L(\mathbf{Y}|\phi, \sigma^2))$. Solving

$$\frac{\partial(\log L(\mathbf{Y}))}{\partial\sigma^2} = \frac{\partial(\log L(\mathbf{Y}))}{\partial\phi} = 0 \quad (4)$$

for σ^2 and ϕ yields after some algebra approximative likelihood estimates of ϕ and σ^2 given as

$$\hat{\phi} = \frac{a_4}{a_2 + a_1}; \quad \widehat{\sigma^2} = \frac{a_2 + a_3 - 2\hat{\phi}a_4}{n}. \quad (5)$$

The asymptotic variance estimates of the model parameters can be determined using the Fisher information matrix for the log-likelihood $\log L(\mathbf{Y}|\phi, \sigma^2)$. Using the expectations $E(a_1) = E(a_3) = \frac{\sigma^2}{1-\phi^2}$, $E(a_2) = \frac{(n-2)\sigma^2}{1-\phi^2}$ and $E(a_4) = \frac{(n-1)\phi\sigma^2}{1-\phi^2}$ this leads to

$$I_{\phi, \sigma^2} = -\mathbf{E}_{\phi, \sigma^2} \begin{pmatrix} \frac{\partial^2 \log L(\mathbf{Y}|\phi, \sigma^2)}{(\partial\phi)^2} & \frac{\partial^2 \log L(\mathbf{Y}|\phi, \sigma^2)}{\partial\phi\partial\sigma^2} \\ \frac{\partial^2 \log L(\mathbf{Y}|\phi, \sigma^2)}{\partial\phi\partial\sigma^2} & \frac{\partial^2 \log L(\mathbf{Y}|\phi, \sigma^2)}{(\partial\sigma^2)^2} \end{pmatrix} = \begin{pmatrix} \frac{n-1}{1-\phi^2} & 0 \\ 0 & \frac{n-2}{2\sigma^4} \end{pmatrix} \quad (6)$$

Inverting I_{ϕ, σ^2} yields the asymptotic covariance matrix for $\{\phi, \sigma^2\}$ as

$$I_{\phi, \sigma^2}^{-1} = \begin{pmatrix} \frac{1-\phi^2}{n-1} & 0 \\ 0 & \frac{2\sigma^4}{n-2} \end{pmatrix}. \quad (7)$$

Respective asymptotic confidence intervals (CI) for $\hat{\phi}$ and $\widehat{\sigma^2}$ can now be obtained employing the asymptotic normality of the estimates.

3 Local likelihood and its estimation using structural adaptation

A global parametric structure considered in the previous section is generally too restrictive. To allow for possible variability in the model coefficients we introduce a local parametric structure to the model. The most general way of describing a local model is based on weights assigned to each observation used. Let, for a fixed t , a nonnegative weight $w_i = w_i(t) \leq 1$ be assigned to the observations y_i at $t_i, i=1, \dots, n$. The local, or weighted, maximum log-likelihood estimate is, in this case, introduced as

$$\{\hat{\phi}(t), \widehat{\sigma^2}(t)\} = \arg \sup_{\phi, \sigma^2} \sum_{i=1}^n w_i(t) \log p(y_i|\phi, \sigma^2)$$

where $\log p(y_i|\phi, \sigma^2)$ denotes the contribution to the global log-likelihood function at any given point t_i , given by

$$\begin{aligned} \log p(y_i|\phi, \sigma^2) &= -\frac{1}{2} \log(2\pi\sigma^2) - \\ &\quad -\frac{1}{2\sigma^2} (y_i - \phi y_{i-1})^2; \quad i = 2, \dots, n. \end{aligned} \quad (8)$$

The principal idea of the approach proposed in (Polzehl and Spokoiny, 2000) is to use a structural assumption of local homogeneity to determine data dependent weights $w_i(t_j)$ that define the local likelihood. Weights $w_i(t_j)$ are obtained using a sequence of likelihood ratio tests for the hypothesis that the parameters at times t_j and t_i coincide.

For any local model characterized by a set of weights $W(t) = (w_1(t), \dots, w_n(t))$, the local log-likelihood defined by these weights is

$$\begin{aligned} \mathcal{L}(W(t), \phi, \sigma^2) &= -\frac{N(t)}{2} \log(2\pi\sigma^2) - \\ &\quad -\frac{1}{2\sigma^2} (R_1(t) + \phi^2 R_2(t) - 2\phi S(t)), \end{aligned} \quad (9)$$

were

$$\begin{aligned} N(t) &= \sum_{j=1}^n w_j(t); & S(t) &= \sum_{j=1}^n w_j(t) y_j y_{j-1}; \\ R_1(t) &= \sum_{j=2}^n w_j(t) y_j^2; & R_2(t) &= \sum_{j=2}^n w_j(t) y_{j-1}^2. \end{aligned} \quad (10)$$

Maximization of this weighted log-likelihood with respect to ϕ and σ^2 yields the local Maximum Likelihood estimates

$$\hat{\phi}(t) = \frac{S(t)}{R_2(t)} \quad \text{and} \quad \widehat{\sigma^2}(t) = \frac{R_1(t)R_2(t) - S(t)^2}{N(t)R_2(t)}. \quad (11)$$

Let us assume that the time series is stationary within some local vicinity $U(t_i)$ of time t_i . If we knew the neighborhood $U(t_i)$ in advance, we would define local weights as $w_{ij} = w_j(t_i) = I_{t_j \in U(t_i)}$, where I denotes the indicator function, and use them to estimate the model parameters. Such

knowledge is usually not available, but information on a suitable local model $W(t_i)$ can be inferred from estimates $\{\hat{\phi}(t), \hat{\sigma}^2(t)\}$. Assume that we have estimates in all observed time points t_j . We can use this information to infer on the neighborhood $U(t_i)$ by testing the hypothesis

$$H : \{\phi(t_i), \sigma^2(t_i)\} = \{\phi(t_j), \sigma^2(t_j)\}. \quad (12)$$

A weight w_{ij} can be assigned based on the value of a test statistic T_{ij} , assigning zero weights if $\{\phi(t_i), \sigma^2(t_i)\}$ and $\{\phi(t_j), \sigma^2(t_j)\}$ are significantly different.

This can be embedded into an iterative procedure. At each iteration (k) we restrict positive weights to observations at times t_j with $|t_j - t_i| < h^{(k)}$, starting with a very small initial bandwidth ($h^{(0)}$) and increasing the bandwidth with each iteration. Testing along the entire time-series yields a set of weights $W(t_i)$ that determines a local model in t_i and hence a new estimate of $\{\phi(t), \sigma^2(t)\}$ in t_i .

We define a test statistics T_{ij} by a local likelihood ratio test for the hypothesis of equal parameters at times t_i and t_j . Given estimates $\{\hat{\phi}_i, \hat{\sigma}_i^2\}$ and $\{\hat{\phi}_j, \hat{\sigma}_j^2\}$ at times t_i and t_j , the difference between the local log-likelihoods evaluated for the two estimates employing the weighting scheme W_i is

$$\begin{aligned} T_{ij} &= \mathcal{L}(W, \hat{\phi}_i, \hat{\sigma}_i^2, \hat{\phi}_j, \hat{\sigma}_j^2) \\ &= \frac{N_i}{2} \log\left(\frac{\hat{\sigma}_j^2}{\hat{\sigma}_i^2}\right) - R_{1i}\left(\frac{1}{2\hat{\sigma}_i^2} - \frac{1}{2\hat{\sigma}_j^2}\right) - \\ &\quad - \frac{1}{2\hat{\sigma}_i^2}(\hat{\phi}_i^2 R_{2i} - 2\hat{\phi}_i S_i) + \frac{1}{2\hat{\sigma}_j^2}(\hat{\phi}_j^2 R_{2i} - 2\hat{\phi}_j S_i). \end{aligned} \quad (13)$$

The value of T_{ij} is used to define a statistical penalty. That is, if T_{ij} is sufficiently larger than some prescribed value λ , the parameters at time t_i and t_j are significantly different and therefore the corresponding weight is set to zero. The process of finding the local models and estimating the local likelihoods is implemented in an iterative procedure, formally presented below in Sect. 5.

Approximate confidence limits for the parameters $\{\phi, \sigma^2\}$ at every t_i can be established using an equivalent of Eq. (7) for the local likelihood and assuming asymptotic normality of the estimates. This yields approximate variances of $\hat{\phi}(t)$ and $\hat{\sigma}^2(t)$ as

$$\begin{aligned} V_{\hat{\phi}}(t) &= \frac{\sum_{j=2}^n w_j(t)^2}{N(t)^2} (1 - \hat{\phi}(t)^2) \quad \text{and} \\ V_{\hat{\sigma}^2}(t) &= \frac{\sum_{j=2}^n w_j(t)^2}{N(t)^2} 2\hat{\sigma}^4(t) \end{aligned} \quad (14)$$

The correspondent $100(1 - \alpha)\%$ CIs are given by $\hat{\phi}(t) \pm Z_{\alpha/2} \sqrt{V_{\hat{\phi}}(t)}$ and $\hat{\sigma}^2(t) \pm Z_{\alpha/2} \sqrt{V_{\hat{\sigma}^2}(t)}$ on $\hat{\phi}(t)$ and $\hat{\sigma}^2(t)$, respectively, where $Z_{\alpha/2}$ denotes the $\alpha/2$ quantile of the standard Gaussian distribution. These intervals are approximate in the sense that they are centered at mean parameter values within the set of time points with positive weights $w_i(t_j)$.

They are, as usual for any nonparametric estimates, not adjusted for a potential bias. Additionally the asymptotic for σ^2 is very slow.

Note that in a homogeneous situation with a free extension of every local model during the iterative process, local MLE's $\{\hat{\phi}(t), \hat{\sigma}^2(t)\}$ and their variances converge asymptotically to their global MLE estimates, as given by Eqs. 5 and 7.

4 Definition of weights

For every pair (i, j), the weight w_{ij} is constructed on the base of two independent quantities: a location penalty $l_{ij} = |\rho(t_i, t_j)/h|^2$ and a statistical penalty $s_{ij}^{(k)} = T_{ij}/\lambda$. Here h and $\rho(t_i, t_j) = |t_i - t_j|$ denote, respectively, the bandwidth and the Euclidean distance between the design points t_i and t_j . The parameter λ can be treated as the critical value for the test statistics T_{ij} and defined empirically to satisfy a propagation criteria in the homogeneous case (see below). The weights are thereby defined as

$$w_{ij} = K_{\text{loc}}(l_{ij}) K_{st}(s_{ij}) \quad (15)$$

with K_{loc} and K_{st} being two positive kernel functions. We present the definitions of kernels later in Sect. 6.

5 Numerical implementation

The computational steps to analyze a time-series using the proposed technique are as follows:

1. Initialization

Specify the initial bandwidth ($h^{(0)}$) and compute, for every $i=2, \dots, n$, the weights $w_{ij}^{(1)} = K_{\text{loc}}(l_{ij}^{(1)})$ and statistics $N_i^{(1)}, S_i^{(1)}, R_{1i}^{(1)}, R_{2i}^{(1)}$ as defined in Eq. (10). Obtain initial estimates of $\phi_i^{(1)}$ and $\sigma_i^{2(1)}$ using Eq. (11).

2. Adaptation

At each iteration step k compute for every pair $i, j=2, \dots, n$ the location and statistical penalties

$$\begin{aligned} l_{ij}^{(k)} &= ((t_i - t_j)/h^{(k)})^2 \\ s_{ij}^{(k)} &= T_{ij}^{(k-1)}/\lambda \end{aligned}$$

and obtain a new set of weights as $w_{ij}^{(k)} = K_{\text{loc}}(l_{ij}^{(k)}) K_{st}(s_{ij}^{(k)})$. This specifies the local model at t_i by $W_i^{(k)} = \{w_{i2}^{(k)}, \dots, w_{in}^{(k)}\}$.

3. Local estimation

Compute new estimates of the sought parameters $\phi_i^{(k)}$ and $\sigma_i^{2(k)}$ (Eq. 11) for iteration step k . Increase the actual bandwidth $h^{(k)}$ by some specified factor $a > 1$ to widen the search radius for the next iteration.

4. Stopping

Stop the procedure if the bandwidth reaches the specified limit h_{\max} , otherwise continue with the adaptation step.

6 Choice of parameters

The proposed method involves several parameters and kernel functions that are to be specified. The most important is λ which scales the statistical penalty and defines its contribution to the weighting scheme at each particular point. If λ is too large, this contribution is negligible and leads to a kernel estimate with bandwidth h_{\max} . Too small values of λ , in turn, lead to overpenalization and may result in a random segmentation. We therefore suggest to choose a minimal value of λ satisfying a propagation condition (Polzehl and Spokoiny, 2006). This implies, in a completely homogeneous situation, a free extension of every local model during the iterative process. More specifically, if $\phi(t) \equiv \phi$ and $\sigma^2(t) \equiv \sigma^2$, we request at each iteration step k and for a specified constant $\alpha > 0$ that

$$E|\hat{\phi}^{(k)} - \phi| \leq (1 + \alpha)E|\check{\phi}^{(k)} - \phi| \quad (16)$$

$$E|\hat{\sigma}^{(k)} - \sigma| \leq (1 + \alpha)E|\check{\sigma}^{(k)} - \sigma|. \quad (17)$$

Here $\check{\phi}^{(k)}$ and $\check{\sigma}^{(k)}$ denote the nonadaptive kernel estimate of ϕ employing the bandwidth $h^{(k)}$ from step k . Then, if h_{\max} is sufficiently large, the resulting local estimates of ϕ and σ^2 will employ the same local weighting scheme and will with a high probability coincide in every point with the global estimates. Note that the sought λ does not depend on the unknown parameter ϕ and therefore can be approximately found by simulations. For convenience, the value of λ is introduced in the programming code via the quantile of the χ^2 distribution with two degrees of freedom and probability p_λ . A value of $p_\lambda = 0.7$ was found to obey the propagation condition for a prespecified value of $\alpha = 0.1$ by simulation in a homogeneous situation.

The minimal bandwidth, $h^{(0)}$, which determines the method's resolution, should be reasonably small. The minimal value of $h^{(0)}$ in the model is 2, to ensure identifiability of parameters. Nevertheless, to improve stability we suggest using $h^{(0)} = 3$, which is set as a default in the implementation. The maximal bandwidth h_{\max} controls the numerical complexity of the algorithm. Ideally, a sufficient value of h_{\max} is comparable to the size of the maximal homogeneous region in the analyzed data. However, since this is not known a priori, h_{\max} equal to the length of the time-series can be used instead. It ensures a convergence of the algorithm in terms of convergence of the statistical penalties $K_S(s_{ij})$ to 1, if t_i and t_j are in the same homogeneous region, or 0, if they belong to regions with different parameters, respectively. Factor a controls the growth rate of the local neighborhood in the course of the iterative process for every point x_i . By default we use the value of $a = 1.25$.

The kernel functions used in constructing the local weighting scheme should satisfy two main criteria: be non-negative and non-increasing. For calculating the statistical penalty we use $K_{st}(u) = e^{-u} I_{u \leq 5}$. A good alternative is $K_{st}(u) = I_{u \in [0, 1/4]} + 4/3(1 - u) I_{u \in [1/4, 1]}$ leading to better sensitivity and an earlier stabilization of the estimates. The localization kernel used in calculating the location penalty is $K_{loc}(u) = (1 - u)_+$. Numerical experiments with different kernel designs have demonstrated that the shape of the location kernels has only minor influence on the resulting estimates.

7 Case studies

We will in this section present results obtained by applying the AWS method to both synthetic and real data sets. The goal is to test the behavior of the algorithm and demonstrate the overall performance of the method. We first run a series of simulations using synthetic data assessing how a goodness of fit varies depending on the properties of the test series (Sect. 7.1). Then two simple examples with synthetic time-series where the parameters ϕ and σ^2 are piecewise constant and linearly changing are presented in Sect. 7.2. The other three examples in Sects. 7.3, 7.4 and 7.5 show the application of the proposed method to time-series widely used in climate studies: the NAO and ENSO *Niño* - 3 indices and GRIP oxygen isotope data (see further details below).

Appropriateness of the AR(1) model is often checked by default in geophysical applications. We fitted the global AR(1) model to the sample time-series using the *armcov* MatLab function utilizing a modified covariance method. The cumulative periodogram test for randomness of the residuals (Box et al., 1994) did not reject stationarity of the proposed model for describing the analyzed time-series of NAO and *Niño* - 3, although it was not the case for the last example of oxygen isotope record from Greenland ice core shown in Sect. 7.5. This is in contradiction to our findings indicating that the default test is not sensitive enough.

The results of our analysis are visualized in a four-components visual display, see Figs. 4–8. It comprises the analyzed time-series itself (top left panel), estimated ϕ and σ^2 of a signal (top right and bottom left panels, respectively) and a sum of weights $N(t)$ at the last iteration step (bottom right panel).

7.1 Performance analysis

To evaluate the performance of the algorithm we ran a series of numerical simulations. The test data was a composite signal constructed concatenating two equal length AR(1) sequences, such that $\phi_2 = \phi_1 \pm \Delta\phi$, where $\phi_1 \in U(-1, 1)$ and $\phi_2 \in (-1, 1)$. For each numerical experiment two parameters were assigned: the length of the composite time-series n and the difference $\Delta\phi$ of the autocorrelation coefficients.




Fig. 1. A family curve showing the 95% quantiles of the distribution of 1000 AICC-based goodness of fit values (see text for details) for AWS estimates of the autocorrelation structure in synthetic series. The series of length n are constructed by concatenating equal length segments with $\sigma^2=1$ and AR(1) autocorrelation coefficient differing in $\Delta\phi$ (shown from the right) between the segments.

We ran 80 experiments in total for $\{n=100,200,\dots,1000; \Delta\phi=0.1,0.2,\dots,0.8\}$ using $M=1000$ independent generated series for each pair $\{n, \Delta\phi\}$. For convenience σ^2 was set to 1 and $E(y_0)=0$.

A goodness of fit measure is established on the basis of the corrected Akaike’s information criterion (AICC, Burnham and Anderson, 2002). We ran the program twice for each particular synthetic series with $p_\lambda=0.7$ and 1.0, which in the second case implies a global MLE estimate of $\{\phi(t), \sigma^2(t)\}$. The corresponding $AICC_{AWS}$ and $AICC_g$ measures were calculated using RSS_{y_i} , the residual sum of squares between the series and its fitted values

$$AICC = 2k + n[\log(\frac{2\pi RSS_{y_i}}{n}) + 1] + \frac{2k(k+1)}{n-k-1},$$

where k denotes the effective number of parameters in the fitted model. For the global AR(1) model we have $k=2$ while for the AWS approach this number is random. Instead of this random value we used the effective number of parameters in the true underlying model which is 4. This number is larger than or equal to the expected effective number of parameters, so that our estimate will in most cases be larger or equal to the AICC using the correct value of (k). Each numerical experiment with $\{n, \Delta\phi\}$ yielded in total 1000 values of the test statistics $AICC_{AWS} - AICC_g$ scaled, for convenience, on the length of the series n . We then used the 95% quantile of the empirical distribution of $(AICC_{AWS} - AICC_g)/n$ as a conservative estimate of the relative skill of the method in identifying the autocorrelation structure of the data. Note that Fig. 1 shows that the AICC-derived goodness of fit is preferentially negative, suggesting that the choice of the proposed technique within this particular ensemble of numerical




Fig. 2. Same as in Figure 1 but for the goodness of fit measure based on RMSE of the estimated ϕ (see text for details).

experiments is at least as relevant as the use of the common global MLE method. Usually for small n and small $\Delta\phi$ we do not detect nonstationarity, and the resulting estimate coincides with the global estimate. What we see in these cases in Fig. 1 is the effect of overestimating $AICC_{AWS}$ by using $k=4$ instead of the correct $k=2$ if nonstationarity was not detected.

This approach however does not demonstrate if the method does detect the presence of jumps in the parameters of the modeled series. Using our prior knowledge of the experimental design we define a second goodness of fit measure from the RMS error of the estimated ϕ relative to the true one. For each run of AWS in $m=1, \dots, M$ with $\{n, \Delta\phi\}$, $RMSE_m^{(n, \Delta\phi)}$, the root mean square error in ϕ of the estimated model $\{\hat{\phi}(t), \hat{\sigma}^2(t)\}$ was calculated. A simple approach to quantify the capability of the method to capture the sudden change in ϕ for a given $\{n, \Delta\phi\}$, is to compare the $RMSE^{(n, \Delta\phi)}$ with $\Delta\phi$. The respective goodness of fit statistics is defined as

$$g_{(n, \Delta\phi)} = \frac{\Delta\phi}{2RMSE^{(n, \Delta\phi)}} \tag{18}$$

with $g > 1$ being a critical value, implying that the procedure has successfully separated the two segments with different correlation structures.

Figure 2 displays the family curve with estimated values of $Q_{RMSE^{(n, \Delta\phi)}}^{0.95}$ for different values of $\Delta\phi$. Here $Q_{RMSE^{(n, \Delta\phi)}}^{0.95}$ denotes the 95% quantile of the empirical distribution of $RMSE^{(n, \Delta\phi)}$ for a given $\{n, \Delta\phi\}$. We see that for a relatively short (about 200 points) time-series the method guarantees separation only if the step in the autocorrelation coefficient is sufficiently large (above 0.7). These numerical estimates are in reasonably good agreement with a simple inference about the expected performance of the method. These can be drawn from the asymptotic 95% confidence intervals for ϕ in




Fig. 3. Double width of the 95% confidence interval ($CI_{95\%}$) for the global maximum likelihood estimate of ϕ for different values of a time-series length n and with autocorrelation coefficient ϕ .

case of the global MLE (see Eq. 7 and Fig. 3). One should note, however, that this numerical experiment disregards the actual values of ϕ_1 and ϕ_2 , whereas the width of the confidence interval for $\hat{\phi}$ varies substantially with ϕ and n . Figure 2 therefore provides the mean estimates of the goodness of fit as the whole range of possible (ϕ_1, ϕ_2) is considered.

7.2 Synthetic time-series

A simple synthetic series containing $n=1500$ data points is constructed by concatenation of segments of different lengths with locally constant ϕ and σ^2 . The AWS results are shown in Fig. 4. The figure demonstrates the ability of the method to successfully localize discontinuities in the model parameters by a propagation of weights within homogeneous regions. Note that despite some discrepancy between the true (dotted lines) and estimated (solid lines) values of ϕ and σ^2 , a good agreement is generally reached.

Figure 5 shows the second example where the method is applied to a series of length $n=1000$ with $\sigma^2=1$ and ϕ changing linearly from -0.99 to 0.99 . For $h_{\max}=n$ the method yields piecewise constant estimates of ϕ and σ^2 , in accordance with the structural assumption of local stationarity. Note that the magnitudes of jumps in ϕ and the correspondent segment lengths are close to the minimal detectable ones for the considered $\frac{\Delta\phi}{n}$, as Fig. 2 shows.

7.3 Reconstructed winter NAO index

Figure 6 shows the AWS analysis of the winter (December through March) North Atlantic Oscillation (NAO) index reconstruction for the period 1500–1997 published in (Luterbacher et al., 2002) (the series is available in the World Data Center for Paleoclimatology at <http://www.ncdc.noaa.gov/paleo/wdc-paleo.html>).

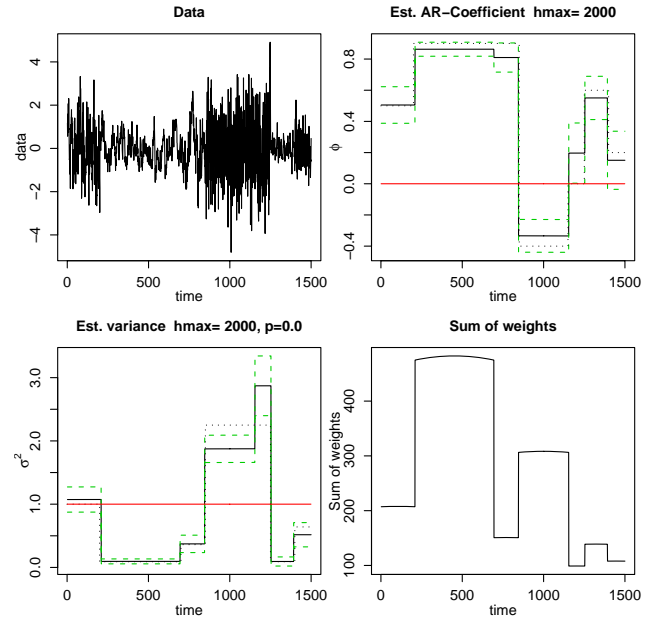


Fig. 4. Synthetic time-series with piecewise constant ϕ and σ^2 (top left); AWS estimates of the autocorrelation coefficient (solid lines, top right) and variance (solid lines, bottom left); sum of weights, $N(t)$ (bottom right). Dotted lines show the true values of ϕ and σ^2 used when modeling the time-series. Dashed lines outline the confidence limits for the estimates of ϕ and σ^2 . A p-value of the Shapiro-Wilk test for normality of the residuals is shown above the left bottom panel.

The NAO is the dominant pattern of atmospheric circulation variability over the North Atlantic basin, having large impacts on weather and climate in the North Atlantic region and surrounding continents. The NAO index, which is defined as the time-averaged difference of sea level pressure (SLP) between Iceland and the Azores, reflects the strength of the westerly across the Atlantic basin into Europe (note that there also exist alternative definitions of the NAO index, see for example van Loon and Rogers, 1978).

Luterbacher et al. (2002) provide a proxy- and early instrumental data based reconstruction of the seasonal and monthly NAO indices. We used winter means for the earlier part of the record for 1500–1668 and derived mean December through March indices from the respective monthly means for the rest of the period.

The estimated ϕ is close to zero for the first four hundred years of the reconstruction period. Around the turn of the 20th century, however, the lag-1 autocorrelation coefficient undergoes an excursion from zero towards more positive values. A tendency for the winter NAO to show a behavior close to a Gaussian noise during the 19th century have already been reported earlier. Hurrell and van Loon (1997) came to this conclusion when analyzing a spectrum of the




Fig. 5. Same as in Fig. 4 but for synthetic time-series with $\sigma^2=1$ and ϕ changing linearly from -0.99 to 0.99 .

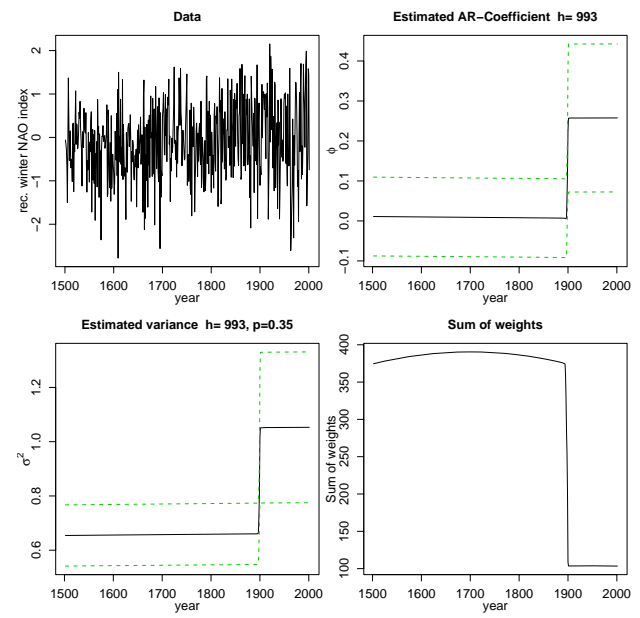


Fig. 6. Reconstructed winter NAO index (top left); AWS estimates of the autocorrelation coefficient (top right) and variance (bottom left); sum of weights, $N(t)$ (bottom right). Dashed lines outline the confidence limits for the estimates of ϕ and σ^2 .

instrumental station-based winter-mean NAO SLP index for the shorter period of 1864–1996. Weak positive ϕ in the 20th century reflects a tendency for the NAO to have a slightly “red” spectrum, what can be attributed to the enhanced variability at decadal scale in the 20th century (Hurrell and van Loon, 1997). Note that the change in $\hat{\phi}$ is accompanied by the increase in the estimated variance, which can also be ascribed to strengthening of the decadal variability in the NAO. A simple AR(1) model fitted to the data, despite being adequate for the analyzed time-series on average, is not capable of capturing the quasi-oscillatory behavior in the data. Yet it may provide an indication that the character of the dependence in the analyzed series has changed.

7.4 Reconstructed boreal winter Niño – 3 index

The next example shows the application of AWS to the reconstructed boreal winter Niño – 3 index (Mann et al., 2000) covering the period 1650–1980 (data from <http://www.ncdc.noaa.gov/paleo/wdc-paleo.html>). The Niño – 3 index is based on the eastern tropical Pacific sea surface temperatures and serves as one of the indicators of the global El Niño/Southern Oscillation (ENSO) variability. Positive and negative values of the Niño – 3 index indicate El Niño (warm) and La Niña (cold) episodes, respectively.

Figure 7 reveals a substantial decrease in $\hat{\phi}$ between the two intervals of not rejected local stationarity with respect to the AR(1) model fitted to the reconstructed Niño – 3 index. The change most likely occurred around 1800, and is accompanied by a rise in the estimated value of σ^2 . We suggest

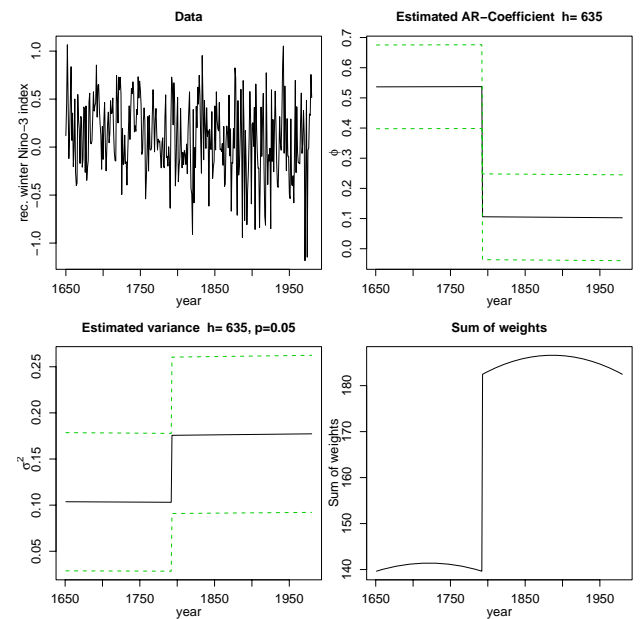


Fig. 7. Same as in Fig. 6 but for the reconstructed boreal winter Niño – 3 index.

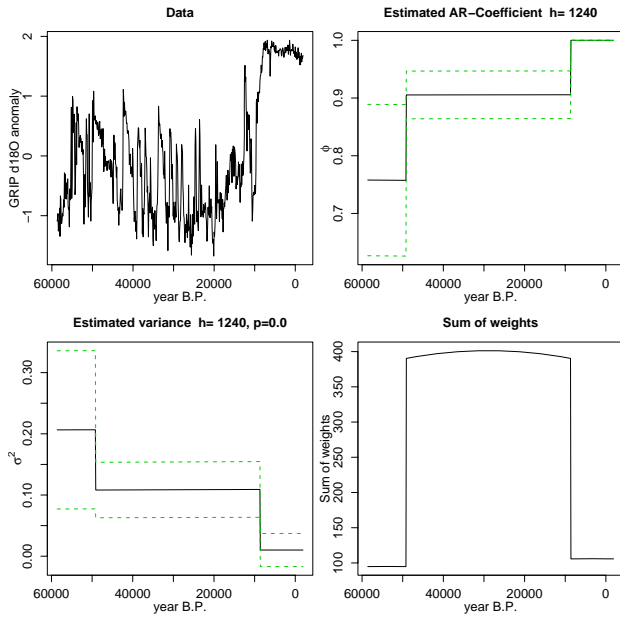


Fig. 8. Same as in Fig. 6 but for the resampled GRIP $\delta^{18}O$ series.

that this provides an indication of a change in the character of the dependence in the correlation structure of the analyzed series. During the first half of the considered period the reconstructed *Niño* – 3 index exhibits a negative trend which levels out in the early 19th century. We note that the presence of this long-term trend component is not consistent with the simple model fitted to the data and partly accounts for the elevated values of $\hat{\phi}$ before 1800.

The tendency towards decreased ENSO variability before 1850 has been argued for in previous proxy-based ENSO studies (see for example Stahle and Cleaveland, 1993). Mann et al. (2000), in turn, came to a similar conclusion having applied the evolutive spectral analysis to the same *Niño* – 3 index reconstruction. The analysis revealed an enhanced interannual variability after the mid 19th century, which also agrees well with our estimates.

7.5 GRIP oxygen isotope series

In the last example we apply the AWS technique to GRIP (Greenland Ice Core Project) ice core $\delta^{18}O$ series (Johnsen, 1999) which serves, to a large extent, as an indicator of past temperature changes at the core site. The analyzed part of the record covers some 60 000 years and is originally unevenly sampled in time. We resampled the time-series using 100 year bins and calculated the normalized anomalies. The time increment of 100 years for the procedure is chosen so that at least one data point falls within the resampling interval in the oldest part of the series under consideration. In the upper part of the core the data density is higher, varying be-

tween 25 data points per bin in the 20th century to 5 at the termination of the last glacial period around 12 000 years BP.

Figure 8 suggest an increase in the serial autocorrelation coefficient GRIP $\delta^{18}O$ with the onset of the Holocene accompanied by a decrease in the estimated variance. This result represents a combined effect of actual climate variability and a bias introduced in the series by data interpolation. On one hand, the Holocene is characterized by a more stable climate compared with the last glacial period punctuated by a series of abrupt warmings – the so-called Dansgaard-Oeschger oscillations (Dansgaard et al., 1993). On the other hand the resampling is known to alter the series autocovariance structure by bringing additional dependence to the data (Schulz and Stategger, 1997). The latter effect will naturally be more pronounced in the uppermost part of the ice core series where the sampling density in the time domain is higher and each point in the resampled series is an average over a number of original observations. As a reference we estimated the autocorrelation using the RedFit package (Schulz and Mudelsee, 2002) which fits the AR(1) model directly to an unevenly spaced time series. The separate analysis of the Holocene and glacial parts of the original GRIP $\delta^{18}O$ series yields the values of ϕ of 0.01 and 0.87, respectively, which are below the corresponding AWS estimates.

8 Conclusions

We presented a method which employs the idea of structural adaptation for fitting an AR(1) model to a time-series. The approach utilizes an assumption of a local stationary AR process. This is used to simultaneously generate local weighting schemes, i.e. a local model, and to estimate the parameters of the AR(1) model as functions of time by weighted maximum likelihood. The proposed procedure leads, for a large maximum bandwidth (h_{\max}) to a local constant approximation of these parameter functions. This seems appropriate if periods of local stationarity are separated by shorter periods of rapid change. The approach implicitly provides a test for global stationarity, i.e. leads to the global AR(1) model if stationarity cannot be rejected.

An implementation of the AWS method is available from the authors as a package (acoraws) for the R-Project for Statistical Computing (R Development Core Team, 2005).

Acknowledgements. The authors thank WIAS for hosting D.Divine. This study was financially supported by the Norwegian Research Council, projects 160008/v30 and 176872/v30.

Edited by: J. Kurths

Reviewed by: two anonymous referees

References

- Box, E., Jenkins, G., and Reinsel, G.: Time series analysis: Forecasting and control, Prentice-Hall, Englewood Cliffs, 1994.
- Brockwell, P. and Davis, R.: Time series: theory and methods, Springer, New York, 1998.
- Burnham, K. and Anderson, D.: Model selection and multimodel inference: a practical-theoretic approach, Springer, 2002.
- Chen, C., Harde, W., and Unwin, A.: Handbook of Data Visualization, Springer, 2008.
- Dansgaard, W., Johnsen, S., Clausen, H., Dahl-Jensen, D., Gundestrup, N., Hammer, C.U. and Hvidberg, C., Steffensen, J., Sveinbjornsdottir, A., Jouzel, J., and Bond, G.: Evidence for general instability of past climate from a 250-kyr ice-core record, *Nature*, 364, 218–220, doi:10.1038/364218a0, 1993.
- Hasselmann, K.: Stochastic climate models: Part I. Theory, *Tellus*, 28(6), 473–485, 1976.
- Hurrell, J. and van Loon, H.: Decadal variations in climate associated with the North Atlantic Oscillation, *Climate Change*, 36, 301–326, 1997.
- Johnsen, S.: GRIP Oxygen Isotopes, PANGAEA, doi:10.1594/PANGAEA.55091, 1999.
- Luterbacher, J., Xoplaki, E., Dietrich, D., Jones, P., Davies, T., Portis, D., Gonzalez-Rouco, J., von Storch, H., Gyalistras, D., Casty, C., and Wanner, H.: Extending North Atlantic Oscillation Reconstructions Back to 1500, *Atmos. Sci. Lett.*, 2, 114–124, 2002.
- Mann, M., Bradley, R., and Hughes, M.: Long-term variability in the El Nino Southern Oscillation and associated teleconnections, in: *El Nino and the Southern Oscillation: Multiscale Variability and its Impacts on Natural Ecosystems and Society*, edited by Diaz, H. and Markgraf, V., Cambridge University Press, Cambridge, UK, 2000.
- Polzehl, J. and Spokoiny, V.: Adaptive weights smoothing with application to image restoration, *Journal of the Royal Statistical Society*, 62(B), 335–354, 2000.
- Polzehl, J. and Spokoiny, V.: Local likelihood modeling by adaptive weights smoothing, Preprint 787, WIAS, 2002.
- Polzehl, J. and Spokoiny, V.: Propagation-separation approach for local likelihood estimation, *Probab. Theory Related Fields*, 135, 335–362, 2006.
- R Development Core Team: R: A Language and Environment for Statistical Computing, Tech. Rep. ISBN 3-900051-07-0, R Foundation for Statistical Computing, Vienna, Austria, 2005.
- Schulz, M. and Mudelsee, M.: REDFIT: estimating red-noise spectra directly from unevenly spaced paleoclimatic time series, *Computers&Geosciences*, 28, 421–426, 2002.
- Schulz, M. and Stettgen, K.: SPECTRUM: spectral analysis of unevenly spaced paleoclimatic time series, *Computers&Geosciences*, 23, 929–945, 1997.
- Stahle, D. W. and Cleaveland, M. K.: Southern Oscillation Extremes Reconstructed from Tree Rings of the Sierra Madre Occidental and Southern Great Plains., *J. Climate*, 6, 129–140, 1993.
- van Loon, H. and Rogers, J.: The seesaw in winter temperatures between Greenland and Northern Europe. Part I: general description, *Monthly Weather Review*, 106, 296–310, 1978.