11-15-1999

Evaluating Maximum Likelihood Estimation Methods to Determine the Hurst Coefficient

Christina Marie Kendziorski
Marquette University

J. B. Basingthwaighte
University of Washington - Seattle Campus

Peter J. Tonellato
Marquette University, peter.tonellato@marquette.edu

Published in final edited form as:

Physica A. 1999 November 15; 273(3-4): 439–451. doi:10.1016/S0378-4371(99)00268-X.

Evaluating maximum likelihood estimation methods to determine the Hurst coefficient

C.M. Kendzior^{a,*}, J.B. Bassingthwaite^b, and P.J. Tonellato^a

^aDepartment of Mathematics, Statistics, and Computer Science, Marquette University, Milwaukee, WI 53233, USA

^bDepartment of Bioengineering, University of Washington, Seattle, WA 98195, USA

Abstract

A maximum likelihood estimation method implemented in S-PLUS (*S-MLE*) to estimate the Hurst coefficient (H) is evaluated. The Hurst coefficient, with $0.5 < H < 1$, characterizes long memory time series by quantifying the rate of decay of the autocorrelation function. *S-MLE* was developed to estimate H for fractionally differenced (fd) processes. However, in practice it is difficult to distinguish between fd processes and fractional Gaussian noise (fGn) processes. Thus, the method is evaluated for estimating H for both fd and fGn processes. *S-MLE* gave biased results of H for fGn processes of any length and for fd processes of lengths less than 2^{10} . A modified method is proposed to correct for this bias. It gives reliable estimates of H for both fd and fGn processes of length greater than or equal to 2^{11} .

Keywords

Time series analysis; Autocorrelation function; Fractional ARIMA process; Fractional Gaussian noise; Hurst coefficient; Long memory

1. Introduction

The analysis of time series has proven useful in providing insight into and information about the dynamics of many complex systems. We define a discrete time series of length N by $\{x_1, x_2, \dots, x_N\}$ where each x_i is a realization of the random variable X_i for $i = 1, 2, \dots, N$. When referring to a general time series, we will use $\{x_t\}$ to denote the series and $\{X_t\}$ to denote the underlying stochastic process from which the series was obtained. Although the indexing variable t can refer to dimensions such as space, we will assume that the observations are ordered in time, and that the intervals are uniform.

Of particular interest and importance in many scientific disciplines are stochastic processes which exhibit long memory characteristics [1–3]. A process is long memory if its correlations decay to zero so that the sum of the autocorrelations is infinite.

Specifically, a stationary stochastic process X_t ($E[X_t] = \mu$ and $\text{var}[X_t] = \sigma^2$) is a long memory process if there exists a real number $\alpha \in (0, 1)$ and a positive constant c_ρ such that

$$\lim_{k \rightarrow \infty} \frac{\rho(k)}{c_\rho k^{-\alpha}} = 1, \quad (1)$$

where $\rho(k) = (E[X_t - \mu, X_{t-k} - \mu]) = \sigma^2$, the autocorrelation function. Since $0 < \alpha < 1$, $\sum_{k=-\infty}^{\infty} \rho(k) = \infty$ for a long-memory process. A long-memory process can also be defined in terms of power-law scaling in the frequency domain (as opposed to the correlation domain given in Eq. (1)). In particular, the stationary process described above is long memory if there exists a real number $\beta \in (0, 1)$ and a positive constant c_h such that

$$\lim_{f \rightarrow 0} \frac{h(f)}{c_h |f|^{-\beta}} = 1, \quad (2)$$

where $h(\cdot)$ represents the spectral density of X_t . Since $h(f) = \sigma^2 (2\pi)^{-1} \sum_{k=-\infty}^{\infty} \rho(k) e^{ikf}$, Eqs. (1) and (2) are theoretically equivalent (see Theorem 2.1 of [1]). However, this theoretical equivalence does not necessarily hold given a particular finite data set which may exhibit clear power-law scaling in one domain (correlation or frequency), but not the other. As the statistical procedures based on analysis in the correlation and frequency domains typically use the underlying data in quite different ways, analysis in each domain will highlight different aspects thereof [4]. For example, there is a long memory process with power-law scaling in its spectral density function, but the power-law scaling is limited to the low range of lags k in the autocorrelation function $\rho(k)$ [5].

Because much of classical time series-related statistical theory applies only to processes with summable correlations, the application of such theory to long memory processes can lead to erroneous statistical results. Thus, techniques to generate, identify, characterize, and analyze long-memory time series are rapidly being developed and applied. The particular types of long memory fractal processes discussed above which demonstrate slightly different self-similarity features were introduced by Hurst [6] and Mandelbrot and van Ness [7] and are given broad context in Beran [1]. The Hurst coefficient, H , is one of the most commonly used parameters to characterize the rate of decay of the spectrum or autocorrelation function of a long-memory process. In 1990, Marsh, Osborn and Cowley reported that the logarithm of the arterial pressure spectrum obtained from dog data decays as $f^{-\beta}$ (i.e., Eq. (2) holds) [8]. Note that β is related to the Hurst coefficient by $\beta = 2H - 1$, where β is taken from the slope of $h(f)$ in the lower-frequency range. The authors noted that the quantification provided by estimation of β , or equivalently H , may be useful in identifying differences between states of the blood pressure regulation system. Similarly, others have noted that estimation of certain fractal dimensions, which are also directly related to H , from physiological time series may prove useful in the diagnosis and treatment of disease [9]. Thus, considerable effort has been expended to develop techniques to estimate the Hurst coefficient and evaluation of such techniques is essential.

These techniques include Hurst's rescaled range (R/S) analysis [6], dispersional analysis [10], scaled windowed variance (SWV) methods (reviewed in [12]), autocorrelation analysis [11], Fourier spectral analysis [13,14], and maximum likelihood estimation (MLE) methods [1]. The evaluation of these as well as other techniques has been and continues to be addressed. It has been determined that R/S analysis, autocorrelation analysis, and Fourier spectral analysis based on the periodogram should all be avoided because they are poor H -estimators. Specifically, Bassingthwaite and Raymond have shown that Hurst's rescaled range estimates of H are highly biased and variable [15]. Schepers, van Beek and Bassingthwaite have shown that estimates based on autocorrelation analysis of series with

power law spectra are also highly biased toward $H = 0.5$ [16]. Fourier spectral analysis based on the periodogram has high variance in its estimates from fractional Gaussian noise (fGn) processes compared to other more reliable estimators such as the dispersional analysis technique, Disp [17], and is highly biased for fractional Brownian motion (fBm) processes [13].

It is important to note that each of these assessments was based on artificial known fractal series generated by the spectral synthesis method (SSM) described by Peitgen and Saupe [18]. SSM generates fractional Brownian motion (fBm) processes; and fractional Gaussian noise (fGn) signals are derived directly from fBm by taking the difference between adjacent points. SSM produces a series where the first and last points are considered neighboring points, so that the beginning and the end segments are correlated. Generating time series several times longer than the desired length and taking a randomly chosen segment of the desired length minimizes this problem [17]. Other generating methods such as the successive random addition (SRA) method can also be used [18]. However, neither method generates processes with the correct power spectral density and autocorrelation [19] and it has been recommended that both techniques be avoided [20].

A method developed by Davies and Harte [21] and modified by Percival [22] generates fGn signals whose population statistics of the autocorrelation function agree exactly with those of Eq. (1), and finite signals generated by this algorithm will have sample statistics that agree with those of fGn within limits imposed by sampling variability. In this case fBm can be obtained from the cumulative sums of fGn. In their review of three variants of SWV methods, Cannon et al. [12] generated fGn signals using the Davies–Harte method. It was shown that estimates obtained using SWV methods have a lower bias and variance than those obtained from R/S, periodogram, or autocorrelation analysis [12]. The SWV methods were as reliable as dispersional analysis; however, it should be noted that dispersional analysis can only be applied to fGn, while SWV must be applied to fBm. This is because the respective methods of analysis are based on fGn and fBm. This is stressed in [5] where Raymond and Basingthwaight derive the dispersional analysis and scaled windowed variance methods from the correlation function of fGn and demonstrate mathematically the errors that arise by misapplying each of these methods.

A general method of evaluation that has received relatively little attention is maximum likelihood estimation (MLE). The main objective of this work is to evaluate the approximate MLE algorithm implemented in the statistical analysis package S-PLUS (*S-MLE*). *S-MLE* is intended to estimate H for a fARIMA(0, d , 0) process, also called a fractionally differenced (fd) process. Similar to fGn, the fd process also demonstrates asymptotic self-similarity in both its power spectral density and autocorrelation function, which is given by

$$\rho_{fd}(k) = \frac{\Gamma(1-d)\Gamma(k+d)}{\Gamma(d)\Gamma(k+1-d)} \quad (3)$$

where $d = H - 0.5$, k is the lag, and $\Gamma(x)$ denotes the gamma function [1]. However, given a finite data set, it is difficult, if not impossible, to distinguish between a fd process and other long-memory processes such as fGn. Thus, it is important to evaluate the behavior of the MLE algorithm applied to both fd and fGn processes. The paper is organized as follows. First, fd processes, fGn processes, and the Hurst coefficient are defined and discussed. Second, *S-MLE* is reviewed. Third, estimates of H for varying lengths of the processes are evaluated by measuring the bias and variance of the estimates using *S-MLE*. The results of this analysis suggest that a method be devised to better estimate H for fGn. A modified *S-MLE* method is then proposed to obtain more accurate estimates of H for fGn signals while

maintaining reliable estimation of H for fd processes. Finally, the proposed method is evaluated.

2. Long-memory processes and the Hurst coefficient

2.1. Fractional Gaussian noise and fractionally differenced processes

A fd process is a natural extension of the autoregressive integrated moving average (ARIMA) models introduced by Box and Jenkins [23]. An ARIMA model of order (p, d, q) can be written as

$$\phi(B)(1 - B)^d X_t = \psi(B)\varepsilon_t, \quad (4)$$

where $\phi(z) = 1 - \sum_{i=1}^p \alpha_i z^i$, $\psi(z) = 1 + \sum_{i=1}^q \beta_i z^i$, $B^k(X_t) = X_{t-k}$, $d \in \mathbb{Z}^+$, and $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$. In 1980, a fractional value of d was considered by Granger and Joyeux, giving rise to the fractional autoregressive integrated moving average (fARIMA) models [24]. A fARIMA process X_t of order (p, d, q) is defined to be the stationary solution of Eq. (4) where $-0.5 < d < 0.5$ ($d = H - 0.5$). We refer to a fARIMA process of order $(0, d, 0)$ as a fd process. A fd process is long memory for $0 < d < 0.5$ ($0.5 < H < 1$).

A detailed discussion of fGn is given in Beran [1]. In short, if Y_t is a self-similar process with stationary, Gaussian increments, then the distribution of $X_t = Y_t - Y_{t-1}$ is fully specified by the mean and autocorrelation function which is given by

$$\rho_{fGn}(k) = \frac{1}{2}[(k+1)^{2H} - 2k^{2H} + (k-1)^{2H}] \quad (5)$$

for $k \geq 0$ and $\rho_{fGn}(k) = \rho_{fGn}(-k)$ for $k < 0$. Thus, for each $H \in (0, 1)$, there is exactly one Gaussian process X_t that is the stationary increment of a self-similar process Y_t . We refer to this process as fGn. In practice, one obtains one of many finite realizations of a fGn process; and sample statistics will agree with those of the sampled fGn process within limits imposed by sampling variability. For $H \in (0.5, 1)$, the correlations of fGn, given by Eq. (5), satisfy Eq. (1); thus, a fGn process with $H \in (0.5, 1)$ is long memory [1]. An intuitive motivation for fGn is discussed by Bassingthwaite et al. [25].

2.2. Similarities between fGn and fd processes

As stated above, for $0.5 < H < 1$, both fd and fGn processes are long memory. In fact, both autocorrelation functions (ACFs), $\rho_{fd}(k)$ and $\rho_{fGn}(k)$, are asymptotically proportional to k^{2H-2} [1]. This similarity in the asymptotic decay rates of the ACF is equivalent to a similarity between the spectra near the origin. In particular, the spectrum, $h_{fGn}(\omega)$, of a long-memory fGn process is very well approximated near the origin by $c_h |\omega|^{1-2H}$ for some $c_h > 0$ [1]. The spectrum of a fd process is also approximately proportional to $|\omega|^{1-2H}$ near the origin since

$$h_{fd}(\omega) = \frac{\sigma_\varepsilon^2}{2\pi} |1 - e^{i\omega}|^{1-2H} \text{ and } |1 - e^{i\omega}| = 2 \sin\left(\frac{1}{2}\omega\right)$$

and

$$\lim_{\omega \rightarrow 0} \frac{2 \sin(\frac{1}{2}\omega)}{\omega} = 1 [1].$$

Because the fd and fGn processes are similar theoretically, it is not surprising that it is difficult to distinguish between the two given a finite data set. Standard methods to identify an unknown stochastic process include a visual analysis of the data in the time, correlation, and frequency domains. This often involves generation and interpretation of time plots, ACFs, and spectra [4,26,27]. Fig. 1 demonstrates the difficulty in using these standard visual methods to identify differences between a fd and fGn process. The upper panels of Fig. 1 give the time plots of a fd and fGn process, which are visually indistinguishable. The decay in the ACFs shown in the middle panel are equally similar; both show a slow decay in the ACF with deviation from the theoretical ACF at larger lags. Finally, both spectra attenuate near the origin and both have points scattered around a line with negative slope at higher frequencies. Similar structure is also observed in the partial ACFs (not shown). Given any two finite realizations of fd and fGn, one may be able to identify subtle differences between the processes; however, refined quantitative tests which are capable of consistently identifying such differences in all sets of realizations have yet to be developed. Since it is difficult to detect fd from fGn a priori, it is important that any algorithm used to characterize properties of processes that are similar to fd or fGn be applicable to both.

2.3. The Hurst coefficient

The above definitions indicate the way in which H is used to define fd and fGn processes. H quantifies the rate of decay of the autocorrelations of both processes, and may be used to compare different time series or to suggest how a stochastic process might be modeled. H is related to the lag k autocorrelation coefficient of fGn, $\rho_{fGn}(k)$, by Eq. (5), and for a fd process $\rho_{fd}(k)$, by Eq. (3). $S\text{-MLE}$ will be evaluated in the range $0 < d < 0.5$ ($0.5 < H < 1$), where both fd and fGn processes are long memory.

3. S-PLUS maximum likelihood estimation (S-MLE) method

The $S\text{-MLE}$ algorithm was developed to analyze fARIMA(p, d, q) processes (recall fd is fARIMA($0, d, 0$)). $S\text{-MLE}$ is developed based upon the relationship between the log-likelihood for the fARIMA(p, d, q) and for the ARMA models. This relationship can be described by considering the ARMA(p, q) process given by Eq. (4) with $d = 0$. Let

$\hat{X}_t^{t-1} = E[X_t | X_{t-1}, \dots, X_1, \alpha_1, \dots, \alpha_p, \beta_1, \dots, \beta_q]$ denote the conditional mean one-step ahead prediction of X_t based on the data $X_1, X_2, \dots, X_{t-1}$, and let

$\sigma_{\varepsilon}^2 g_t = \text{var}[X_t | X_{t-1}, \dots, X_1, \alpha_1, \dots, \alpha_p, \beta_1, \dots, \beta_q]$ denote the conditional variance of $\hat{X}_t^{t-1}$. The prediction errors are defined by $e_t = X_t - \hat{X}_t^{t-1}$. Let $L = L(X_1, \dots, X_n; \bar{\theta})$ denote the likelihood where $\bar{\theta} = [\alpha_1, \dots, \alpha_p, \beta_1, \dots, \beta_q]$. Then,

$$-2 \log L(X_1, \dots, X_n; \bar{\theta}) = n \log(2\pi\sigma_{\varepsilon}^2) + \sum_{t=1}^n \log g_t + \frac{1}{\sigma_{\varepsilon}^2} \sum_{t=1}^n \frac{e_t^2}{g_t}. \quad (6)$$

The log-likelihood for the fARIMA(p, d, q) model can be computed exactly using Eq. (6) where $\hat{X}_t^{t-1}$ and g_t are given by Haslett and Raftery [28]. Specifically,

$\hat{X}_t^{t-1} = \phi(B)\psi(B)^{-1} \sum_{i=1}^{t-1} \gamma_{ti} X_{t-i}$ where γ_{ti} are the partial linear regression coefficients for the fARIMA ($0, d, 0$) process and $g_t = \prod_{i=1}^{t-1} (1 - \gamma_{ii}^2)$. The explicit form of the partial linear regression coefficients is given in Hosking as

$$\gamma_{ti} = -\frac{t!}{i!(t-i)!} \frac{(i-d-1)!(t-d-i)!}{(-d-1)!(t-d)!} \text{ and } \gamma_{ii} = \frac{d}{t-d} \quad [29].$$

The prediction errors e_t are then computed using an algorithm by Ansley based on the Choleski decomposition of the covariance of the process X_t [30]. A major practical problem with maximum likelihood estimation based on this likelihood is that the required CPU time is $O(n^2)$. Thus, an approximation described in Haslett and Raftery (Section 4.3) is used [28]. In short, Hosking has shown that for $1 \ll i \ll t$, as $i, t \rightarrow \infty$ with $it^{-1} \rightarrow 0$; $\gamma_{it} \sim -i^{-(d+1)/2} \Gamma(-d)$. The approximation consists of taking this asymptotic relationship to hold exactly for $j > M$ where M is some integer and defining the γ_{ij} to be constant for $M < j \leq t-1$. This approximation is implemented in *S-MLE* with $M = 100$.

4. Evaluation of *S-MLE*

4.1. Generating fractionally differenced processes and fractional Gaussian noise

The fd processes were generated by the fractional ARIMA algorithm (“arima.fracdiff.sim”) in *S-PLUS*. This algorithm uses the prediction error decomposition to generate X_t from its conditional distribution given the previous values [31]. The fGn signals analyzed were generated by an algorithm designed by Davies and Harte [21] and modified by Percival [22]. The code is implemented in *fgp*, available via ftp from the National Simulation Resource website (www.physio.me.org/software/fractal). The Davies–Harte algorithm generates signals whose population statistics agree exactly with those of fGn as defined by Eq. (5) and will have sample statistics that agree with those of fGn within limits imposed by sampling variability which depend only on series length.

4.2. Experimental design

Numerical experiments were carried out to evaluate the MLE method implemented in *S-PLUS*. First, the bias that occurs when the MLE method is applied to processes other than fd processes was measured by testing the method on both fd and fGn signals. In particular, for a given Hurst coefficient H , where $H \in \{0.5, 0.6, \dots, 0.9\}$, 100 fd and 100 fGn signals were generated for each length N , where $N \in \{2^6, 2^7, \dots, 2^{14}\}$. The mean-squared error (MSE) was chosen as the criterion of method reliability as it takes into account both estimate bias and variance. For each series length, the average MSE of the estimates refers to the MSE averaged over all H . Similarly, average standard deviation and average bias refer to averages taken over all H . The same methodology was applied to evaluate the “modified” MLE method. Results are reported in terms of H . Recall that $d = H - 0.5$.

4.3. Numerical evaluation of *S-MLE*: mean and bias of estimates of H

S-MLE gave reliable estimates (average MSE < 0.0006) of H for fd signals of length greater than or equal to $N = 2^{10}$ but were biased (average MSE > 0.002) for fGn signals of every length considered, as Fig. 2 shows. This is expected as *S-MLE* was developed to estimate H for a fd, not fGn, process. The average standard deviations of the estimates from both fd and fGn processes decay at identical rates from approximately 0.08 for $N = 2^6$ to approximately 0.005 for $N = 2^{14}$. Thus, the differences in the average MSE of the estimates from fd and fGn processes are due to differences in the bias of the estimates. Specifically, the average bias in the fd estimates ranges from 0.004 to 0.0 for $N = 2^6$ and 2^{14} , respectively. In contrast, the average bias in the estimates from fGn processes actually increases with increasing N . In particular, the range of average bias in the fGn estimates is 0.0004 to 0.002 for $N = 2^6$ and 2^{14} , respectively. It is important to note that the average bias in the fGn estimates reaches approximately 0.002 at $N = 2^{10}$ and remains near that value for increasing N .

5. A modified MLE method for fd applied to fGn

Since the MLE method assumes a fd process, the bias in the ML estimates of H for fGn processes of every length considered confirms the presence of structural features in fGn that

are distinct from those in fd processes. The apparent approach of the bias to a non-zero value (0.002) as the length of the realization increases supports this contention. Because the MLE method in this case assumes a fd process, any structure in the fGn process that cannot be modelled by a fd process (with the same H) adversely affects the estimates. This is because H is the only parameter used to describe the fd process; and, therefore, all structure is accounted for by this one parameter. To address this issue, the fARIMA(p, d, q) model with $p > 0$ is proposed as the underlying model of the modified MLE method. It is assumed that the addition of autoregressive parameters will provide increased flexibility in modelling the features of fGn that are not captured by the fd model. If the bias in estimates of H from fGn processes is in fact due to structure present in fGn that is not in fd, then incorporation of additional parameters into the fARIMA model used in the MLE method may reduce the bias as the additional parameters may account for structure unique to fGn.

The average MSE was evaluated for the MLE method with the underlying model fARIMA($p, d, 0$) for $p = 1, 2, 3, 4$. The results are given in Table 1. As Table 1 shows, the average MSE increases when $p = 1$ (as compared to $p = 0$) for both fd and fGn processes of length $N = 2^{10}$. However, the average MSE decreases by an order of 10 in the fGn estimates of length $N = 2^{11}$. In addition, for both the fd and fGn processes, there is a decrease in the average MSE with increasing N . Similar characteristics are observed for models of higher orders and, for reasons of parsimony, the underlying model of order (1, $d, 0$) is used. The *modified MLE* (*M-MLE*) method refers to the *S-MLE* method adapted to correct for the observed bias in estimates of fGn. The underlying model in *M-MLE* is a fARIMA(1, $d, 0$). Using $p = 1$ specifies the order of the autoregressive (AR) component of a fARIMA model, and incorporation of this term allows for flexible modelling of short term properties of the process [1]. Presumably a viable alternative approach also capable of modifying the short memory properties would be to use a moving average component with $q = 1$.

5.1. Numerical evaluation of M-MLE method using the fd model: mean and bias of estimates of H

The results of the M-MLE algorithm for increasing H and N are shown in Fig. 3. The *M-MLE* method gave reliable estimates (average MSE < 0.0009) of H for both fd and fGn signals of length greater than or equal to $N = 2^{11}$. As in the case of *S-MLE*, the average standard deviations of the estimates from both fd and fGn processes decay at identical rates and thus, the decrease in the average MSE is due to a decrease in average bias. In contrast to the *S-MLE* method, the average bias for estimates from both processes approaches zero with increasing N .

6. Discussion

Most of the methods that are available for estimating H are developed to estimate H from a fGn or fBm process. As a result, to measure long-memory properties of a real signal, one must first identify the process as noise or motion before a particular method is used. Cannon et al. [12] give an example of the errors in the estimates if an erroneous identification is made by using the SWV method (which is developed to estimate H from fBm) applied to fGn and DISP (which gives reliable estimates for fGn) used on fBm. The H estimates from the SWV method are highly biased toward zero; and for every input H from 0.1 to 0.9, the estimated H obtained from DISP is between 0.9 and 1.0. Raymond and Basingthwaight [5] demonstrate theoretically why this is the case. Fortunately, such errors are not often encountered in practice as it is possible to identify differences between fGn and fBm in many cases. In particular, the former process is stationary, whereas the latter is not. Although detecting stationarity can be difficult due to long-memory properties which may appear as non-stationarities, methods do exist to classify fGn and fBm processes [12].

In contrast, the differences between fd and fGn processes cannot be readily identified. As a result, it is quite possible in practice that a technique to estimate H , which assumes a fd process, is applied to a fGn. As demonstrated here, incorrect estimates are obtained when this is the case, even for relatively large sample sizes. To address this issue, we have proposed a method which provides robust estimates of H for both fd and fGn processes of sufficient length. The bias in the estimates from the method decreases with increasing sample sizes, as expected if appropriate. The method should not be applied to short series, however, as a dramatic increase in bias is observed for both fd and fGn processes. This may be due to the fact that for short series, the distinction between the short- and long-term properties of the process are not as apparent. In this case estimation of the AR parameter, which models short term properties, is adversely affected by the long memory properties of the series. Fig. 4 shows that this is the case. The average estimated AR parameter is shown as a function of series length for both fd and fGn. For short series, the average AR parameter from fd processes is non-zero. It is not until the sample size is relatively large that a zero autoregressive parameter is estimated. For fGn, the autoregressive parameter does not approach zero even for series of very large length, except for a purely Gaussian random process with $H = 0.5$. This supports the contention that the bias in the estimates of H from fGn was due, at least in part, to structure unique to fGn and that such structure is accounted for by addition of the AR parameter into the fd based MLE method.

Although incorporation of the AR parameter reduces the bias in the long-memory parameter toward zero for long fd and fGn series, for short series, the AR parameter is incorrectly estimated as non-zero in the fd process (Fig. 4 left panel) which increases the error in estimating H from fd. Constraining the value of the autoregressive parameter may reduce the error; and it may be possible to determine the value of the constraint which provides maximum flexibility in modelling fGn, but minimizes the error in describing a fd process by a fARIMA(1, d , 0) model. An understanding of the relationship between fGn and fARIMA parameters could lead to a robust test for distinguishing between the two processes, thereby facilitating the use of already established and tested methods of analysis. It should be noted that *M-MLE* utilizes an already implemented method developed for fd processes, thus suggesting that other methods to estimate H which assume some underlying parametric form may be modified if necessary to accommodate other processes not assumed during development. In particular, it has been observed (Cannon, unpublished) that fGn-based MLE methods provide robust estimates of H from fGn series, as expected. Further analysis is required to determine the applicability of the method (or a related modified method) to fd processes. Finally, in reality, it is often the case that a signal of interest is neither fd nor fGn, but shares characteristics of both. Since the *M-MLE* method provides robust estimation of H for long signals, it should be appropriate for long linear stationary signals that are similar to fd or fGn. The applicability of *M-MLE* to linear stationary processes other than fd and fGn is currently under investigation.

7. Conclusion

The *S-MLE* method was evaluated and shown to provide unbiased estimates of H for fd processes of length 2^{11} or greater. However, *S-MLE* gives biased estimates for fd processes of length less than 2^{10} and for fGn processes of every length considered. The *M-MLE* method is a robust method for estimating H from both fd or fGn processes of length greater than or equal to 2^{11} ; but is much less reliable than the *S-MLE* method for processes of length less than 2^{11} .

References

1. Beran, J. Statistics for Long-Memory Processes. New York: Chapman & Hall; 1994.

2. Granger CW. *Econometrica*. 1966; 34:150.
3. Hosking JR. *Water Resources Res.* 1984; 20:1898.
4. Diggle, PJ. *Time series: a Biostatistical Introduction*. New York: Oxford University Press; 1990.
5. Raymond GM, Bassingthwaighte JB. *Phys. A*. 1999; 265:85.
6. Hurst HE. *Trans. Am. Soc. Civ. Eng.* 1951; 116:770.
7. Mandelbrot BB, van Ness JW. *SIAM Rev.* 1968; 10:422.
8. Marsh DJ, Osborn JL, Cowley AW Jr. *Am. J. Physiol.* 1990; 258:F1394. [PubMed: 2337155]
9. Wagner CD, Persson PB. *Am. J. Physiol.* 1995; 268:H621. [PubMed: 7864187]
10. Bassingthwaighte JB. *News Physiol. Sci.* 1988; 3:5. [PubMed: 20871797]
11. Bassingthwaighte JB, Beyer RP. *Physica D*. 1991; 53:71.
12. Cannon MJ, Percival DB, Caccia DC, Raymond GM, Bassingthwaighte JB. *Physica A*. 1997; 241:606. [PubMed: 22049250]
13. Fougere PF. *J. Geophys. Res.* 1985; 90:4355.
14. Schmittbuhl J, Vilotte JP, Roux S. *Phys. Rev. E*. 1995; 51:131.
15. Bassingthwaighte JB, Raymond GM. *Ann. Biomed. Eng.* 1994; 22:432. [PubMed: 7998689]
16. Schepers HE, van Beek JHGM, Bassingthwaighte JB. *IEEE Eng. Med. Biol.* 1992; 11:57.
17. Bassingthwaighte JB, Raymond GM. *Ann. Biomed. Eng.* 1995; 23:491. [PubMed: 7486356]
18. Peitgen, HO.; Saupe, D. *The Science of Fractal Images*. New York: Springer; 1988.
19. Percival, D.; Walden, AT. *Spectral Analysis for Physical Applications: Multitaper and Conventional Univariate Techniques*. Cambridge: Cambridge University Press; 1993.
20. Caccia D, Percival D, Cannon M, Raymond G, Bassingthwaighte J. *Physica A*. 1997; 246:609. [PubMed: 22049251]
21. Davies RB, Harte DS. *Biometrika*. 1987; 74:95–101.
22. Percival D. *Comput. Sci. Stat.* 1992; 24:534.
23. Box, GP.; Jenkins, G. *Time Series Analysis: Forecasting and Control*. San Francisco: Holden-Day; 1970.
24. Granger CW, Joyeux R. *J. Time Ser. Anal.* 1980; 1:15.
25. Bassingthwaighte, JB.; Liebovitch, LS.; West, BJ. *Fractal Physiology*. New York: Oxford University Press; 1994.
26. Anderson, OD. *Time Series Analysis and Forecasting: The Box–Jenkins approach*. London: Butterworth; 1976.
27. Chatfield, C. *The Analysis of Time Series*. London: Chapman & Hall; 1996.
28. Haslett J, Raftery AE. *JRSS.* 1989; C38:1.
29. Hosking JR. *Biometrika*. 1981; 68:165.
30. Ansley CF. *Biometrika*. 1979; 66:59.
31. S-PLUS Guide to Statistical and Mathematical Analysis, Version 3.3, StatSci Division. Seattle, Washington: MathSoft, Inc.; 1995.

Fig. 1. Time plots (upper), ACFs (middle, log-log scale), and power spectra (lower, log-log scale) for a fd (left) and fGn process (right) where $d = 0.2$ and $H = 0.7$, respectively. The theoretical forms of the ACFs for fd and fGn are given by the dotted lines in the plots of the sample ACFs. Recall that $d = H - 0.5$. In each simulation, $N = 2^{12}$. The spectra are plotted as functions of frequency, f , where $f = \omega/2\pi$.

Fig. 2. Means of *S-MLE* estimates of the Hurst coefficient for 100 time series of fd (left) and fGn (right) for $H = 0.5, 0.6, \dots, 0.9$, and for each series of length $N = 2^6, 2^7, \dots, 2^{14}$. Departure from the dashed lines shows bias in the means of the estimates. One-sided standard deviations of the estimates are given by the perpendicular lines for $H = 0.5, 0.7$ and 0.9 .

Fig. 3.

Means of M -MLE estimates of the Hurst coefficient for 100 time series of fd (left) and fGn (right) for $H = 0.5, 0.6, \dots, 0.9$, and for each series of length $N = 2^6, 2^7, \dots, 2^{14}$. Departure from the dashed lines shows bias in the means of the estimates. One-sided standard deviations of the estimates are given by the perpendicular lines for $H = 0.5, 0.7$ and 0.9 .

Fig. 4. Average of the AR parameter, α_1 , for 100 time series of fd (left) and fGn (right) for $H=0.5, 0.6, \dots, 0.9$, and for each series of length $N=2^6$ to 2^{14} . The value of α_1 was obtained via the MLE estimates using the *M-MLE* method.

Table 1

Average MSE of Hurst coefficient estimates obtained by fitting fARIMA models of order $(p, 1, 0)$, $p = 0, 1, 2, 3, 4$, to fd (upper) and fGn (lower) processes of length $N = 2^6, 2^7, \dots, 2^{14}$

Series length, N		2^6	2^7	2^8	2^9	2^{10}	2^{11}	2^{12}	2^{13}	2^{14}
fd Processes										
0	0.01287	0.00643	0.00242	0.0011	0.00058	0.00028	0.00013	0.00006	0.00004	
1	0.03961	0.0233	0.00803	0.00395	0.00157	0.00074	0.00035	0.00018	0.00011	
2	0.05081	0.03643	0.01626	0.00768	0.00298	0.00129	0.00064	0.00031	0.00017	
3	0.05424	0.04165	0.02368	0.01279	0.00471	0.00211	0.00092	0.00041	0.00024	
4	0.0539	0.04728	0.03131	0.01685	0.00694	0.00273	0.0012	0.00055	0.00032	
p	2^6	2^7	2^8	2^9	2^{10}	2^{11}	2^{12}	2^{13}	2^{14}	
fGn Processes										
0	0.00786	0.00389	0.00265	0.00229	0.00219	0.00245	0.00248	0.0025	0.00245	
1	0.04297	0.02803	0.01344	0.00537	0.00227	0.00097	0.00049	0.00025	0.00011	
2	0.05024	0.03711	0.01876	0.00828	0.00353	0.0015	0.00061	0.00037	0.00018	
3	0.05333	0.04419	0.02734	0.01343	0.00563	0.0022	0.00087	0.00051	0.00028	
4	0.05494	0.04924	0.03375	0.01745	0.00765	0.00287	0.00123	0.00061	0.00039	