Physica A 389 (2010) 5573–5580

Contents lists available at ScienceDirect

Physica A

journal homepage: www.elsevier.com/locate/physa

Preservation of long range temporal correlations under extreme random dilution

Dror Mirzayof, Yosef Ashkenazy *

Solar Energy and Environmental Physics, BIDR, Ben-Gurion University, Israel

ARTICLE INFO

Article history:
Received 28 April 2010
Received in revised form 25 June 2010
Available online 6 September 2010

Keywords:
Long range correlations
Random dilution
Time series
Heart interbeat interval
DFA

ABSTRACT

Many natural time series exhibit long range temporal correlations that may be characterized by power-law scaling exponents. However, in many cases, the time series have uneven time intervals due to, for example, missing data points, noisy data, and outliers. Here we study the effect of randomly missing data points on the power-law scaling exponents of time series that are long range temporally correlated. The Fourier transform and detrended fluctuation analysis (DFA) techniques are used for scaling exponent estimation. We find that even under extreme dilution of more than 50%, the value of the scaling exponent remains almost unaffected. Random dilution is also applied on heart interbeat interval time series. It is found that dilution of 70%–80% of the data points leads to a reduction of only 8% in the scaling exponent; it is also found that it is possible to discriminate between healthy and heart failure subjects even under extreme dilution of more than 90%.

1. Introduction

Natural time series often vary on many time scales resulting in long range temporal correlations (e.g., [1]). By long range temporal correlations we refer, e.g., to the cases in which the auto-correlation function decays as a power-law, indicating that even temporally far apart data points are statistically correlated. The power-laws are characterized by the so called scaling exponents. The scaling exponents may reveal the linear and nonlinear nature of the time series (e.g., [2,3]). Scaling exponents have been used to characterize different states of natural systems, such as distinguishing between healthy and heart failure subjects (e.g., [4,5]), characterizing different human gait activities [6], and characterizing interbeat interval time series during different sleep phases [7]. In addition, they have been used in other disciplines such as economics (e.g., [8]) and geoscience (e.g., [9–12]).

In many cases, there are difficulties in extracting an evenly sampled natural time series. For example, in the case of heart interbeat intervals, sometimes there are beats that are not annotated as normal beats and must be excluded [5]. In other cases, the noise level is large such that there is a need to exclude the "noisy" data points. In addition, instrumental problems sometimes lead to gaps in time series. Also, in some cases, the nature of the studied time series is such that the time series is uneven; this is the case for example for earthquake data, in which events above a certain threshold are considered (e.g., [13–15]). The above causes and others result in time series with missing data points.

The aim of the present work is to study the effect of randomly missing points on the estimation of the scaling exponents. To achieve this goal we generate artificial long range correlated time series with known scaling exponents. We then randomly "dilute" these time series and check the effect of the dilution on the scaling exponent. Motivated by the above (previous paragraph), we study two dilution procedures, the first with random deletion of data points and the second with deletion of data points above a certain threshold. We also apply random dilution on heart interbeat interval time series

* Corresponding author.
E-mail address: ashkena@bgu.ac.il (Y. Ashkenazy).

0378-4371/$ – see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.physa.2010.08.035

and show that it is possible to discriminate between healthy and heart failure subjects using the scaling exponents, even under drastic dilution. In a recent work [16], the effect of extreme truncation on the scaling exponent of correlated and anti-correlated time series was analyzed; we became aware of this work only in a late stage of the review process.

The effect of randomly removing segments from synthetic time series on the detrended fluctuation analysis (DFA) scaling exponent was studied by Chen et al. in Ref. [17]. It was shown that anti-correlated time series exhibit crossover (from anti-correlated to white noise scaling exponent) after applying the "cutting out" procedure. When this procedure was applied on correlated time series, the exponent was preserved. In the analysis of Ref. [17] up to 50% of the time series was randomly excluded. The present study extends the results of Ref. [17] as follows: (i) We study the effect of dilution on the scaling exponents using both Fourier transform and DFA techniques. (ii) We study the effect of extreme dilution of up to 99% of the time series. (iii) We apply the dilution procedure on heart interbeat interval time series and show that it is possible to discriminate between healthy and heart failure patients even under extreme dilution of more than 90%. (iv) We apply an additional dilution procedure in which only data points that are above a certain threshold are considered; this dilution may be associated with earthquake data (see, e.g., Ref. [15]).

The paper is organized as follows: We first provide in Section 2 a brief background for the methods used in this study (namely, Fourier transform, DFA, and artificial data generation techniques); then the effect of random dilution on artificial time series is studied (Section 3); following is an application to heart interbeat interval time series (Section 4); a summary concludes the paper (Section 5).

2. Methods

2.1. Algorithms for the estimation of scaling exponents

There are several methods to measure and characterize the scaling properties of long range correlated time series $\{x_i\}$. Below we briefly describe a few of them.

One basic way to estimate scaling exponents is through the autocorrelation function $A(l)$ (e.g., [1,18]):

$$A(l) = \frac{1}{N-l}\sum_{i=1}^{N-l} x_i x_{i+l}, \tag{1}$$

where N is the number of data points in the time series. In some cases, the autocorrelation function follows a scaling law such that $A(l) \sim l^{-\gamma}$ with $0 < \gamma \leq 1$, where γ is called the *scaling exponent*. Uncorrelated (or short range) time series have exponent $\gamma = 1$, while smaller γ values indicate stronger correlations.

There are two main difficulties in estimating a scaling exponent using the autocorrelation function: (i) The autocorrelation function is often noisy and thus the uncertainty associated with the estimation of the scaling exponent is large, and (ii) the range of the estimated scaling exponent is limited; i.e., γ is bounded between zero and one.

Thus, it is more common to estimate scaling exponents using the Fourier power spectrum $P(f)$

$$P(f) = |\mathcal{F}\{x_i\}|^2, \tag{2}$$

where $\mathcal{F}\{x_i\}$ is the discrete Fourier transform of the series x_i. When $P(f) \sim f^{-\beta}$, the time series follows a scaling law and β is the Fourier transform scaling exponent. $\beta = 0$ indicates that the time series is uncorrelated; a long range correlated time series is characterized by a scaling exponent $\beta > 0$ and a long range anti-correlated series is identified by $\beta < 0$. In principle, unlike the exponent γ of the autocorrelation function, the β exponent is not limited to a specific range.

In spite of the advantages of the Fourier power spectrum in estimating scaling exponents, the Fourier transform technique can lead to less reliable estimation when trends exist in the data, as might be the case for natural time series. The DFA [4,2,7] aims to overcome this deficiency. In DFA a time series $\{x_i\}$ is first integrated: $y(k) = \sum_{i=1}^{k}[x_i - \bar{x}]$. Then, the integrated series is divided into boxes of equal length (or scale) s, and a piecewise continuous polynomial trend $y_s(k)$ (i.e. a trend which is continuous within each box) is fitted to the series. Finally, for each scale s, the detrended fluctuation function $F(s)$ is computed:

$$F(s) = \sqrt{\frac{1}{N}\sum_{k=1}^{N}[y(k) - y_s(k)]^2}. \tag{3}$$

A DFA procedure using a polynomial trend $y_s(k)$ of order n is also called DFA of order n and denoted as DFAn. Time series with long range power-law temporal correlations will have the fluctuation function $F(s) \sim s^{\alpha}$, where α is the scaling exponent. For uncorrelated data, the integrated series is a random walk and thus $\alpha = \frac{1}{2}$. Correlated time series have $\alpha > \frac{1}{2}$ while anti-correlated time series have $\alpha < \frac{1}{2}$. We note that the α exponent is always larger than zero and the upper bound is determined according to the order of the polynomial fitting used in the DFA procedure. The scaling exponents obtained by DFA1 (DFA with linear fitting) are limited by two and that of DFAn are limited by $n + 1$. In this study DFA2 is used.

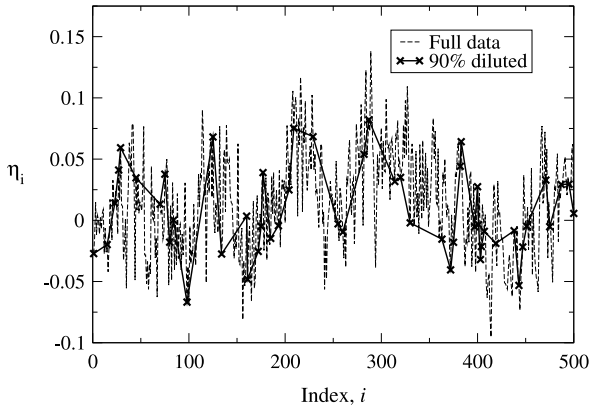

Fig. 1. Example of an artificial dataset before (dashed line) and after (solid line) dilution of 90% of the data points. The original dataset has a scaling exponent $\alpha = 0.9$.

The scaling exponents of the autocorrelation function, the Fourier power spectrum,and the DFA (γ, β, and α respectively) are related to each other as follows [4,18]:

$$\alpha = \frac{1}{2}(\beta + 1) = 1 - \frac{\gamma}{2}. \tag{4}$$

Below we will mainly use the α DFA exponent as it seems to be the most widely used exponent. We use the above relations to express the different scaling exponents in terms of the α exponent.

2.2. Generation of artificial data

Following Ref. [19], a *normally* distributed fractal time series with a given α exponent can be generated as follows:

i. Generate an uncorrelated, normally distributed random series x_i (white noise).
ii. Fourier transform x_i: $\hat{X}_\omega = F\{x_i\}$, where ω is the frequency.
iii. Introduce long range correlation to the series by multiplying the Fourier coefficients $\hat{X}_\omega$ by factors that rescale with exponent $\beta/2$: $\hat{Y}_\omega = \frac{1}{\omega^{\frac{\beta}{2}}}\hat{X}_\omega$. To prescribe a DFA exponent α we use the relation $\beta = 2\alpha - 1$, following the relation given in Eq. (4).
iv. Inverse Fourier transform the series from (iii), $y_i = F^{-1}\{\hat{Y}_\omega\}$.

Generation of an *exponentially* distributed time series with a given α exponent can be done using the following scheme:

i Create a correlated *normally* distributed time series (using the scheme described above).
ii Sort the Gaussian time series in ascending order, keeping pointers to return to the original ordering.
iii Generate an exponentially distributed uncorrelated time series and sort it in ascending order.
iv Reorder the series from (iii) using the Gaussian data reverse ordering pointers from (ii).

The restriction of exponential distribution may sometime lead to correlated time series, with a scaling exponent that is not exactly equal to the scaling exponent of the Gaussian distributed time series from (i); however, see [20] for a more advanced technique for generating exponentially distributed correlated time series.

2.3. Dilution procedure

We applied and tested the effect of both random and threshold-based dilutions on the scaling exponents of time series.

Random dilution of p% of the data is done by passing over all data points, deleting each point with probability $d = \frac{p}{100}$ or leaving it with probability $1 - \frac{p}{100}$. An example of a long range correlated artificial time series (with an exponent $\alpha = 0.9$) before and after random dilution of 90% data points is shown in Fig. 1.

For threshold-based dilution of p% of a dataset of size N, we first set the threshold t by sorting the data points, and setting t to be the value of point number $\lfloor \frac{p}{100}N \rfloor$; then all the points with values smaller than t are deleted from the original (unsorted) series.

3. Effect of dilution on artificial time series

We use two schemes of data generation and dilution. The first is random deletion of points from correlated, normally distributed time series. This scheme corresponds to random loss of data points. The second scheme is motivated by

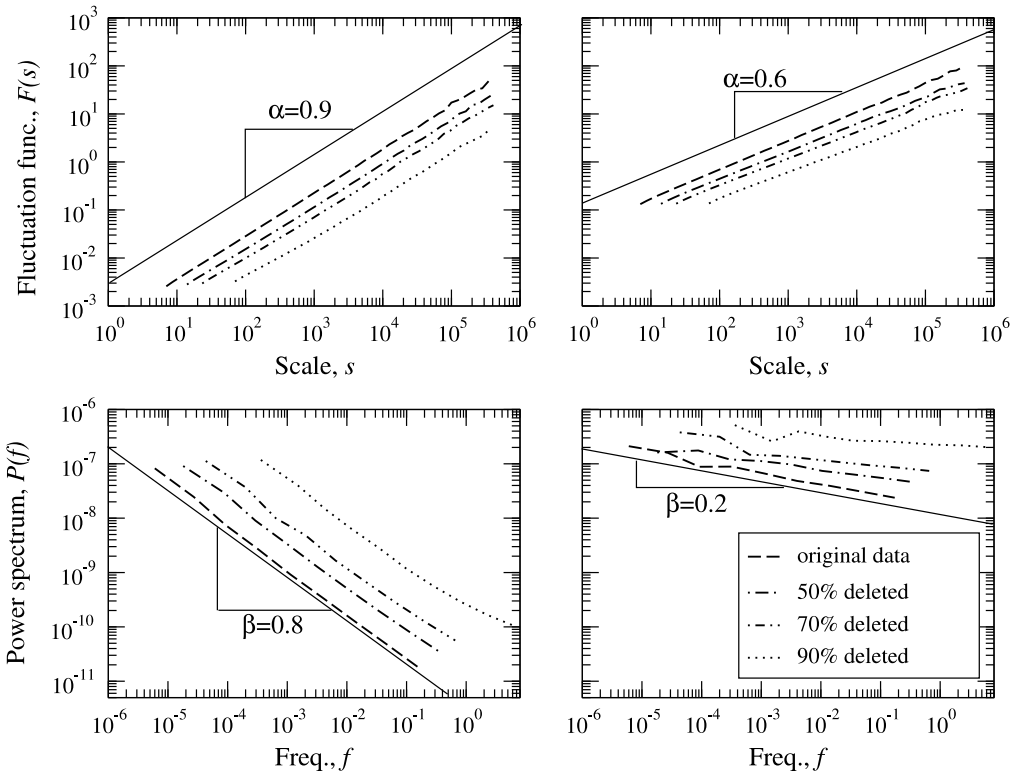

Fig. 2. Preservation of scaling exponents under dilution of data, for two datasets of $2^{21} = 2097,152$ data points, as manifested in the linear curve on log–log plots of the scaling functions for diluted and undiluted time series. The left and right columns of the figure correspond to artificial datasets with scaling exponents of $\alpha = 0.9$ and $\alpha = 0.6$ respectively; the upper panels depict the DFA fluctuation function vs. window scale [$F(s)$], while the lower panels depict the Fourier power spectrum vs. frequency [$P(f)$] for the two datasets. It is evident that all the log–log plots are close to being linear and the slopes do not change significantly with dilution of the datasets. Still, for smaller α, the linearity of the $P(f)$ log–log plot starts to become overly sensitive to data dilution.

earthquake data analysis [13,21,22,14]; here exponentially distributed data sets are considered where all points that are smaller than some threshold are excluded (as in, Ref. [15]). The second scheme is therefore a threshold-based dilution of points from a correlated, exponentially distributed time series.

3.1. Random dilution of normally distributed data

We computed the scaling functions $F(s)$ and $P(f)$ for *normally* distributed artificial time series with different values of α, generated as described in Section 2.2. For each series we then repeated the computation after *random* dilutions of 50%–99% of the data points, as described in Section 2.3.

The four panels of Fig. 2 depict the log–log plots of two different time series with different scaling exponents. In the upper (lower) panels, the plots of $F(s)$ ($P(f)$) are shown; the left panels are for a time series with $\alpha = 0.9$ ($-\beta = -0.8$), and the right are for $\alpha = 0.6$ ($-\beta = -0.2$). In each panel there is one graph of the original (undiluted) series and three graphs of diluted series that correspond to dilutions of 50%, 70% and 90% of the data points. An additional continuous straight line is added as a reference to each panel, plotted with the expected slope. It is evident from this figure that the curves of the diluted series are approximately linear, with the approximately expected slopes; the DFA curves exhibit better agreement with the undiluted curve compared to the curves of the Fourier power spectra $P(f)$. In addition, $P(f)$ performance seems to be worse for the smaller α exponent.

The procedure described above was repeated with time series prepared with different values of α (ranging from 0.5 to 1.5), for each of which an ensemble of ten different realizations was generated. The slope for each series, indicating its scaling exponent α (or β), was extracted using linear regression of $\log F(s)$ vs. $\log s$ (or $\log P(f)$ vs. $\log f$). Each of these series was reanalyzed after random deletion of data with dilution rates d in the range 50%–99%. The β exponents were translated to their equivalent α using Eq. (4). For each combination of exponent α and dilution d, the average and standard deviation of the scaling exponent over the ensemble of ten realizations were computed. These statistics are presented in Fig. 3, in two different ways for each of the two algorithms. On the left, the dependence of the computed α on the level of dilution is presented for 11 different original α values. On the right are plots of the computed α vs. the value of α with which the time series was prepared. The upper panels correspond to α computed directly by DFA, while the lower panels show the Fourier

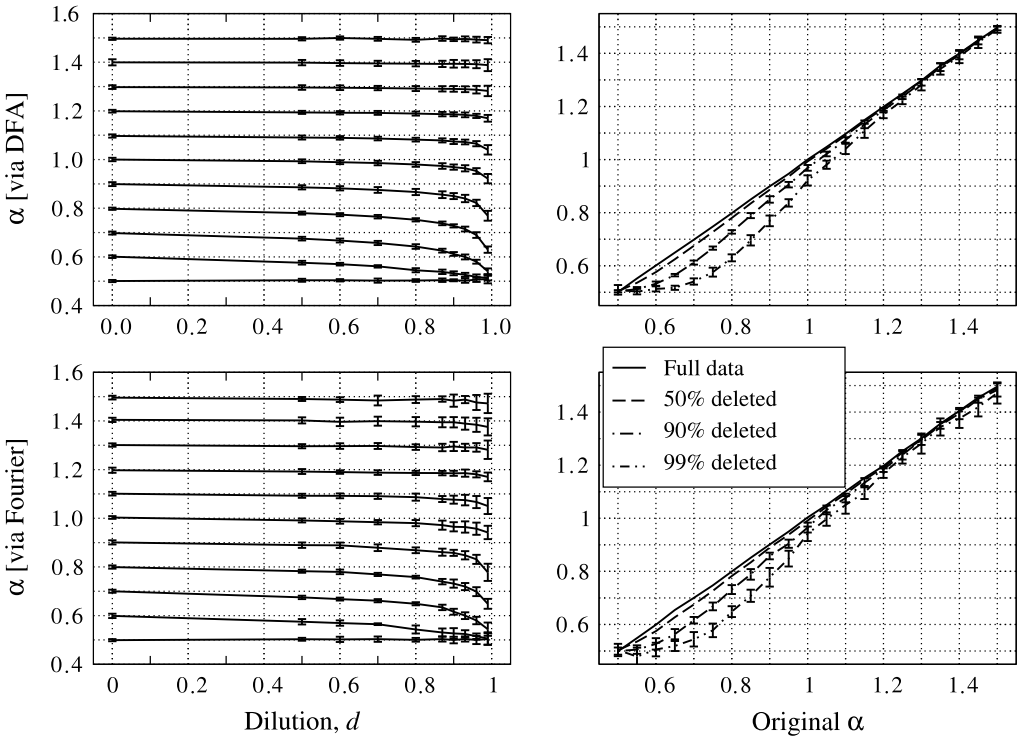

Fig. 3. The effect of different rates of random dilution on the scaling exponents: the curves are based on normally distributed artificial data generated with different α values. The β exponents of the Fourier power spectrum were converted to the corresponding α using Eq. (4). For each original value of α the measured exponents were averaged for 10 realizations, each of $2^{21} = 2097,152$ data points.

power spectrum exponents expressed in terms of α using Eq. (4). It is evident that for time series generated with $\alpha \geq 1.0$ and diluted up to 50%, both algorithms yield exponents that are negligibly different than those of the original undiluted series.

We repeated the above experiments on much shorter time series of length $2^{15} = 32,768$ data points (see Fig. 4) and obtained similar results as for the much longer time series studied above.

3.2. Threshold-based dilution of exponentially distributed data

The experiments described in the previous subsection were repeated for *exponentially* distributed artificial data (described in Section 2.2) with *threshold-based* dilution (described in Section 2.3).

The results are shown in Fig. 5, in which the panels are equivalent to those of Fig. 3. It can be seen that the exponents computed by both algorithms are systematically smaller then expected (although the difference is small), even for undiluted series. In contrast to the case of randomly diluted normal series (Fig. 3), the exponents are less preserved. Nevertheless, for α in the range $0.5 < \alpha < 1.5$ and dilution $d \leq 0.7$ (70%), the reduction due to dilution seems to be, to a good approximation, linear in d, and it might indicate the possibility to easily reconstruct the original exponent if one can estimate the rate of dilution. The results presented in this sub-section are consistent with a recent study of earthquake time series in which it was shown that the scaling exponent is preserved when considering data points above a certain earthquake magnitude [15].

4. Effect of random dilution on heart interbeat interval time series

We examined a database of 30 subjects, of which 18 were healthy and 12 were heart failure patients; these time series were downloaded from the physionet web-page (www.physionet.org) [5] and were studied in the original DFA paper [4] and many other studies (e.g., [3,23]). The time series span an approximately 24 h period. The scaling exponents were approximated for window scales larger than 20 beats, a scale that is larger than the crossover in the DFA curves for interbeat interval time series.

We tested the effect of data dilution on the classification between the heart failure and the healthy individuals. We computed separately the mean and standard deviation of α for the healthy and heart failure groups. The same was repeated after diluting the time series. The results are shown in the left panels of Fig. 6. We note that DFA separates well between the two groups, and does not lose its separation ability even when almost all the data is randomly diluted. The Fourier transform loses its ability to separate between the two groups for dilution levels larger than 90%.

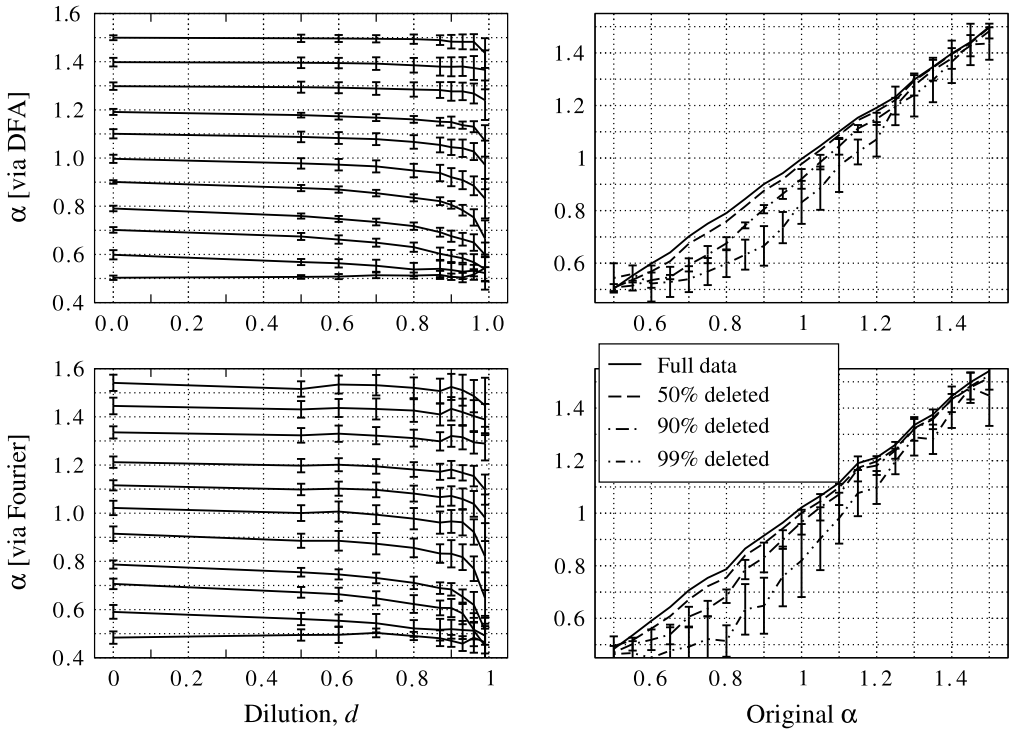

Fig. 4. Same as Fig. 3 but for time series of length 32,768 data points.

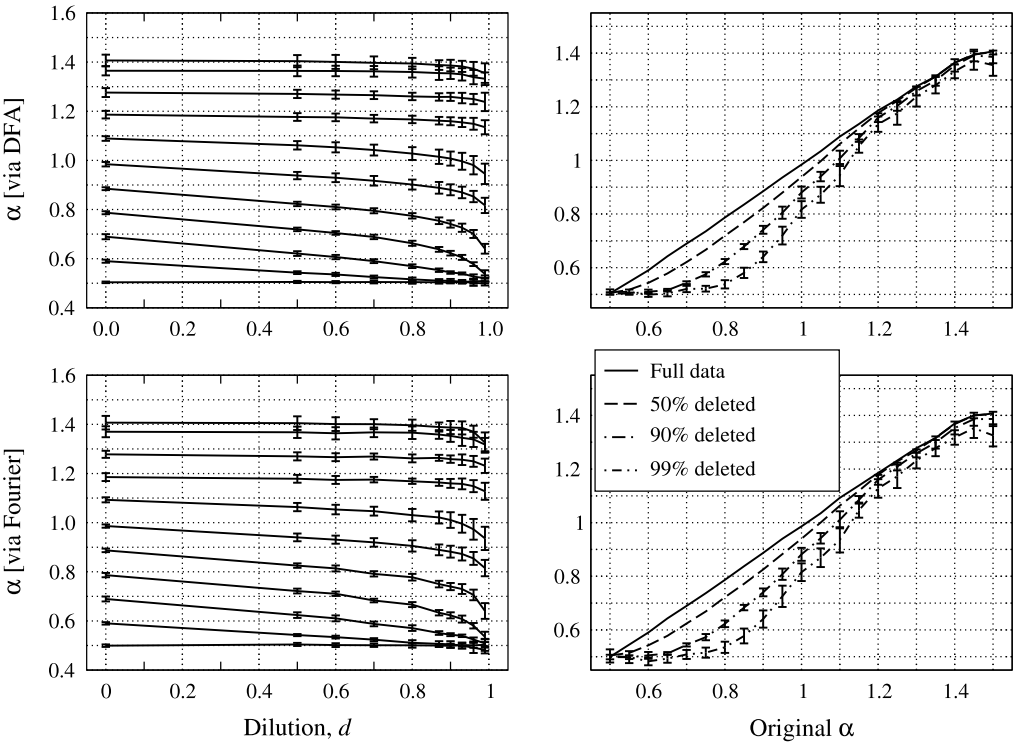

Fig. 5. Same as Fig. 3, but for artificial *exponentially* distributed datasets, using *threshold-based* dilution (i.e., data points that are smaller than the threshold are excluded).

Another perspective regarding the effect of dilution on scaling exponents of the heart interbeat interval time series can be obtained by computing the relative change $\frac{\alpha-\alpha_0}{\alpha_0}$ in the scaling exponent α due to dilution, where α_0 is the exponent

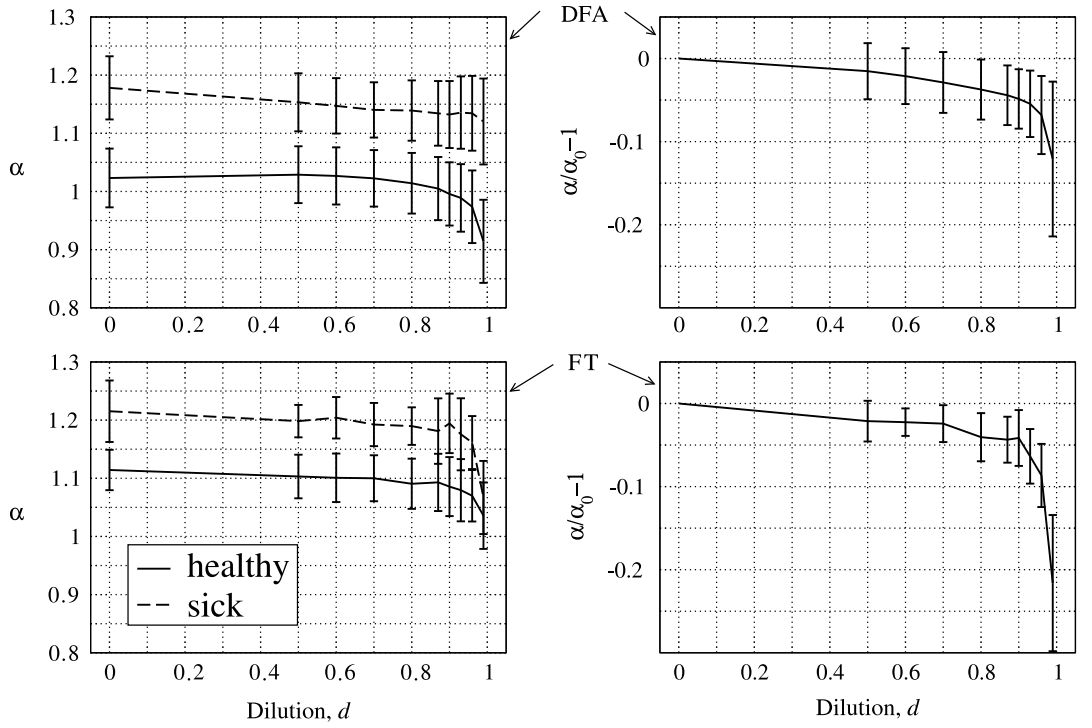

Fig. 6. The effect of dilution on the scaling exponents of heart interbeat interval time series, each spanning approximately 24 h. The left panels depict the mean scaling exponent ± one standard deviation of the healthy and heart failure groups, for different dilution levels; it is evident that both methods, DFA (upper panel) and Fourier transform (lower panel) yield good separation between the two groups even for high level of dilution. The healthy group consists of 18 individuals while the heart failure group consists of 12 individuals. The right panels depict the relative shift in the scaling exponent $\frac{\alpha-\alpha_0}{\alpha_0}$ for the healthy individuals as computed by DFA (up) and Fourier transform (down), where α_0 is the original exponent without dilution.

without dilution; this was applied to the group of 18 healthy subjects. The results are presented as a function of the dilution level d in the right panels of Fig. 6. It can be seen that for dilutions up to 70%–80% both algorithms yielded a reduced mean exponent by no more than 8%, with the Fourier power spectrum being more robust than DFA.

5. Summary

We investigated the effect of random exclusion of data points on the scaling properties computed by DFA and Fourier power spectrum. We examined artificially generated normally distributed time series under random dilution, motivated by random corruption of data points. It was found that for scaling exponent α in the range $\alpha \geq 1.0$ and dilution level up to 50%, the effect of dilution is very small. For smaller α values and/or higher dilution levels (at least up to 80%) the offset due to dilution is still minor, with no significant difference between DFA and Fourier transform techniques. The original α seems to be reconstructable if the dilution level is known. It is noteworthy that even for dilution as large as 90% the offset in the DFA scaling exponent α from the original value does not exceed 10%–15%.

Another type of artificial data and dilution scheme we tested was exponentially distributed time series and deletion of all data points below a certain threshold, motivated by earthquake data analysis. Here, the reconstruction of α was imprecise even for the undiluted series (most probably as a result of the restriction to exponential distribution of the time series), the effect of dilution is much stronger, but still systematic.

We also examined heart interbeat interval data, in which the relative change due to dilution was averaged for an ensemble of 18 healthy individuals. Our results indicate that for dilution levels less than 80%, the error in α does not exceed 8%. Another aspect is the preservation of the separation between the scaling exponent of a group of healthy subjects and heart impaired individuals. In this respect DFA seems to have an advantage over the power spectrum method, and the separability is preserved to a good extent even when 99% of the data are randomly excluded.

At a first glance the results reported here seem surprising. However, since the scaling exponent reflects the "fractal" nature of the time series, and since a basic feature of scale free time series is its statistical "self-similarity", the reported results are less surprising. From this perspective it is easier to understand why non-stationary time series (with $\alpha > 1$) exhibit better preservation of the scaling exponent under dilution, since their profile preserves its structure even after dilution; however, this is not the case for stationary time series.

There are known techniques to compute the power spectrum of unevenly sampled time series, such as the Lomb–Scargle periodogram technique [24,25]. In addition, it is possible to extend DFA to deal with unevenly sampled time series by

performing least-squares fitting where the x-axis values are assumed to be unevenly spaced. We applied both techniques on diluted time series and found similar scaling exponents to those calculated by conventional techniques.

In many cases, significant effort is devoted to extract as many reliable data points as possible. The results reported here indicate that, at least from the scaling exponent perspective, such an effort is not necessary, since the scaling exponents are hardly affected when a small proportion of the data points is excluded.

References

[1] A. Bunde, S. Havlin (Eds.), Fractals in Science, 2nd ed., Springer, Springer, Berlin, 1996.
[2] C.-K. Peng, S. Buldyrev, S. Havlin, M. Simons, H.E. Stanley, A.L. Goldberger, Mosaic organization of DNA nucleotides, Phys. Rev. E 49 (1994) 1685–1689.
[3] P.Ch. Ivanov, L.A.N. Amaral, A.L. Goldberger, S. Havlin, M.G. Rosenblum, Z.R. Struzik, H.E. Stanley, Multifractality in human heartbeat dynamics, Nature 399 (1999) 461–465.
[4] C.-K. Peng, S. Havlin, H.E. Stanley, A.L. Goldberger, Quantification of scaling exponents and crossover phenomena in nonstationary time series, Chaos 5 (1995) 82–87.
[5] A.L. Goldberger, L.A.N. Amaral, L. Glass, J.M. Hausdorff, P.C. Ivanov, R.G. Mark, J.E. Mietus, G.B. Moody, C.-K. Peng, H.E. Stanley, PhysioBank, PhysioToolkit, and PhysioNet: components of a new research resource for complex physiologic signals, Circulation 101 (2000) e215–e220. Circulation Electronic Pages: http://circ.ahajournals.org/cgi/content/full/101/23/e215.
[6] J.M. Hausdorff, C.-K. Peng, Z. Ladin, J.Y. Wei, A.L. Goldberger, Is walking a random walk? Evidence for long-range correlations in the stride interval of human gait, J. Appl. Physiol. 78 (1995) 349–358.
[7] A. Bunde, S. Havlin, J.W. Kantelhardt, T. Penzel, J.H. Peter, K. Voigt, Correlated and uncorrelated regions in heart-rate fluctuations during sleep, Phys. Rev. Lett. 85 (2000) 3736–3739.
[8] F. Wang, K. Yamasaki, S. Havlin, H.E. Stanley, Scaling and memory of intraday volatility return intervals in stock markets, Phys. Rev. E 73 (2006) 026117.
[9] D.L. Turcotte, Fractals and Chaos in Geology and Geophysics, 2nd ed., Cambridge University Press, Cambridge, 1997.
[10] J.D. Pelletier, Analysis and modeling of the natural variability of climate, J. Clim. 10 (1997) 1331–1342.
[11] J.W. Kantelhardt, E. Koscielny-Bunde, D. Rybski, P. Braun, A. Bunde, S. Havlin, Long-term persistence and multifractality of precipitation and river runoff records, J. Geophys. Res. 111 (2006) D01106.
[12] P. Huybers, W. Curry, Links between annual, Milankovitch and continuum temperature variability, Nature 441 (2006) 329–332.
[13] P. Bak, K. Christensen, L. Danon, T. Scanlon, Unified scaling law for earthquakes, Phys. Rev. Lett. 88 (2002) 178501.
[14] V. Livina, S. Havlin, A. Bunde, Memory in the occurrence of earthquakes, Phys. Rev. Lett. 95 (2005) 208501.
[15] S. Lennartz, V.N. Livina, A. Bunde, S. Havlin, Long-term memory in earthquakes and the distribution of interoccurrence times, Europhys. Lett. 81 (2008) 69001.
[16] Q.D.Y. Ma, R.P. Bartsch, P. Bernaola-Galvan, M. Yoneyama, P.C. Ivanov, Effect of extreme data loss on long-range correlated and anticorrelated signals quantified by detrended fluctuation analysis, Phys. Rev. E 81 (2010) 031101.
[17] Z. Chen, P.C. Ivanov, K. Hu, H.E. Stanley, Effect of nonstationarities on detrended fluctuation analysis, Phys. Rev. E 65 (2002) 041107.
[18] M. Makse, S. Havlin, M. Schwartz, H.E. Stanley, Method for generating long-range correlations for large systems, Phys. Rev. E 53 (1996) 5445–5449.
[19] T. Kalisky, Y. Ashkenazy, S. Havlin, Volatility of linear and nonlinear time series, Phys. Rev. E 72 (2005) 011913.
[20] J.F. Eichner, J.W. Kantelhardt, A. Bunde, S. Havlin, Extreme value statistics in records with long-term persistence, Phys. Rev. E 73 (2006) 016130.
[21] A. Corral, Local distribution and rate fluctuations in a unified scaling law for earthquakes, Phys. Rev. E 66 (2003) 031305.
[22] A. Corral, Long-term clustering, scaling and universality in the temporal occurrence of earthquakes, Phys. Rev. Lett. 92 (2004) 108501.
[23] Y. Ashkenazy, P.Ch. Ivanov, S. Havlin, C.-K. Peng, A.L. Goldberger, H.E. Stanley, Magnitude and sign correlations in heartbeat fluctuations, Phys. Rev. Lett. 86 (2001) 1900–1903.
[24] N.R. Lomb, Least-squares frequency analysis of unequally spaced data, Astrophys. Space Sci. 39 (1976) 447–462.
[25] W.H. Press, B.P. Flannery, S.A. Teukolsky, W.T. Vetterling, Numerical Recipies, Cambridge University Press, 1990.