PERGAMON

Computers in Biology and Medicine 29 (1999) 393–406

Computers in Biology and Medicine

www.elsevier.com/locate/compbiomed

Determining the Hurst exponent of fractal time series and its application to electrocardiographic analysis

Paolo B. DePetrillo[a,*], d'Armond Speers[b], Urs E. Ruttimann[c]

[a]*Laboratory of Clinical Studies, Section of Clinical Science, Unit of Clinical and Biochemical Pharmacology, Bethesda, MD, USA*

[b]*Intelligent Systems, Thomson Labs, Rockville, MD, USA*

[c]*Section of Brain Electrophysiology and Imaging, Division of Intramural Clinical and Biological Research, National Institute on Alcohol Abuse and Alcoholism, National Institutes of Health, Bethesda, MD, USA*

Received 15 December 1998; received in revised form 28 June 1999; accepted 9 August 1999

Abstract

An alternative regression-based method for estimating the Hurst coefficient of a fractal time series is proposed. A formal mathematical description of the methodology is presented. The geometric relationship of the algorithm to the family of self-similar fractal curves is outlined. The computational structure of the algorithm is optimal for generation of real-time estimates of H. We show that the method can be applied to biologically-derived time series such as the cardiac interbeat interval and we obtain estimates of H from several diverse electrocardiographic data sets. © 1999 Published by Elsevier Science Ltd. All rights reserved.

Keywords: Heart; Nonlinear dynamics; Fractals; Electrocardiography; Software

1. Introduction

Measurement of the signal dynamics of electrocardiographic interbeat interval (IBI) time series have been shown to be a powerful means of assessing the influences of autonomic tone on cardiac function [1,2]. Methodological problems remain in obtaining reliable estimates of these measures. In time-domain analyses, measures of dispersion, such as the standard

* Corresponding author. Tel.: +1-301-496-9420; fax: +1-301-402-0445.
E-mail address: pbdp@helix.nih.gov (P.B. DePetrillo)

0010-4825/99/$ - see front matter
PII: S0010-4825(99)00018-9

deviation, increase with increasing data length, making cross-study comparisons difficult. Dispersional measures also do not take into account the degree of temporal autocorrelation. Problems due to autocorrelations can be avoided by conducting the analysis in the frequency domain, using a Fast-Fourier Transform (FFT). However, FFT assumes that for the epoch investigated the time series remains stationary. This assumption is less likely to hold as longer

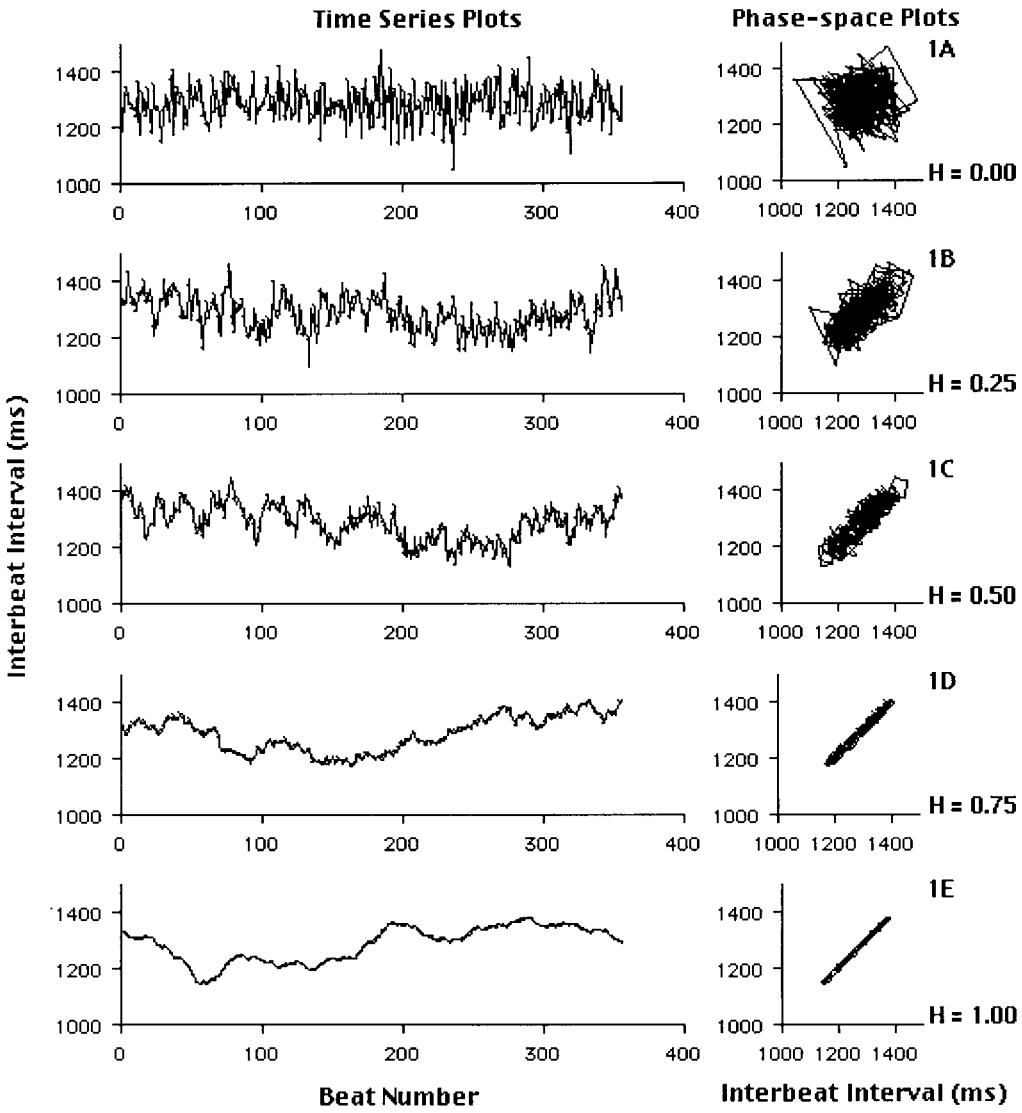

Fig. 1. Interbeat interval time series constructed using a spectral synthesis method to have defined noise characteristics, as noted by the H values. The means of the interbeat intervals were all set to 1267 and the standard deviations to 63, congruent with human subject data. The time series are plotted against beat number in the left column and in phase space in the right column.

time intervals are sampled. On the other hand, short data lengths are also problematic because the contributions of low frequencies to the overall power spectrum cannot be properly estimated. Consequently, the outcomes of applying either methodology are not easily interpreted when the recording time is relatively short, (<5 min), when the length of data recording varies significantly between individuals or if the time series is non-stationary.

Mandelbrot and Van Ness [3] formally related the Hurst exponent (H) to fractional Brownian motion and thus made the H available as a descriptor of the scaling behavior of fractal curves, called self-similarity. This was exploited to derive a direct geometric method of estimating H from a time series of interbeat intervals (IBI). The method is shown to be robust even for relatively short (data length < 1000) time series obtained from human subjects.

2. Background

The application of fractal models to biologically derived signals such as the electrocardiographic interbeat interval time series has greatly facilitated statistical analysis of these complex systems. Many biological signals demonstrate long-term correlations and (1/f)-type spectral behavior which greatly complicates their analysis. The development and study of a family of Gaussian random functions designated as fractional Brownian motions (fBm) by Mandelbrot and Van Ness [3], laid the mathematical foundations for the analysis of biological signals exhibiting such fractal, or long-range correlations.

The Hurst exponent, H, characterizes a fBm as follows. If we take a Gaussian function $B(t)$, where t is usually time and $B(t)$ is the magnitude at time t, then $[B(t_2) - B(t_1)]$ are the increments of this function with the variance:

$$\mathrm{var}[B(t_2) - B(t_1)]^2 \propto |t_2 - t_1|^{2H} \quad \text{with } 0 < H < 1$$

where the variance is determined from many samples of $B(t)$. Also note that when $H = 0.5$, the variance is now proportional to the time increment, which is what we expect from a randomly generated Gaussian function where the magnitude of the future increments is completely independent of past increments. However, for H values greater than or less than 0.5, there is an increasing statistical dependence of future increments on past increments.

We show a graphic demonstration of the meaning of H as it applies to a set of time series in Fig. 1. Visually inspecting a time-series with $H = 1.0$ in Fig. 1E, we see an overall relative smoothness to the graph of IBI versus time. As H tends towards 0, trends are more rapidly reversed, as shown in Fig. 1A, where there are large variations between adjacent values, which give it a very irregular look. When $H = 0.5$, as seen in Fig. 1C, the magnitude of the sequential points of the time series are independent and therefore uncorrelated. The time series exhibits properties of a random walk. Thus, H values approaching 0.5 from either extreme (0 or 1) are symptomatic of a breakdown in the long-range correlations of the signal. Also note that while the means and standard deviations of all the time series are equivalent (mean = 1267 and S.D. = 63), the graphs are dissimilar in their appearance, since time-domain measures lose all the information related to phase.

3. Description of method

We wanted to determine H for interbeat interval time series. The most powerful and direct methods rely on algorithms that operate on self-similar fractal curves, where variable scales are in the same units, such as distance in the xy-plane. However, interbeat interval time series are described by two scales: the time scale and the heart beat number. If these time series were fractal, they would be self-affine fractals, i.e. there would exist at least two scaling parameters, one for each scale, rather than one scaling parameter as would be seen for a self-similar fractal.

To circumvent this problem, we mapped the points of the time series into phase-space, as shown in Fig. 1. This embedding procedure results in a curve defined in n-space which now has the same units (time) for all scales. Using the following procedures, we take advantage of this fact to derive both the fractal dimension as well as the Hurst exponent of the time series.

3.1. Dimensional embedding procedure

A phase space is constructed by first choosing the dimensional embedding constant, n. The coordinates of each point within this so-called phase-space are determined by the values of n consecutive samples of the time series. For example, when $n = 3$, the time series is represented in three-dimensional space as a set of points such that each point $P = (x,y,z) = (S_i,S_{i+1},S_{i+2})$ where the x-axis represents the IBI in milliseconds of the sample S_i of the time series $T(S_1,S_2,S_3, \ldots S_i, \ldots S_N)$ and the y- and z- axes, respectively, represent the magnitude of the IBI of the phase advanced (by 1 and 2, respectively) elements of the time series. A graphical example of this procedure for $n = 2$ of IBIs derived from human subject data is shown in Fig. 2.

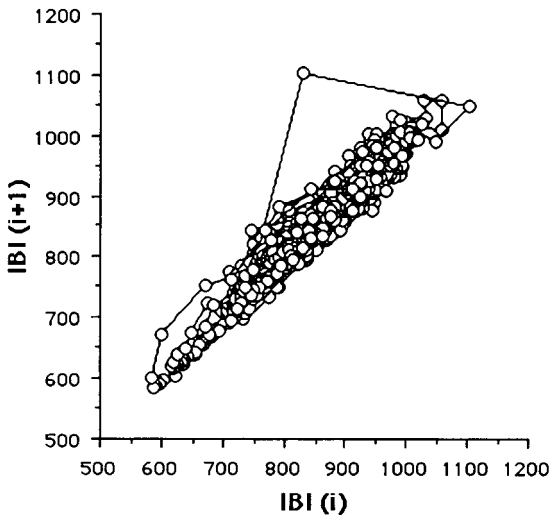

Fig. 2. Phase-space plot of IBI (1000 data points) from a healthy human subject obtained during mild exertion (walking). The coordinates on the x-axis represents the original time-series in milliseconds. The coordinates on the y-axis represent the phase-advanced value of the time series, also in milliseconds. Consecutive points are connected by a straight line.

3.2. Rescaling

Rescaling is accomplished by computing the sum of the lengths of all vector segments in n-space that connect each succeeding point on the n-dimensional curve. For any embedding dimension, we can determine the Euclidean length L of the curve for any sampling resolution r as $L(r)$. When $r = 1$, $L(r)$ is calculated by the summation of the length of the curve joining every point in phase-space giving L_{max}, when $r = 2$, $L(r)$ is calculated by summing the length of the curve joining every other point in phase space. This process is continued for increasing values of r to yield a relationship between $L(r)$ and r. For succeeding computations with a different r a new curve is thus generated, connecting points along the original curve, but sampled at a lower resolution.

In general, the relationship between the length $L(r)$ and the sampling resolution is given by:

$$L(r) = \sum_{i=0}^{I}\left[\sum_{k=0}^{n-1}(S_{ir+k} - S_{ir+r+k})^2\right]^{1/2} \tag{1}$$

with $r = 1,2,3,\ldots M$ and I is the integer part of $(N - r - n + 1)/r$. N is the number of data points of the IBI time series vector $T(S_1,S_2,S_3,\ldots S_i,\ldots S_N)$ and M is the maximum number of scales. The choice of M is dictated by the length of the data set. Empirically, we have found that the relationship $\log[L(r)/r]$ vs $\log(r)$ remains linear when $M < \sqrt{N}$. If M is set too high, there will be some deviation from linearity at higher values of r in the relationship, which will bias the results. The calculation of $L(r)$ for a three dimensional phase-plot is thus:

$$L(r) = \sum_{i=0}^{I}\sqrt{[s(ir) - S(ir + r)]^2 + [S(ir + 1) - S(ir + r + 1)]^2 + [S(ir + 2) - S(ir + r + 2)]^2} \tag{2}$$

3.3. Choice of embedding dimension

Since each point in the phase-space contains interval information with respect to adjacent

Table 1
Dependence of H on embedding dimension

Embedding dimension	$H \pm$ S.E.M.		
	expected $H = 0.00$	expected $H = 0.50$	human subject
6	0.00 ± 0.01	0.49 ± 0.01	0.29 ± 0.03
5	0.00 ± 0.01	0.48 ± 0.01	0.29 ± 0.03
4	0.01 ± 0.01	0.46 ± 0.01	0.29 ± 0.03
3	0.00 ± 0.01	0.46 ± 0.01	0.29 ± 0.03
2	0.01 ± 0.02	0.45 ± 0.01	0.30 ± 0.03
1	0.02 ± 0.03	0.43 ± 0.02	0.31 ± 0.03

elements of the time series, the choice of the embedding dimension depends on the dynamics of the time series. The best choice of the embedding dimension would correspond to the largest number of elements which encompass the "memory" of the time series. For a true fractional Brownian process, this is known to be infinite [3]. However, for human IBI time series, the memory of the process appears to degrade rapidly. The appropriate embedding dimension can be experimentally determined by testing for convergence of the slope of the log–log plot of $L(r)/r$ vs r, as will be described below. Choosing higher embedding dimensions than is required by the dynamics of the time series exponentially increases computing time, while not significantly altering the final value of D and therefore H. As shown in Table 1, an optimal embedding dimension for these particular data from human subjects was found to be $n = 3$.

3.4. Estimation of H

The length of a self-similar fractal curve is given by the following relationship [3]:

$$L(r) = (L_{\max})r^{(1-D)} \quad r > 0 \quad 0 < D < 2 \tag{3}$$

where the length $L(r)$ is measured as the Euclidean distance between sampled values of the time series, r is the sampling resolution or 'yardstick size' and D is the similarity dimension. When $r = 1$, L is computed using every sample value of the time series and reaches a maximum length $L_{\max}$. For $r = 2$, every second sample is used in the length computation and $L < L_{\max}$ because 'peaks' and 'valleys' of width < 2 do not contribute to the total length. With increasing values of r, progressively wider local maxima and local minima are disregarded and, if the curve is self-similar, the total length decreases according to Eq. (3). The log transformation of Eq. (3) yields:

$$\log[L(r)/r] = \log[L_{\max}] - D[\log(r)], \quad r > 0 \tag{4}$$

which has the form of a linear equation, with intercept $L_{\max}$ and slope $-D$.

Hence, if the curve exhibits self-similar scaling behavior, a plot of $\log[L(r)/r]$ vs $\log(r)$ will result in a straight line, with $-D$ as the slope. The similarity dimension D of the curve is related to the Hurst exponent by $H = 2 - D$ [4–6].

We can use this process to determine if the curve associated with the interbeat interval time series is self-similar. If the curve is self-similar, then we can derive the fractal dimension D and the Hurst exponent of the curve corresponding to the original time series.

4. Results

The method was applied to electrocardiographic data from 9 healthy subjects. Data sets for subjects 1 through 4 were obtained from Y. Yamamoto[1], data for subject 5 was obtained from this laboratory and data sets from subjects 6, 7 and 8 were obtained from W. Poplawska[2].

[1] Y. Yamamoto, ftp://psas.p.u-tokyo.ac.jp/

[2] J.J. Zebrowski, http://www.mpipks-dresden.mpg.de/~ntptsa/Data/Zebrowski-D/

Subject 4 was engaged in moderately heavy exercise, while subject 5 was examined during very heavy exercise (5a) and at rest (5b). All data were analyzed using software incorporating the algorithm, developed on Microsoft (R) Windows NT (TM), using Microsoft (R) Visual C++ 6.0 and Microsoft (R) Foundation Classes and Templates (MFC&T) version 6.0.

Benchmark tests were run on a Toshiba Satellite Pro 490XCDT with an Intel Family 6 266 MHz processor and 64 MB RAM, running under Windows 95 (TM) Workstation 4.0 SP3. For benchmarking purposes, a series of 50,000 data points was used. The dimensional embedding was set to 6 and the maximum scale M was set to $\sqrt{N}$, where N is the window size. The interval between windows was set to 1. The benchmark measures the time the program took to compute H for all windows across the time series. The software calculated H 49,900 times for a window size of 100 in 59 s. For a window size of 1000, H was calculated 49,000 times in 9 min and 41 s.

Results for the calculation of H derived from the relationship $H = 2 - D$, are shown in Table 2. A plot of $\log[L(r)/r]$ vs. $\log(r)$ is linear, as shown for subject 6 in Fig. 3. Analysis of data sets from other subjects resulted in qualitatively similar linear plots.

The values of H obtained in resting subjects 1, 2, 3, 5a, 7 and 8 are consistent with values of H obtained by coarse-graining spectral analysis (CGSA) [7]. The higher value of H seen in subject 5b is consistent with increased H values obtained during exercise [7].

4.1. Tests using synthetic time series with defined parameters

For these tests, all the synthetic IBI time series were computed using a spectral synthesis method [1] from the Time-Series Statistical Analysis System, TSAS 3.01.01b (TSAS) (7) compiled and run on a DEC 3000/600S AXP. This module conveniently allows varying the length, mean, standard deviation and spectral characteristics of the resulting time series, using a spectral synthesis method. The spectral exponent β is used to characterize noise having a nonflat Fourier spectrum, which is inversely proportional to the frequency $1/f^{\beta}$. Since for these characteristic noises, a log–log plot of the power spectral density vs. the frequency results in a slope of $-\beta$, the spectral exponent is sometimes also called the spectral slope.

Table 2
Results of analysis of human subject EKGs

Subject ID	$H \pm$ S.E.	M	R^2	IBI mean $\pm$ S.D. (ms)
1	0.07 ± 0.02	31	0.998	1034 ± 91
2	0.14 ± 0.02	31	0.997	984 ± 94
3	0.08 ± 0.02	31	0.998	1092 ± 107
4 (moderate exercise)	0.16 ± 0.02	41	0.995	578 ± 44
5a (rest)	0.09 ± 0.02	31	0.996	1219 ± 73
5b (heavy exercise)	0.47 ± 0.03	20	0.995	289 ± 106
6	0.164 ± 0.005	100	0.999	858 ± 191
7	0.148 ± 0.005	100	0.999	741 ± 87
8	0.156 ± 0.003	100	0.999	913 ± 129

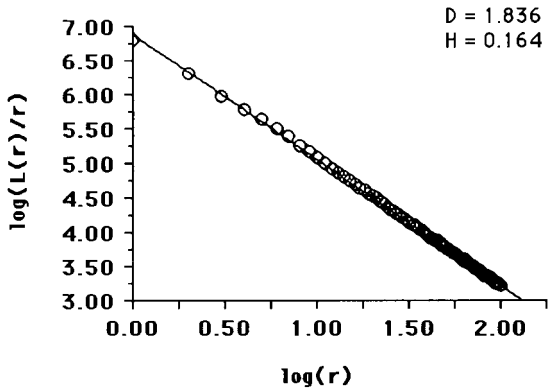

Fig. 3. Plot of log[L(r)/r] vs. log(r), calculated from a set of 98355 IBI data points obtained from subject ID 6, with embeddingdimension = 3 and $M = 100$. The linear slope suggests that a power-law relationship is present, supporting the assumption of self-similarity of the IBI time-series for the range of scales investigated.

To test if the method is capable of characterizing the behavior of an IBI series with defined spectral characteristics, we used computer-generated IBI time series data with mean IBI = 1000 ms, S.D. = 100 and N = datalength = 1000. The spectral slopes in the synthesis procedures were set to 0.0; 0.5; 1.0; 1.5; 2.0; 2.5; 3.0. D was obtained from Eq. (3) with $L(r)$ estimated from Eq. (1) with an embedding dimension of 6 and $M = 31$. As shown in Fig. 4, as the spectral slopes are incremented from 0 to 3, H varies from 0 to 1 as follows: $H = 0.00$;

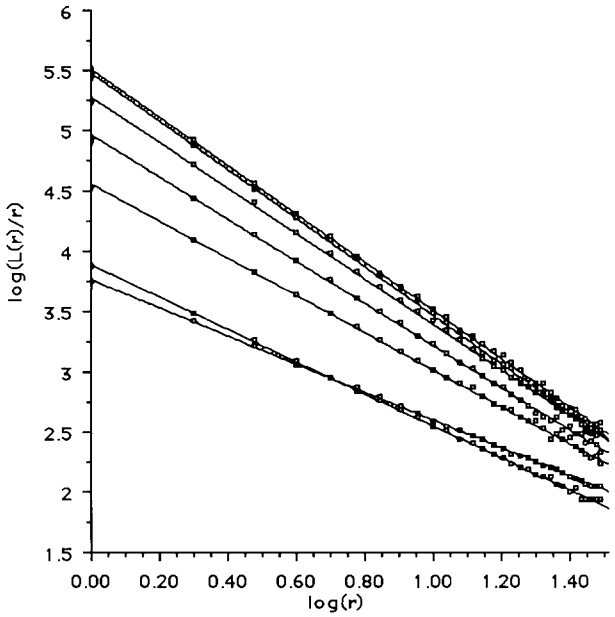

Fig. 4. Computation of D for seven synthesized IBI time-series of 1000 data points with the same mean (1000) and standard deviation (100) but differing spectral slopes. As the spectral exponent is varied from 0 to 3 in increments of 0.5, the slope −D, decreases from −2 to about −1, where D is the fractal dimension of the curve. $H = 2 - D$.

0.04; 0.10; 0.26; 0.46; 0.67; 0.83. The R^2 for all the regression lines was 1.00. The divergence from the expected $H = 1$ with the spectral slope $\beta = 3$ is discussed below.

4.2. Effects of statistical moments of the time series on H

We tested if varying the statistical moments of the time series while holding the spectral noise characteristics steady would affect the calculation of H. We generated the appropriate time series of datalength = 1000 and varied the statistical moments while maintaining the spectral slope at 0 or 2. As shown in Table 3, D was calculated using $M = 31$ and $n = 3$. The results show good agreement with theoretical values of $H = 0$ for a spectralslope = 0 and $H =$ 0.5 for a spectralslope = -2, suggesting that the method is not sensitive to differences in the mean and standard deviation of the time series. Note that for $H = 0.5$ there was a bias, in that an underestimation of the true value occurred with an embedding dimension of $n = 3$. The reason for this is discussed in the next section.

4.3. Effects of embedding dimension on the calculation of H

The results of altering the embedding dimension on the calculation of H are shown in Table 1. The relationship of the estimated H and the embedding dimension for a time series of 1000 IBI data points synthesized with a spectral slope of 0 (expected $H = 0$) and a spectral slope of 2, (expected $H = 0.5$) appear in the first and second columns. Data from a human subject is shown in the third column all calculated with $N = 1000$ and $M = 31$ with mean IBI $\pm$ S.D. in milliseconds was 864.29 $\pm$ 76.39.

We note from Table 1 that with the synthetic data, a higher embedding dimension is required to achieve less bias in the estimation of H. This is expected, because the synthetic data has a longer 'memory' and is thus closer to an ideal fractal. However, for the human subject data presented, an embedding dimension of 3 was sufficient for convergence.

When the embedding dimension is >1, we are examining the scaling behavior of the time series with respect to phase, as well as time. Given these results, it appears that the influence of the magnitude of previous IBIs on a particular IBI is exerted within 2 to 3 heart beats, at least within the data lengths examined in the present study of up to 1197 data points. An embedding dimension of 3 was used for calculating H in all the human subject examples, as all human subject electrocardiographic data analyzed resulted in H values that converged with $n = 3$.

Table 3
Effect of signal characteristics on computed H-values

Mean IBI (ms)	S.D.	CV	$H = 0$	$H = 0.5$
1034	91	0.09	0.00	0.47
834	73	0.09	0.00	0.47
634	56	0.088	0.00	0.45
1034	182	0.18	0.00	0.42
1034	365	0.35	0.01	0.45

4.4. Data length and calculation of H

The effects of varying data length on determination of H was tested by synthesizing a set of time series letting H vary from 0 to 1 in increments of 0.25, with mean of 1000 and S.D. 200. For each value of H and N, 5 time series and the mean and S.D. of the estimated value of H are reported. The embedding dimension was $n = 6$. To avoid problems with bias inherent in the spectral synthesis method, only the initial N data points of the generated time series were used for the analysis, the total length of the original time series being 10 N. The results are shown in Fig. 5.

5. Discussion

For short series with $N \leq 1000$, the estimate of H obtained using the current method is more accurate than estimates obtained with other regression-based methods, including the power spectral density, discrete wavelet transform and dispersive analysis methodologies [8,9]. With $N \geq 500$ points, there is reasonable convergence of the parameter H to values determined by the generating algorithm, except for a bias towards lower values with $H < 0.5$. At this time, it is unclear whether this represents a bias inherent in the methodology for the generation the time series, or whether the method is less robust when $H > 0.5$. Nevertheless, since the method was developed for the study of IBI time series and since such time series appear to result in H values < 0.5, the methodology is adequate for the analysis of electrocardiographic data.

Differences in the mean and S.D. of the signal do not significantly alter the estimated value of H, nor do variations in data length. These strengths allow the estimation of H from relatively short data sets, such as might be obtained from 3–5 min of electrocardiographic recordings. This allows investigators to compare and contrast data sets from different subjects relatively unaffected by minor differences in data collection times.

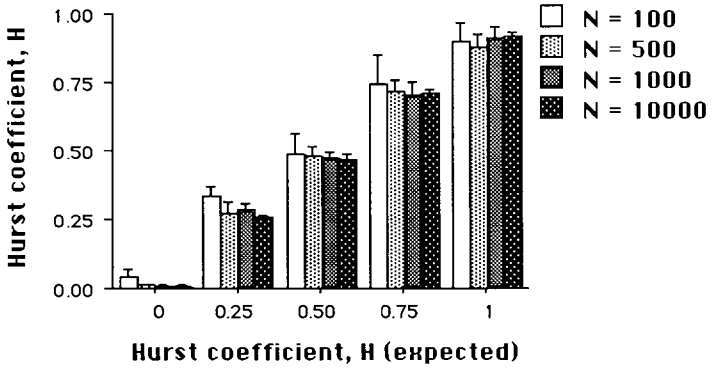

Fig. 5. Accuracy for the estimation of H based on varying data length. The values represented on the y-axis show the calculated values, compared to values shown on the x-axis which were used to generate the time-series. The bar plots show the mean values of H each for $n = 5$ time-series, the error bars representing the standard deviation.

An important implication of this method relates to the problem of measuring the continually changing underlying dynamics of many biologically generated time series. Since any estimation of H must of necessity be a time-average value, the ability to discriminate variations in the dynamics of time series within resolutions of 100–1000 data points is an advantage in the study of biological systems. A graphic example of the changing internal dynamics of human electrocardiographic interbeat interval time series is seen in Fig. 6. Fig. 6 shows the results of opening and closing the eyes, as well as walking, on the value of H calculated for a window size of 1000 data points. This illustrates the sensitivity of the methodology. There is no difference in heart rate between the eyes open and eyes closed conditions. However, a large difference in the noise characteristics of the time series appears as shown by the decrease in H associated with the eyes closed condition.

The characterization of biological signals such as cardiac IBI time series can potentially be improved by assessing the internal underlying dynamics of such time series at relatively high resolutions. Studying changes of the dynamic properties might provide some insight into the generation of deterministic chaos resulting from the interplay of oscillatory homeostatic feedback loops, as they appear to operate in the generation of interbeat-intervals [10].

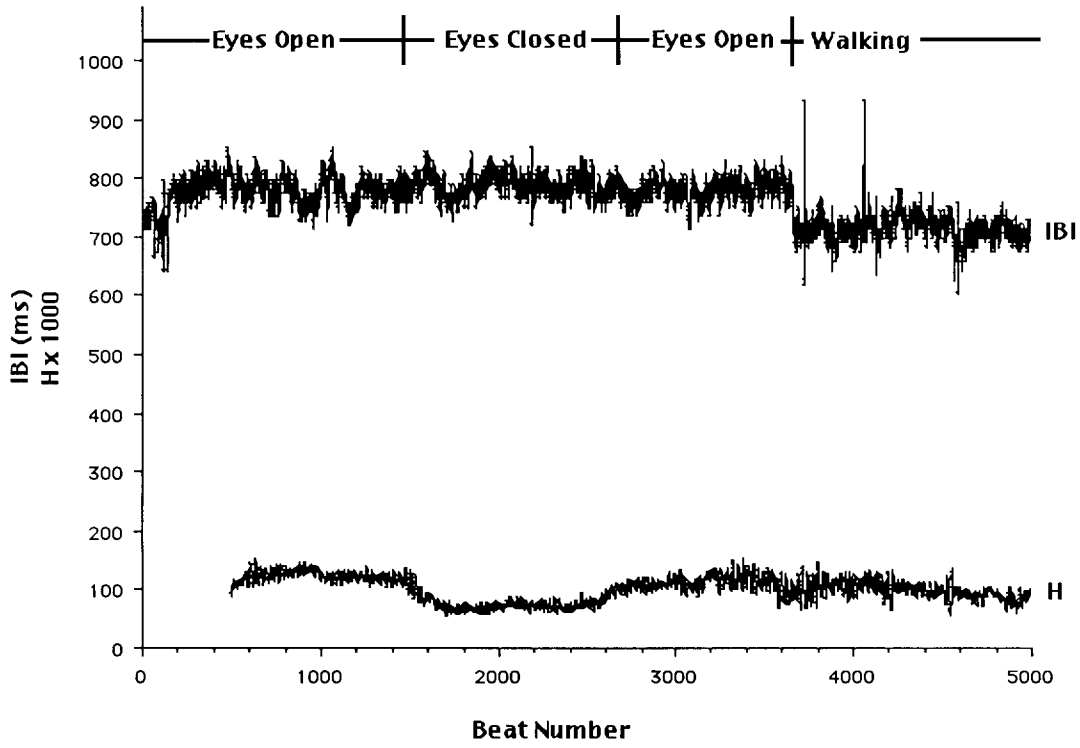

Fig. 6. Interbeat interval (IBI) data from a human subject, in milliseconds, is plotted concurrently with H under eyes open and eyes closed conditions and while ambulating. H values presented are multiplied by 1000 so that they will scale with IBI values. H has been calculated using a sliding window of 1000 data points, with $M = 31$ and $n = 3$.

One of the drawbacks of the methodology is the need for the operator to choose a proper embedding dimension for the time series. In part, the choice of embedding dimension is dependent on the complexity of the underlying dynamics, which may change with time. The choice is straightforward for relatively short time series, since the embedding constant can be incrementally changed and H calculated after each increment. At some point, further increases in embedding dimension do not appreciably change the value of H, as shown in Table 1. For longer time series, it is possible that epochs with very complex internal dynamics might co-exist and be contiguous to epochs with low complexity. H could thus be overestimated or underestimated if an 'average' embedding dimension were chosen based on the results of the whole time series. While overestimating the embedding dimension does not appear to alter the value of H for some biological time series with finite memory, it is unknown whether more complex biologically generated time series would behave in the same fashion.

In the case of synthetic time series with an 'infinite' memory, the choice of low value for the embedding dimension may result in a bias in the estimation of H. As shown in Table 1, we observed that for synthetic time series with H close to 0.50, unless the embedding dimension was set to at least 6, H tended to be underestimated.

Measures of cardiac complexity are known to be altered in various disease states. Autonomic neuropathy and cardiomyopathy are associated with altered electrocardiographic signal dynamics [11]. More importantly, alterations in signal complexity have also been found to be predictive of ventricular tachycardia after myocardial infarction and in patients with coronary artery disease [12–14].

In summary, this method for estimating the Hurst coefficient of a fractal time series is shown to be relatively accurate and unbiased when applied to a physiological time-series, especially for short data lengths and when $H < 0.5$. Inasmuch as the structure of the algorithm allows the streaming of input and output data, the development of research tools for real-time monitoring of any time series, including biologically generated time series such as the IBI, is facilitated.

6. Summary

The analysis of the signal dynamics of electrocardiographic interbeat interval (IBI) time series is made difficult by the fractal nature of these time series, as the variance of the mean IBI increases with increasing numbers of data points. An alternative method is presented which extracts the Hurst exponent of the time series, H. This method avoids problems with nonstationarity which result from measurements in the time domain or frequency domain. The method consists of embedding the time series in phase space and analyzing the resulting function as a self-similar fractal curve. The fractal dimension of this curve is found to be D and is related to the Hurst exponent by $H = 2 - D$. Software utilizing the algorithm was developed on Microsoft (R) Windows 95 (TM), using Microsoft (R) Visual C++ 6.0 and Microsoft (R) Foundation Classes and Templates (MFC&T) version 6.0. The method was checked for reliability using spectrally synthesized time series and was shown to be insensitive to data length, as well as nonstationarity of the input time series. The method was also confirmed using IBI time series data derived from human subjects. The resulting H values were

in agreement with values obtained using alternate methodologies and also show that human IBI time series are fractal in nature. A windowing technique was incorporated in the software, which allows the calculation of H for windows of arbitrary data length. When applied to human IBI time series data, this illustrated the changing nature of the internal dynamics of these time series.

Acknowledgements

The authors are indebted to the late Dr. V. Markku Linnoila for his support and encouragement of this work, Dr. Yoshiharu Yamamoto of Tokyo University, Tokyo, Japan and Dr. Wanda Poplawska of the National Institute of Cardiology, Warszawa, Poland, for kindly donating their electrocardiographic data for analysis and Professor Jan J. Zebrowski of the Warsaw University of Technology, Poland, for assistance in accession of data. Information on obtaining software used in performing these analyses can be found at ftp://helix.nih.gov/pbdp/ The software is distributed free of charge for non-commercial biomedical research purposes.

References

[1] Y. Yamamoto, R.L. Hughson, On the fractal nature of heart rate variability in humans:effects of data length and β-adrenergic blockade, American Journal of Physiology 266 (1994) R40–R49.
[2] P.K. Stein, M.S. Bosner, R.E. Kleiger, B.M. Conger, Heart rate variability: a measure of cardiac autonomic tone, American Heart Journal 127 (1994) 1376–1381.
[3] B.B. Mandelbrot, J.W. Van Ness, Fractional Brownian motions, fractional noises and applications, SIAM Review 10 (1968) 422–437.
[4] J. Feder, Fractals, Plenum Press, New York, 1988.
[5] R.F. Voss, Random fractals: characterization and measurement, in: R. Pynn, A. Skjelthorp (Eds.), Scaling Phenomena in Disordered Systems, Plenum Press, New York, 1985.
[6] M.F. Barnsley, R.L. Devaney, B.B. Mandelbrot, H.-O. Peitgen, D. Saupe, R.F. Voss, The Science of Fractal Images, Springer-Verlag, New York, 1988.
[7] Y. Yamamoto, Y. Hoshikawa, M. Miyashita, Effects of acute exposure to stimulated altitude on heart rate variability during exercise, Journal of Applied Physiology 81 (1996) 1223–1229.
[8] J.B. Bassingthwaighte, G.M. Raymond, Evaluation of the dispersional analysis method for fractal time series, Annals of Biomedical Engineering 23 (1995) 491–505.
[9] R. Fischer, M.A. Akay, comparison of analytical methods for the study of fractional brownian motion, Annals of Biomedical Engineering 24 (1996) 537–543.
[10] C.K. Peng, S.V. Buldyrev, J.M. Hausdorff, S. Havlin, J.E. Mietus, M. Simons, H.E. Stanley, A.L. Goldberger, Non-equilibrium dynamics as an indispensable characteristic of a healthy biological system, Integrative Physiological and Behavioral Science 29 (1994) 283–293.
[11] G.C. Butler, S. Ando, J.S. Floras, Fractal component of variability of heart rate and systolic pressure in congestive heart failure, Clinical Science 92 (1993) 543–550.
[12] T.H. Makikallio, T. Seppanen, K.E.J. Airaksinen, J. Koistinen, M.P. Tulppo, C.K. Peng, A.L. Goldberger, H.V. Huikuri, Dynamic analysis of heart rate may predict subsequent ventricular tachycardia after myocardial infarction, American Journal of Cardiology 80 (1997) 779–782.
[13] M.W. Rich, S.S. Jasbir, R.E. Kleiger, R.M. Carney, A. teVelde, K.E. Freeland, Correlation of heart rate variability with clinical and angiographic variables and late mortality after coronary angiography, American Journal of Cardiology 62 (1988) 714–717.

[14] T.H. Makikallio, S. Hoiber, L. Kober, C. Torp-Pedersen, C.K. Peng, A.L. Goldberger, H.V. Huikuri, Fractal analysis of heart rate dynamics as a predictor of mortality in patients with depressed left ventricular function after acute myocardial infarction. TRACE Investigators. TRAndolapril Cardiac Evaluation, American Journal of Cardiology 83 (1999) 36–839.

Paolo B. DePetrillo received an MD degree from Brown University, Providence, RI, USA, in 1981 and is board certified in Internal Medicine and Clinical Pharmacology. He is currently a Senior Investigator the National Institutes of Health. His research interests include the use of cardiac signal dynamics as pharmacodynamic surrogate markers of end-organ drug toxicity and the development and application of nonlinear dynamic techniques to biomedical signal analysis.

d'Armond Speers was born in 1968. He graduated summa cum laude from Western Carolina University, North Carolina, USA, in 1990 with a degree in cognitive psychology and received an M.S. degree in theoretical linguistics from Georgetown University in 1994. While pursuing his doctorate in computational linguistics at Georgetown University, he also worked as a senior researcher at Thomson Labs, designing and prototyping proof-of-concept solutions to problems in information technology. He currently works for Requisite Technology, Inc. as a research linguist. His research interests include computational linguistics and natural language processing. His other interests include the Klingon language.

Urs E. Ruttiman received his B.S.E.E. from the Technikum Burgdorf (HTL) in Switzerland in 1961 and his M.S.E.E from the Swiss Federal Institute of Technology (ETH) in Zurich, Switzerland. He received a Ph.D. in Biomedical Engineering from the University of Pennsylvania in 1972. From 1991 until his death in 1998, he was a staff scientist in the Section of Brain Electrophysiology and Imaging at the National Institute on Alcohol Abuse and Alcoholism, National Institutes of Health, Bethesda, MD, USA. His research interests included use of wavelet techniques in the analysis of PET and fMRI images, as well as development of novel statistical techniques for biomedical signal analysis.