

Long-duration correlation and attractor topology of the heartbeat rate
differ for healthy patients and those with heart failure

Robert G. Turcott and Malvin C. Teich

Department of Electrical Engineering
Columbia University, New York, NY 10027

ABSTRACT

The point process formed by the sequence of human heartbeats exhibits $1/f$ -type fluctuations arising from long-duration power-law correlation. We obtain the normalized coincidence rate $g^{(2)}(\tau)$ of the underlying point process and demonstrate that the correlation is stronger for patients with normal hearts than those with heart failure. This is consistent with the greater rate fluctuations observed in the normal heart. A number of statistical measures are used to establish the existence and reveal the form of the correlation, including rescaled range analysis (R/S), pulse-number distribution (PND), Fano-factor time curve (FFC), and power spectral density (PSD). The normalized coincidence rate is obtained from the FFC. The long-duration, power-law correlation observed in the sequence of heartbeats is similar to that observed at a number of neurophysiological loci in a variety of species. We also obtain the box-counting estimate of the attractor's fractal dimension from a phase-space reconstruction and analysis of the trajectory of the *number* of heartbeat events. This approach reveals that the heartbeats of normal patients exhibit an attractor of higher dimension than those of heart-failure patients. All of these measures have also been applied to normal and heart-failure data for which the interbeat (R-R) intervals have been randomly shuffled. The results are substantially different from those of the unshuffled data, demonstrating that long-duration correlation is not a subtle property of the heartbeat. The various properties we have investigated suggest several conveniently calculated, quantitative indices that indicate to which group a particular data set belongs. Of the 27 data sets examined, our indices correctly classified 25.

1. INTRODUCTION

Scale-invariant fluctuations and power-law correlations have been demonstrated in the time sequence of heartbeats.^{3,7,8,9,24,27,33} Different heartbeat power spectra are associated with various pathological conditions, such as diabetic autonomic neuropathy,^{5,21} uncomplicated essential hypertension,¹³ sudden infant death syndrome,¹⁷ potential for sudden cardiac death,^{7,8,9} severe heart disease,²⁷ and myocardial infarction.³ The studies that have been conducted focus, for the most part, on the relative power in various spectral peaks as an indicator of health.

In this paper, we examine normal and heart-failure data in the context of new analyses designed to quantify power-law fluctuations. We also present evidence for the presence of chaos on long time scales. Both of our approaches yield different results for normal and heart-failure data, and, as such, may provide indices that are useful for characterizing the degree of health of the heart.

It has previously been observed that long-duration, power-law correlation is present in the sequence of action potentials generated by primary auditory (VIII-nerve) neurons in the cat,^{36,37} chinchilla,²⁹ and chicken;³⁰ in lateral superior olivary³⁹ and mesencephalic reticular formation¹¹ neurons in the cat; and in the descending contralateral movement detector of the locust.⁴⁰ At all these loci the correlation extends to long time scales, with the upper limit of observed correlation times imposed only by the duration of the recording.

Fractal fluctuations, which give rise to $1/f$ -type power spectra, are readily apparent in estimates of the instantaneous heart rate.^{8,9,27,33} In Fig. 1, an instantaneous rate estimate is formed by counting the number of contractions in successive counting periods of duration T . Intuitively, one expects that the fluctuations exhibited in the estimate will decrease as T is increased, as is indeed the case for the nonfractal signal shown in Fig. 1(a). Fractal signals, on the other hand, have the property of slow convergence of fluctuations: as the counting time is increased, the fluctuations decrease slowly, as a power law function of the counting time, or persist at the same magnitude, as is apparent in Fig. 1(b). For a true fractal signal, one would not be able to infer the counting time from the appearance of a sample rate function. The data exhibiting the self-similar rate fluctuations presented in Fig. 1(b) are taken from a long-duration recording of a healthy heart. The nonfractal signal in Fig. 1(a) has the same interval density function as the data presented in Fig. 1(b), but each interval is independent of the others so that there is no long-term correlation. The statistical measures we present have been developed to quantify this phenomenon.

The demonstration of chaotic behavior in the heart has been limited to short time scales^{1,31,16} or special experimental conditions.^{6,12,42} We present evidence for a low-dimensional attractor in the time series of heart *rate* data recorded over long time scales. In particular, we find that the box-counting dimension obtained from the reconstructed phase-space trajectory ranges between 1.6 and 3.1 for the data examined, with normal hearts exhibiting larger dimensions.

Previous studies^{18,34,27} have, for the most part, focused on analyses based on the sequence of *intervals* between heartbeats, in which the abscissa of the time series is the interval number. Our analysis is based on rate estimates constructed from the sequence of heartbeats, which preserve time as the abscissa. This approach has the advantage of allowing a direct interpretation of the observed correlation.

In Section 2, we present the theoretical background for the statistical measures we utilize. The results are provided in Section 3, a discussion of the difference between normal patients and those with heart failure is presented in Section 4, and the implications are considered in Section 5.

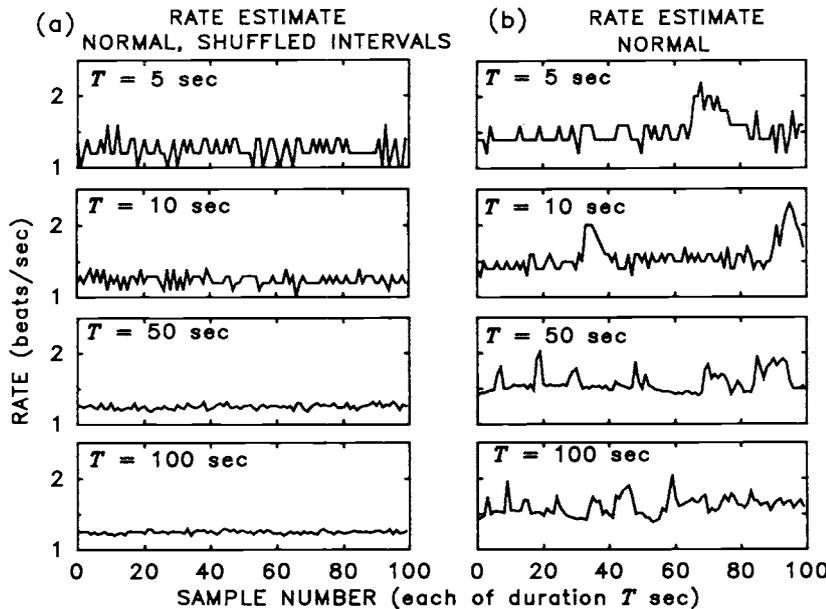


Fig. 1. Instantaneous rate estimates formed by counting the number of beats in successive counting periods. (a) For usual (nonfractal) signals, the estimate converges and the fluctuations (standard deviation of the estimate) decrease with increasing counting time T . (b) The fluctuations in fractal signals converge with a slower power-law dependence on counting time, if they converge at all. Note that for each of the estimates in (b), fluctuations on the order of 0.5 beat/sec persist, even though the counting time is increased by a factor of 20 from the top panel to the bottom. The convergence properties are quantified by the Fano-factor time curve (FFC) presented in later figures. The data in (b) is from a long-duration recording of a normal heart (data set 16265). The data in (a) is from the same recording, but with the intervals randomly reordered, which maintains the same relative frequency of each interval but destroys all long-term correlations. These same two data sets are used in the following figures, and are denoted 'normal', and 'normal, shuffled intervals', respectively.

2. THEORETICAL BACKGROUND AND STATISTICAL MEASURES

We treat the sequence of heartbeats mathematically as an unmarked point process, in which the occurrence time of each contraction (R-phase) is taken to contain the relevant information in the sequence of beats. In this approach the focus is on the long-term behavior of the sequence, allowing the complex waveform of an electrocardiogram-recorded heartbeat (QRS-complex) to be replaced with a single number (the time of occurrence) thereby greatly reducing the computational demand. The occurrence of a contraction at time t_i is therefore represented by an impulse $\delta(t)$ at that time, so that the entire realization is given by

$$s(t) = \sum_i \delta(t - t_i). \quad (1)$$

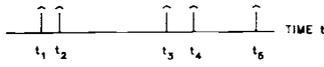
In specifying a particular realization of a point process, then, one need only specify the set of occurrence times $\{t_i\}$ of the events, as illustrated in Fig. 2(a). To gain insight into the mechanism underlying the generation of the point process, we wish to characterize it from the set of occurrence times $\{t_i\}$ obtained from a single recording.

The amplitude distribution and correlation function of a stationary random process provide a partial characterization of its general behavior. For a point process, the amplitude distribution is represented by the rate of occurrence of impulses, which is readily obtained from the data, but the normalized coincidence rate $g^{(2)}(\tau)$ (equivalent to the correlation function of a continuous stochastic process^{36,38}), is generally difficult to estimate. Because of this, alternative representations of the underlying point process are used.

One perspective, which is often used in long-duration heartbeat analysis, considers the sequence of intervals τ_i between adjacent impulses, as shown in Fig. 2(b), rather than the presence or absence of an impulse at each time. In this way the point process $s(t)$ (a sequence of impulses distributed on a continuous time axis) is represented by a discrete-time random process $\{\tau_i\}$ (a sequence of positive real-valued random variables).

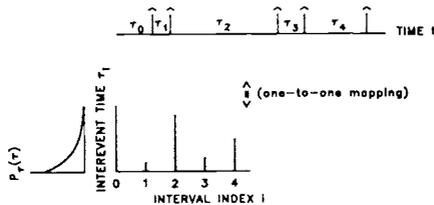
POINT PROCESS

(a)



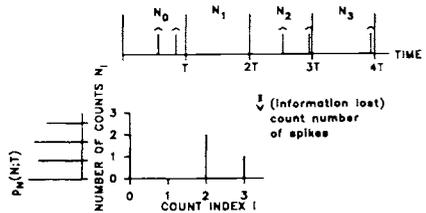
SEQUENCE OF INTEREVENT TIMES $\{\tau_i\}$

(b)



SEQUENCE OF COUNTS $\{N_i\}$

(c)



SEQUENCE OF GENERALIZED COUNTS $\{X_i\}$

(d)

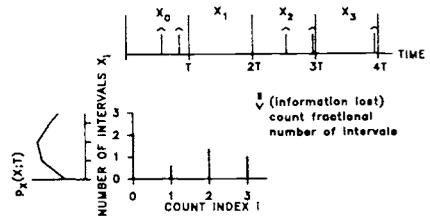


Fig. 2. Analysis schemes for a point process. (a) The sequence of heartbeats is represented by a sequence of idealized impulses, forming a stochastic point process. The complex underlying waveform is thus reduced to a sequence of impulse occurrence times $\{t_i\}$. For convenience of analysis, several alternative representations of the point process are used. (b) A sequence of interevent times $\{\tau_i\}$ is formed from the time between impulses, resulting in a discrete-time, positive real-valued stochastic process. All information contained in the original point process is preserved in this representation, but the discrete-time axis of the sequence of interevent times is randomly distorted relative to the real time axis of point process. (c) The sequence of counts $\{N_i\}$, a discrete time, non-negative integer-valued stochastic process, is formed from the point process by recording the number of events in successive counting windows of length T . (d) Higher resolution in the rate estimate can be obtained by counting the fractional number of intervals. The sequence of generalized counts $\{X_i\}$, a discrete-time, positive real-valued stochastic process, is formed from the point process by recording the fractional number of intervals in successive counting windows of length T . For both the sequence of counts (c) and the sequence of generalized counts (d), the time axis of the point process t is simply related to the discrete time axis of the sequence of counts by $t = Ti$, so that the correlation properties of the point process can be inferred directly from the correlation in the sequence of counts. In both these processes, information is lost in mapping the point process to the sequence. For the sequence of counts (c), the amount lost can be made arbitrarily small by reducing the size of the counting window. The sequence of generalized counts (d) does not have this property.

We employ three statistical measures to characterize the random process comprised of the sequence of interevent intervals $\{\tau_i\}$: the interspike interval (ISI) histogram, rescaled range analysis (R/S), and the power spectral density (PSD) of the sequence of interevent times.

The ISI histogram shows the relative frequency of occurrence of interspike intervals as a function of interval size; it is an estimate of the probability density function of $\{\tau_i\}$. It is, perhaps, the most commonly used of all statistical measures. The ISI histogram readily provides information about the underlying process over short time scales, but all information regarding the order of interval occurrence is lost. Indeed, a reordered sequence of the same intervals will yield an identical ISI histogram since reordering does not affect the relative frequency of occurrence. If no correlation exists among the interspike intervals then the process formed by the sequence of interevent times $\{\tau_i\}$ is renewal, and the point process is completely characterized by its ISI histogram. Conversely, if the process is nonrenewal (i.e., correlation exists among its interevent intervals) then the ISI histogram does not completely characterize the process. In this case, measures that reveal the nature of the correlation provide complementary information to the ISI histogram.

The rescaled range analysis (R/S) provides information about correlation among interspike intervals. It is particularly well suited to processes that exhibit long-term correlation or that have large or infinite variance^{4,14,15,23,37} but has previously not been used in heartbeat analysis. For a block of k intervals, the difference between each interval and the average interval size is obtained and successively added to a cumulative sum. The normalized range $R(k)$ is the difference between the maximum and minimum values that the cumulative sum attains, divided by the standard deviation of the interval size. $R(k)$ is plotted against k . Information about the nature and degree of correlation in the process is obtained by fitting $R(k)$ by the function k^h . If $h > 0.5$ then positive correlation exists among the intervals, while $h < 0.5$ indicates the presence of negative correlation, so that an interval that is larger than the mean tends, on the average, to be preceded by or followed by one smaller than the mean.

Fourier transform methods provide another approach to quantifying the correlation of a stochastic process. The periodogram estimate $S_x(f)$ of the power spectral density (PSD) of the sequence of intervals $\{\tau_i\}$ is given by

$$S(f) \equiv \frac{1}{M} |X(f)|^2, \tag{2}$$

where M is the number of samples and $X_c(f)$ is the discrete Fourier transform of the sequence of intervals.^{26,28} Information can be inferred about the correlation among intervals from the PSD since it is just the transform of the autocorrelation function (ACF) of the interval process.

Since the $\{\tau_i\}$ provide a complete representation of the point process all information can, in principle, be obtained from it. However, the mapping of the "real", continuous time axis of the underlying point process to the discrete time index i of the interevent-time process $\{\tau_i\}$ introduces an intrinsic random distortion: intervals that are close in absolute time may be substantially separated in the sequence $\{\tau_i\}$ (there may be many short intervening intervals) and intervals close in the sequence $\{\tau_i\}$ may be temporally distant in the point process (there may be one long intervening interval). This distortion is manifested by the presence of artifacts in the frequency spectrum of the process.² Because of this some questions are more readily addressed using an approach that estimates the rate of the point process, retaining time as the abscissa.³⁵

Our approach is shown in Fig. 2(c), where the time axis is divided into equally spaced, contiguous counting windows of length T , and the number of impulses in the i^{th} window is counted and denoted N_i . The sequence $\{N_i\}$ forms a discrete-time random process of nonnegative integers. In general, information is lost in forming this process, although for regular point processes the amount lost can be made arbitrarily small by reducing the size of the counting window T . The advantage of this representation lies in the link, provided by the counting time T , between the discrete time axis of the counting process $\{N_i\}$ and the absolute, "real" time axis of the underlying point process. With the process of counts $\{N_i\}$, N_i and N_{i+k} refer to the number of counts in windows separated by precisely $T(k-1)$ seconds, so that correlation in the process $\{N_i\}$ is readily associated with correlation in the underlying point process $s(t)$.

We now turn our attention to three measures of the random process formed by the sequence of counts, or equivalently the rate.

In the same way that the ISI histogram provides an estimate of the probability density function of the amplitude of the interspike interval process $\{\tau_i\}$, the pulse-number distribution (PND) provides an estimate of the probability mass function of the amplitude of the sequence of counts $\{N_i\}$. While the ISI histogram reveals the character of the underlying point process on short time scales (on the order of the average interevent time), the PND reveals the character of the point process on time scales that are of the order of the counting time T . Since the counting time is not intrinsic to the process, but rather is specified externally, the character of the process at arbitrary time scales can be examined. The Fano factor, $F(T)$, defined as the count variance divided by the mean, is a useful index of dispersion in this case since its value for the homogeneous Poisson point process (HPP), which plays the role of the Gaussian process in continuous stochastic process theory, is unity for all T (see Appendix A). A Fano factor greater than unity indicates that the underlying point process is more clustered than the HPP at the given time scale,^{36,24} whereas a Fano factor less than unity indicates that the underlying process is more regular than the HPP, as with a dead-time-modified Poisson point process (DTMP, see Appendix A).

In addition to the Fano factor, the shape of the PND provides information about the underlying process. A higher probability of an even number of counts than an odd number in an interval of length T , for example, implies that events tend to occur in pairs, separated by less than T seconds. In this way the PND at a particular counting time and the Fano factor provide information about the correlation of the point process at that time scale.³⁵

The Fano-factor time curve (FFC), which shows the functional dependence of the Fano factor on the counting time T , provides additional information about the nature of the correlation. For example, an FFC with a fractional power-law dependence on T results from a process with power-law correlation.^{36,38} The relationship between the FFC and the normalized coincidence rate $g^{(2)}(\tau)$ is presented in Appendix B.

Finally, we use phase-space reconstruction and a box-counting algorithm²⁰ to show that the fractal fluctuations might be accounted for by a low-dimensional attractor. This approach exploits the fact that the topological properties of the attractor of a dynamical system can be determined from time series of a single observable.^{10,19,41} Briefly, an m -dimensional vector $\vec{X}_m(t_n) \equiv [x(t_n), x(t_n + l), \dots, x(t_n + (m-1)l)]$ is formed from a discrete-time series $x(t_n)$. The parameter m is the embedding dimension, and l is the lag (usually taken to be the location of the first zero crossing in the autocorrelation function of the time series). As the time t_n elapses, the vector $\vec{X}_m(t_n)$ traces out a trajectory in the m -dimensional embedding space. The box-counting algorithm estimates the (fractal) dimension of the trajectory. For uncorrelated noise, the box counting dimension d_b increases with increasing m . For deterministic systems, in contrast, d_b becomes constant as m becomes larger than the number of degrees of freedom in the system.

The sequence of counts provides an estimate of the instantaneous rate of the point process (after division by the counting time), but it is an estimate that is constrained to take on discrete values. A generalization of this counting estimate is provided by counting the fraction of intervals that fall in each counting window, as illustrated in Fig. 2(d). The use of this measure was initially motivated by an integrate and fire model, where a rate function is integrated until a fixed threshold is reached, triggering an event and resetting the integrator.² For statistical measures such as the FFC and the PSD, the two counting estimates are equivalent for large T . However, analysis of the reconstructed phase space depends on the fine structure of the trajectory, and an integer count estimate (the sequence of counts) would map all vectors of the generalized count estimate to nearby integer lattice points. For

example, for $m = 1$ and T less than the refractory period, the trajectory given by the sequence of counts yields only two points in the embedding space: zero and one. The generalized count estimate provides a continuum of real numbers in the reconstructed trajectory. It is, perhaps, less useful than the usual count in other measures, however, in that it does not permit direct comparison with the results for a homogeneous Poisson point process.

The point-process nature of the sequence of counts for short counting times results in a noise floor in the PSD, apparent at high frequencies.²⁶ For this reason we prefer the sequence of generalized counts for spectral analysis. As with the sequence of intervals, the periodogram estimate $S_x(f)$ of the power spectral density of the sequence of generalized counts $\{x_i\}$ is formed. The quantity $S_x(f)$ is again given by Eq. (2), where M is still the number of samples, but now $X_x(f)$ is the discrete Fourier transform of the sequence of generalized counts. With this PSD estimate, information can be inferred about correlation in real time, as opposed to the interval-based PSD in which time is distorted.

An instantaneous rate estimate can be formed from either the sequence of counts or the sequence of generalized counts by dividing the random variable by the counting time T ; as a result, both of these measures are equivalent to rate estimates. To facilitate the interpretation of results, we carry this out for the sequence of generalized counts $\{X_i\}$, forming a sequence of rate estimates $\{x_i\}$, before performing the box-counting and spectral analyses. However, neither the box-counting dimension d_b nor the form of the PSD are affected by this.

3. RESULTS

Twelve data sets recorded from healthy patients and 15 sets recorded from patients with heart failure were analyzed. Three of the patients with heart failure also had atrial fibrillation. The data were supplied by Ary Goldberger and David Rigney of Beth Israel Hospital, Boston. The recordings were made with a Del Mar Avionics 445 Holter monitor, digitized at 250 Hz, and the beat-to-beat (R-R) intervals were measured automatically with a computer program (Aristotle).²⁵ Information about the data sets is summarized in Table 1. The heart-failure patients were all medicated with the drug Milrinone.

In this section we focus on one recording from a healthy heart (data set 16265) and one recording taken from a patient with heart failure (data set 6796). These two examples are typical of their respective classes of data. We also apply these same measures to the point process formed by shuffling (randomly reordering) the intervals of these data sets. Such random reordering destroys correlation among the intervals, while preserving the interval probability density function. It is equivalent to generating a renewal point process with an identical interevent probability density function as the original, allowing the effects arising from long-term (rate) correlation, and effects arising from the form of the interevent probability density function, to be readily distinguished. Data set 16265 and its shuffled version are the two point processes used in generating Fig. 1 above.

The ISI histograms of normal data (solid curve) and heart-failure data (dotted curve) are presented in Fig. 3. The narrower width of the histogram for the heart-failure patient is consistent with the lower interval variance for data set 6796 [$\text{var}(\tau) = 0.0081 \text{ sec}^2$], as shown in Table 1. The healthy patient generates an excess of both shorter and longer intervals relative to the patient with heart failure [for data set 16265, $\text{var}(\tau) = 0.0291 \text{ sec}^2$, as shown in Table 1]. Since random reordering of the intervals does not alter the relative frequency with which they occur, the ISI histograms of the shuffled data sets are identical to those of the originals.

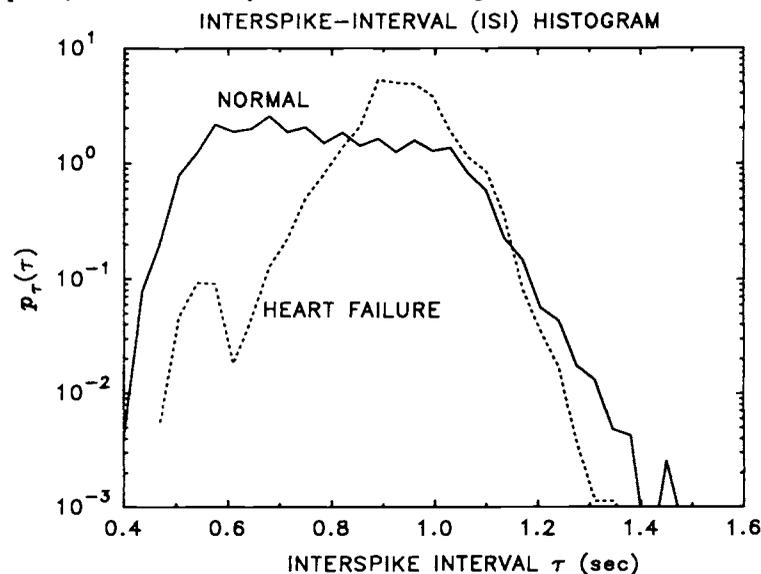


Fig. 3. The interspike interval (ISI) histogram for data from a healthy patient (solid curve) and a patient with heart failure (dotted curve). The variance of the ISI histogram of the heart-failure data is substantially less than that of the normal data.

File Number	Number of intervals	Duration (sec)	Mean Rate (sec ⁻¹)	$\langle \tau \rangle$ (sec)	var(τ) (sec ²)	FFC exponent α	Count PSD exponent β	d_B
NORMAL PATIENTS								
16265	100460	80061.9	1.255	0.7970	0.0291	.953	0.93	2.66
16272	93177	84395.5	1.104	0.9058	0.0202	.872	1.21	2.64
16273	89846	74348.6	1.208	0.8275	0.0212	.872	1.20	2.75
16420	102081	77761.0	1.313	0.7618	0.0102	.897	1.01	2.41
16483	104338	76099.5	1.371	0.7294	0.0079	.941	0.87	2.58
16539	108331	84669.3	1.279	0.7816	0.0225	.873	1.07	2.81
16773	82160	78141.1	1.051	0.9511	0.0600	.959	1.00	2.84
16786	101630	84051.4	1.209	0.8270	0.0134	.925	1.11	3.04
16795	87061	74734.7	1.165	0.8584	0.0448	.943	1.21	2.45
17052	87548	76399.6	1.146	0.8727	0.0251	.843	1.19	2.98
17453	100674	74482.0	1.352	0.7398	0.0106	.921	0.86	2.94
c4	88140	71398.7	1.234	0.8101	0.0172	.872	1.30	2.84
HEART-FAILURE PATIENTS								
6796	75821	71940.9	1.054	0.9488	0.0081	.854	1.40	1.70
7257	118376	71166.4	1.663	0.6012	0.0013	.878	0.93	2.44
8519	80878	71941.4	1.124	0.8895	0.0040	.779	1.48	1.81
8552	111826	71827.0	1.557	0.6423	0.0039	.773	1.76	1.99
8679	119094	71180.1	1.673	0.5977	0.0026	.944	1.35	2.37
8988	118058	71140.3	1.660	0.6026	0.0081	.940	1.28	2.44
9049	92497	71964.8	1.285	0.7780	0.0033	.958	1.35	2.44
9377	90644	71965.0	1.260	0.7939	0.0033	.868	1.24	2.13
9435	114959	71196.7	1.615	0.6193	0.0008	.944	1.42	1.60
9643	148111	72015.6	2.057	0.4862	0.0002	.917	1.44	2.03
9674	115542	71976.3	1.605	0.6229	0.0071	.922	1.57	2.29
9706	115064	71320.0	1.613	0.6198	0.0103	.955	1.21	2.53
9723	115597	71999.9	1.606	0.6229	0.0003	.909	1.75	2.19
9778	93607	71946.3	1.301	0.7686	0.0051	.958	1.69	2.17
9837	115205	71947.0	1.601	0.6245	0.0043	.947	1.62	2.37

Table 1
Characteristics of the data sets investigated.

The rescaled range analysis for each data set is presented in Fig. 4. The circles correspond to the function $k^{1/2}$ for comparison with the data. Positive correlation among interspike intervals is present since $R(k)$ grows more rapidly than $k^{1/2}$. This measure indicates that data from healthy hearts (upper solid curve) has slightly stronger correlation for small numbers of interspike intervals, and slightly weaker correlation for intermediate numbers than unhealthy hearts (upper dotted curve). On the whole, though, it reveals no substantial difference between data from healthy patients and those with heart failure. The processes of shuffled intervals (lower two curves) exhibit a dependence very close to $k^{1/2}$, as expected for renewal point processes.^{4,14,15,23,37}

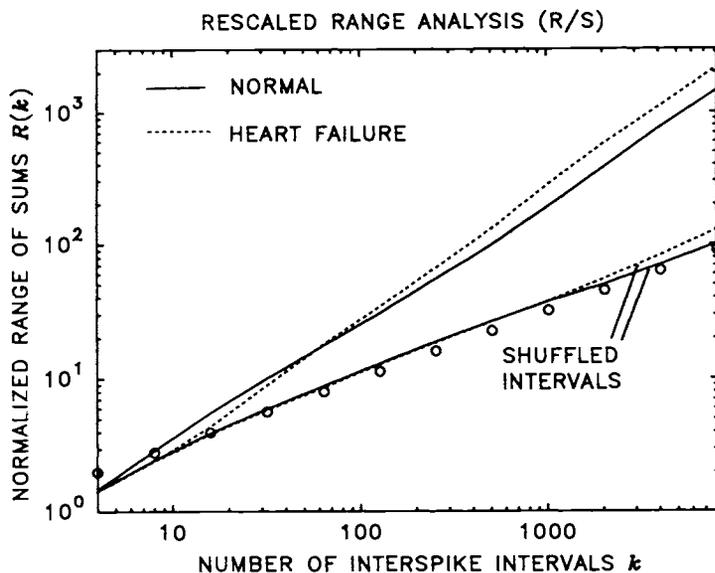


Fig. 4. Rescaled range analysis (R/S) applied to the data from a healthy patient (upper solid) and a patient with heart failure (upper dotted). Positive correlation is present for both data sets. Small collections of intervals for the healthy heart show slightly stronger correlation than those with heart failure, but generally the R/S curves of the two groups of data were essentially indistinguishable. The curves for the shuffled intervals (lower two curves) fall along the curve $k^{1/2}$ (open circles), as will any sequence of independent random variables.

The averaged periodogram estimates of the PSD generated from the sequence of intervals $\{\tau_i\}$ are presented in Fig. 5(a). The power-law form of $S_r(f)$ implies the existence of power-law correlation among the interspike intervals. The correlation is somewhat greater for the normal data set (solid curve). The PSDs of the sequence of shuffled intervals [Fig. 5(b)] are flat, as expected, since the intervals are then independent (i.e., the sequence of shuffled intervals forms a white noise process). This is in contrast to the PSD of the point process, which is the Fourier transform of $g^{(2)}(\tau)$, and which is flat only for a homogeneous Poisson point process. The PSD of a point process cannot, in general, be inferred from the PSD of a sequence of intervals. The variance of the interval process, which is the integral of the interval-based PSD, is smaller for the heart-failure data, in accord with its narrower ISI histogram (see Fig. 3).

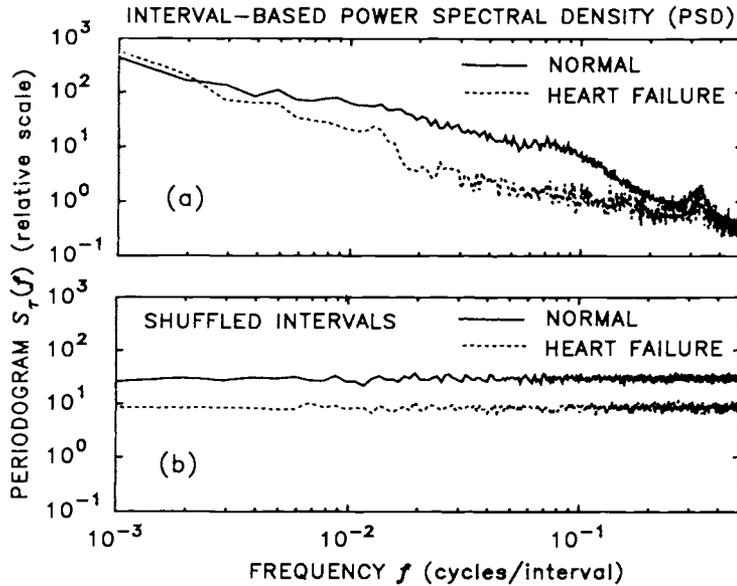


Fig. 5. (a) The averaged periodogram estimate $S_r(f)$ of the power spectral density (PSD) for the sequence of interspike intervals. The divergence of the PSD from a horizontal line indicates that both healthy and heart-failure data have correlated intervals. The correlation for the normal data (solid) is somewhat greater than that for the heart-failure data (dotted). In obtaining the PSD, the data set was partitioned into successive blocks of 1024 intervals. The periodogram of each block was obtained, and the blocks then were averaged. (b) The PSDs of the sequences obtained from the shuffled data are flat, as expected for sequences of independent, identically distributed random variables. The PSD of heart failure data shows less power at all frequencies than does the normal data, which results from the smaller interval variance of the former.

The PND generated with a counting time of $T = 10.0$ sec is shown in Fig. 6. The greater count variance (dispersion in the PND) for the normal data (solid) indicates that there is greater variability in its counts than the heart-failure data (dotted) on this time scale, which in turn implies that the rate exhibits greater variability. This variability cannot be entirely accounted for by the interval variance, since the PNDs of the shuffled data sets show reduced count variance, indicating that the count variance of the original process arises in part from long-term correlation.

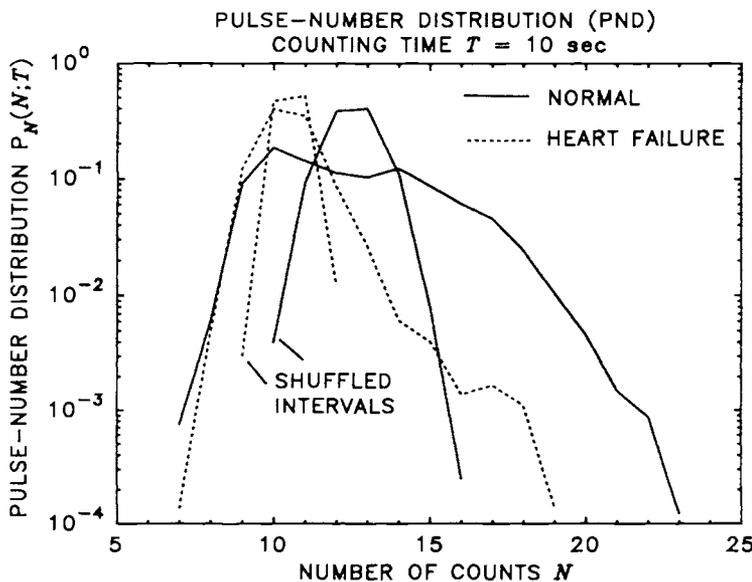


Fig. 6. The pulse-number distribution (PND) for normal (upper solid) and heart-failure (upper dotted) data using a counting time of $T = 10.0$ sec. The data from healthy hearts shows a wider distribution than the heart-failure case, indicating greater clustering, or fluctuations, in the underlying point process. The clustering is a manifestation of correlation. Shuffling the intervals eliminates this correlation, while leaving the ISI histogram unaffected. The substantially narrower PNDs for the shuffled data demonstrate the presence of long-duration correlation.

The FFCs are shown in Fig. 7 (top two curves). In the limit of small T , the sequence of counts becomes a sequence of Bernoulli random variables because of the refractoriness of the heartbeat. With p the probability of observing an event in one of the counting intervals, which goes to zero as $T \rightarrow 0$, we have $\lim_{T \rightarrow 0} F(T) = \lim_{T \rightarrow 0} \frac{p(1-p)}{p} = 1$. Regions where $F(T) < 1$ correspond to time scales over which the sequence heartbeats is more regular than the homogeneous Poisson point process. The FFC dips strongly below unity on time scales where the effect of refractoriness is significant. The ringing in the FFC, particularly apparent in the heart-failure data, is due to the clock-like regularity of the heartbeats on this time scale. The FFC dips to a very low value when the counting time T is such that the number of beats per counting interval is almost constant. A slight increase in T admits an extra count more frequently, resulting in a greater variance-to-mean ratio. The increase of the Fano factor above unity for counting times greater than ≈ 10 sec, as a fractional power-law function $\propto T^\alpha$, indicates that the underlying sequence of heartbeats has greater count variability, and is hence more clustered, than a Poisson process on these time scales. The parameter α , given in Table 1, is constrained to lie between 0 and 1.^{22,36} The greater count variability arises from positive correlation that, for T greater than ≈ 10 sec, overcomes the negative correlation imposed by refractoriness. Since shuffling destroys this long-term, positive correlation, the FFCs for the shuffled processes never rise to unity (bottom two curves).

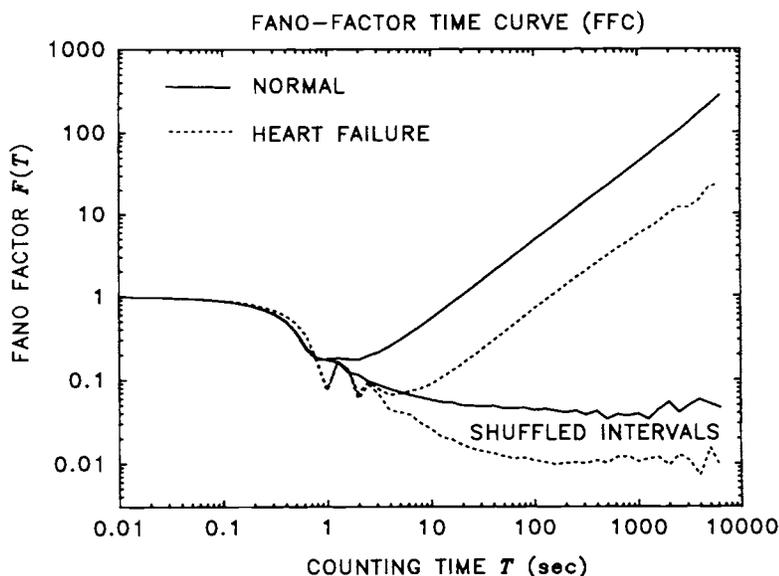


Fig. 7. The Fano-factor time curve (FFC) for normal data (solid) and heart-failure data (dotted). A departure of the FFC from unity indicates the presence of correlation in the sequence of counts. Regions where $F(T) < 1$ correspond to time scales where negative correlation dominates. This negative correlation is imposed in part by the refractoriness of the heart. Positive correlation, reflecting rate fluctuations, causes the FFC to assume values in excess of unity. The FFC appears to follow fractional power law behavior for large values of counting time T , i.e., $F(T) \propto T^\alpha$, implying that the underlying point process has power-law correlation over the corresponding time scales. The larger magnitude of $F(T)$ for the normal data indicates that the healthy heart exhibits stronger correlation, and greater fluctuations, than the sick heart. The difference in the FFCs may provide one way to assess the health of the heart. The FFCs of the renewal processes (lower curves) remain below unity, since no long-duration correlation is present.

The dramatic difference between the FFCs of healthy and sick hearts shown in Fig. 7 is typical of the data sets examined. All the normal data sets exhibit the same relatively strong correlation shown in this Figure. The fact that the FFC for the heart-failure data does not exceed unity until an order of magnitude later in time than the FFC for the healthy heart data indicates that some mechanism, aside from simple refractoriness, is imposing regularity on the point process in the heart-failure case.

In Fig. 8(a) we present the averaged periodograms $S_x(f)$ generated from the rate estimate of the sequence of generalized counts $\{x_i\}$ using a counting time of $T = 10.0$ sec. The divergence of $S_x(f)$ from a horizontal line indicates the presence of correlation among the sequence of generalized counts. In particular, $S_x(f)$ is power law, obeying the form $Af^{-\beta}$ over a range of frequencies. This is consistent with the presence of power-law correlation indicated by the FFC, and reflects the self-similar fluctuations shown in Fig. 1(b). These data therefore represent an example of $1/f$ -type noise (the medium-dashed line is inversely proportional to frequency).

The PSDs $S_x(f)$ of the sequences of generalized counts formed after shuffling the intervals [Fig. 8(b)] are flat, indicating that the counts are approximately uncorrelated. The sequence of counts at this counting time T , then, is a white noise process. The PSDs $S_r(f)$ of the sequences of shuffled intervals (Fig. 5) were also flat, as they will be for any renewal process. The PSDs of the sequence of counts is, in general, only flat for T much greater than the average interevent time. For smaller values of T , the PSD will reflect the correlation in the point process arising from the form of the interval probability density function, even though the intervals have been shuffled. This is because in general (for all point processes other than the Poisson), an interval extending beyond the counting window affects the numbers of counts in succeeding intervals. A white sequence of intervals does not imply that the point process is white. This illustrates one of the dangers of assuming that spectral information obtained from the sequence of intervals is equivalent to the spectral content of the point process. It is not: indeed $S_r(f)$ in Fig. 5 has cycles/interval as its abscissa whereas $S_x(f)$ in Fig. 8 has cycles/sec as its abscissa. Note that both PSDs of the heart failure data have less power at all frequencies than the corresponding PSDs of the normal data. This reflects the lower count variances of the heart-failure data, as is clear from the PNDs shown in Fig. 6.

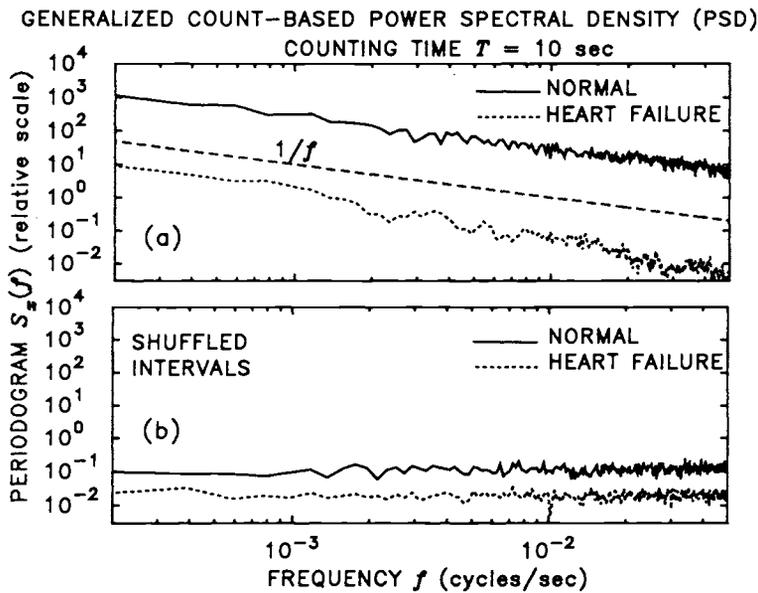


Fig. 8. (a) The averaged periodogram estimate $S_x(f)$ of the power spectral density (PSD) for the sequence of generalized counts. Periodograms with 1024 points were formed from successive segments of data using a counting time $T = 10.0$ sec. The periodograms of each data set were then averaged. The PSD of the data shows power-law behavior, varying as $S_x(f) = Af^{-\beta}$, so that with this counting time the sequence of counts is a $1/f$ -type noise process. The function f^{-1} is included for comparison (medium-dashed curve). (b) For the generalized counts of the shuffled intervals, the PSDs are flat, indicating that the counts are uncorrelated, as is the usual case for renewal processes with T much larger than the average interevent time. This is not true, however, for small T , even though the process is renewal. This distinction with the PSD for the sequence of intervals illustrates one of the dangers of attempting to infer correlation properties of the underlying point process from correlation properties of the sequence of intervals.

In Fig. 9 we present the results of a phase space reconstruction and the estimation of box-counting dimension d_B . The value of d_B increases with the embedding dimension m until m exceeds the number of degrees of freedom of the system. The behavior shown in Fig. 9 is typical of the other data sets, with curves from normal data reaching an asymptote at a higher value of d_B (in this case ≈ 2.6) than curves from the heart failure patients (in this case ≈ 1.7). The values of d_B for the shuffled intervals are higher and appear to continue to increase with m even at the highest values of m considered. Stochastic systems have an infinite number of degrees of freedom; consequently d_B increases indefinitely.

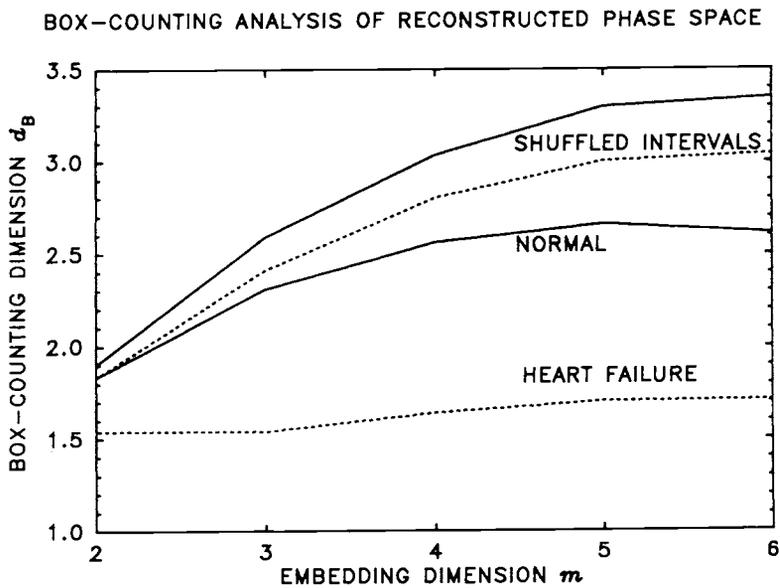


Fig. 9. The box-counting fractal-dimension estimate d_B of the attractor obtained from the sequence of generalized counts, shown as a function of the embedding dimension m . The data presented here are representative of their respective classes, with a larger value attained for the normal data (lower solid curve) than for the heart-failure data (lower dotted curve). For both normal and heart-failure data, d_B reaches its maximum at $m = 5$. The upper two curves show that d_B for the shuffled data increases more rapidly than for the original point process, and continues to rise even for the largest embedding dimensions available given the finite set of data.

The relationship between the FFC and the normalized coincidence rate $g^{(2)}(\tau)$ is presented in Appendix B. We can also forge a connection between the behavior of the FFC with increasing counting time and the box-counting dimension d_B . This is illustrated in Fig. 10, where the reconstructed trajectories for a two-dimensional ($m = 2$) embedding space are shown for a range of counting times. Data from a normal heart (data set 16265) was used in the lower panels, while its shuffled version (equivalent renewal point process) was used in generating the upper panels.

For the renewal point process, the phase space contracts dramatically with increasing counting time, while for the unshuffled data its spatial extent remains largely independent of counting time. This is a direct consequence of the convergence properties of the rate estimates shown in Fig. 1, and quantified by the FFC in Fig. 7.

For a particular value of T and m , d_b provides a measure of the fractal dimension of the set of trajectory points. The Fano factor $F(T)$, on the other hand, provides a measure of the spatial extent of the trajectory in phase space. This spread is independent of the embedding dimension. To generate $F(T)$ from a reconstructed trajectory, as shown in Fig. 10, one would project all points of the trajectory onto a single axis, multiply each of those scalar values by T (since for $F(T)$ the count is the random variable, rather than the rate), and then calculate the variance-to-mean ratio. For a time series to have the property that its reconstructed phase space does not contract (or contracts slowly in power-law fashion) with increasing counting time, or equivalently that its FFC has a power-law dependence on the counting time, it must have long-duration power-law correlation. This correlation may well be manifested by evidence for a low-dimensional attractor in the reconstructed phase space. The lack of such an attractor would imply the lack of long-term correlation, and therefore convergence of the rate estimate with T . In short, then, the appearance of a low-dimensional attractor is a necessary condition for a power-law FFC and self-similar rate fluctuations.

PHASE-SPACE RECONSTRUCTION IN 2-DIMENSIONAL EMBEDDING SPACE

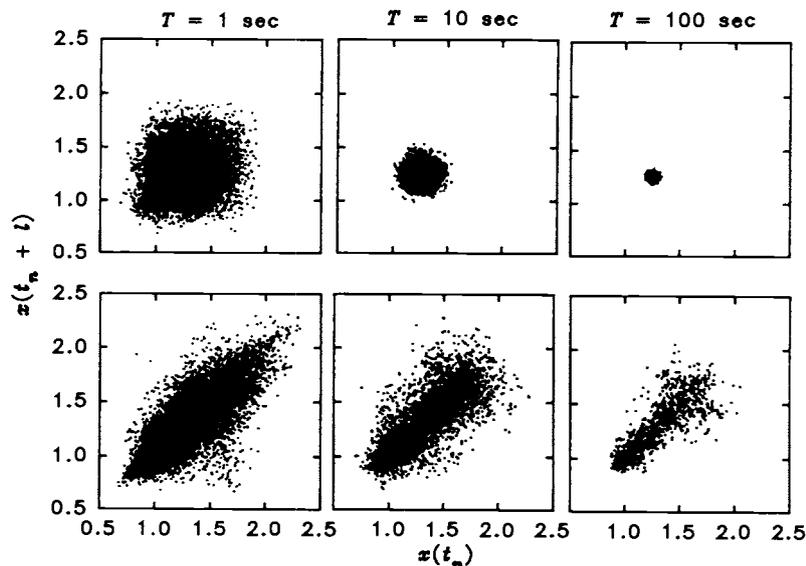


Fig. 10. An example of the reconstructed trajectory for a two-dimensional embedding phase space ($m = 2$), showing the relationship between the FFC and d_b . The box-counting algorithm operates on the set of points $\{\bar{X}_m(t_n)\}$, having nothing to do with the order in which they occur, so that the trajectory is a discrete dust in the phase space. For the nonfractal renewal process (normal data, shuffled intervals, shown in the upper panels), the trajectory contracts with increasing T , while for the normal unshuffled data, which exhibit fractal fluctuations (lower panels), the spatial extent of the trajectory is largely unaffected by increasing T . The FFC measures the rate of contraction with increasing T . The box counting dimension d_b measures the dimension of the dust in the phase space.

4. COMPARISON OF DATA FROM HEALTHY AND HEART-FAILURE PATIENTS

The variance of the intervals, $\text{var}(\tau)$, is shown graphically in Fig. 11(a). The variance of the intervals for data from heart-failure patients is generally less than that for healthy hearts, so that the reduced variability so apparent in heart-failure data on long-time scales is also seen in the interval statistics. However, the reduced interval variance alone is not sufficient to account for the differences in long-term variability shown in the FFC, the count-based PSD, or the box-counting dimension (Figs. 7, 8, 9, respectively). This is evidenced by the dramatically different results for the corresponding renewal point processes obtained by shuffling the intervals.

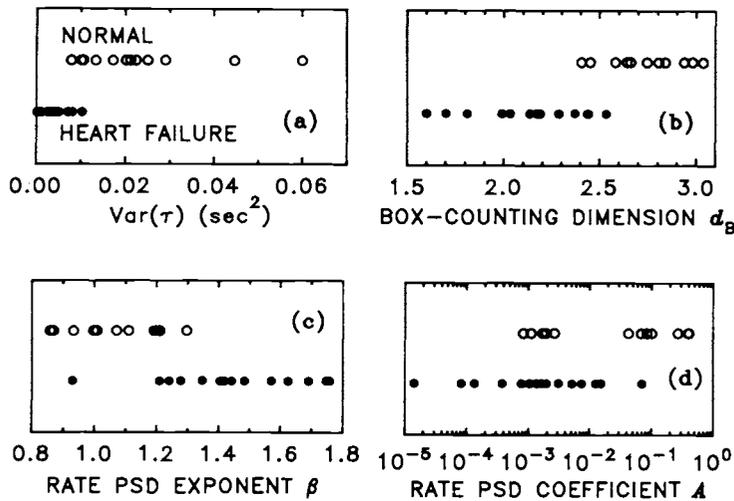


Fig. 11. (a) Graphical presentation of the variance of interspike intervals for data from healthy hearts (open circles) and those with heart failure (filled circles). The heart-failure data tend to exhibit lower variance. (b) The box-counting dimension d_B is typically larger for the normal data (open circles), with two normal sets falling among the heart-failure data (filled circles). The PSDs obtained from the sequence of generalized counts were fit by the power-law function $S_x(f) = Af^{-\beta}$. The normal data (open circles) generally have a lower value for the power law exponent β (c), while having a larger value for the coefficient A (d).

A power-law function $S_x(f) = Af^{-\beta}$ was fit to the PSD obtained from the sequence of generalized counts [Fig. 5(a)] over the region $0.0002 < f < 0.05\text{Hz}$ (cycles/sec) for normal data and $0.0002 < f < 0.01\text{Hz}$ for heart-failure data. The best fitting parameters, β and A , are shown in Figs. 11(c) and (d), respectively, for all the data sets. The PSDs for the healthy data generally exhibit a smaller magnitude of the power-law exponent than for the heart-failure data, and a greater coefficient A , but there is a significant degree of overlap between the two types of data.

The FFC, in contrast, provides a strong distinction between normal and heart-failure data. As shown in Fig. 12, aside from two data sets, the FFCs for healthy and sick hearts are separated, and this separation persists over three decades of counting time. Fitting the FFCs obtained from the data to Eq. (B3) in Appendix B allows us to determine both the underlying coincidence rate $g^{(2)}(\tau)$ and a quantitative index that can be used to predict with which of the two groups a particular data set is associated.

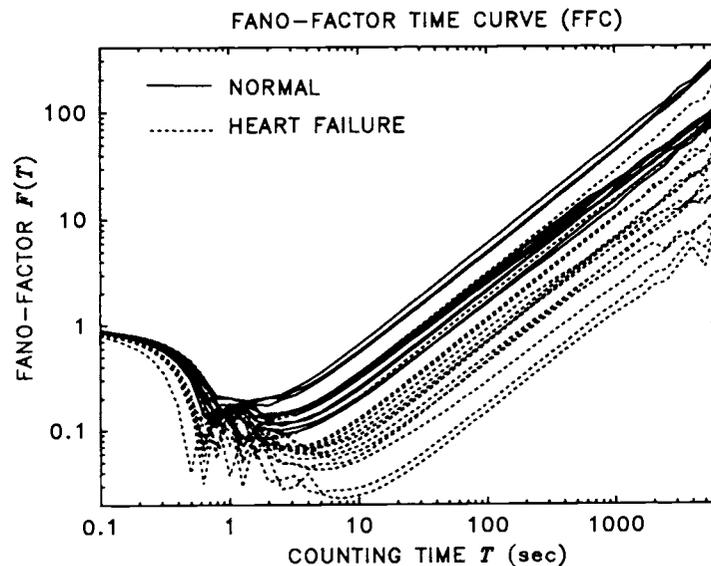


Fig. 12. The FFCs of all data sets presented in a single plot. Except for two curves for heart-failure patients, the two groups of data are well separated. The FFC, then, provides a possible measure for determining the state of health of the heart.

The parameters that provide the best fit Eq. (B3) to the data (μ , α , τ_d , δ , and τ_f) were obtained and are presented graphically in Figure 13. The values corresponding to data sets from healthy hearts are shown as open circles while those from sick hearts are shown as filled circles. None of these parameters, individually, separates the two classes of data, but since the shape of the FFC is implicit in them, some combination of them would provide a quantitative index that separates the two groups (save for the two heart failure sets that fall among the normal ones in Fig. 12).

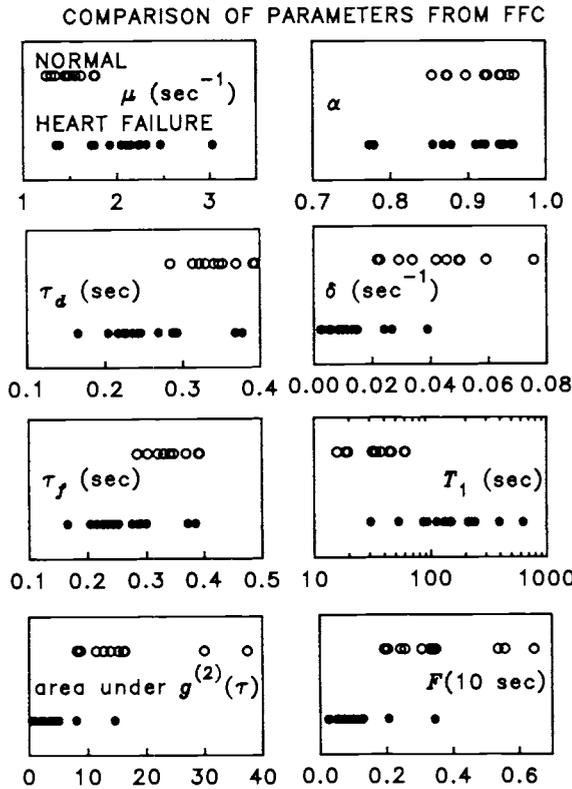


Fig. 13. Graphical presentation of the five parameters describing the FFC of a general stationary point process, obtained by fitting Eq. (B3) to the FFC of each data set (see Appendix B). In addition, T_1 , the counting time at which $F(T) = 1$, and $F(10 \text{ sec})$ are presented for all data sets. Also shown is the area under $g^{(2)}(\tau)$ for $1 \leq \tau \leq 1000 \text{ sec}$ (see Fig. 14).

The underlying normalized coincidence rates were obtained, using the best-fitting parameters provided by Eq. (B3), and are presented in Fig. 14. Aside from the two heart-failure data sets whose FFCs fell among those of healthy hearts, the two groups of data are well separated. The coincidence rates of healthy hearts show stronger correlation than those of sick hearts, indicating that the healthy hearts undergo greater fluctuations of the heart rate. This excess correlation can be quantified by calculating the integral of $g^{(2)}(\tau)$, taken between 1 and 1000 sec, which is shown in the lower left-hand panel of Fig. 13. Also shown in Fig. 13 are T_1 , the counting time at which the FFC exceeds unity, and $F(10)$, the value of the Fano factor at a counting time $T = 10 \text{ sec}$. Any one of these three indices [the integral of $g^{(2)}(\tau)$, T_1 , or $F(10)$] can be used with some confidence to predict whether the corresponding patient is healthy or might be suffering from heart failure. Of the 27 data sets, these three indices misclassify only two heart-failure data sets as healthy.

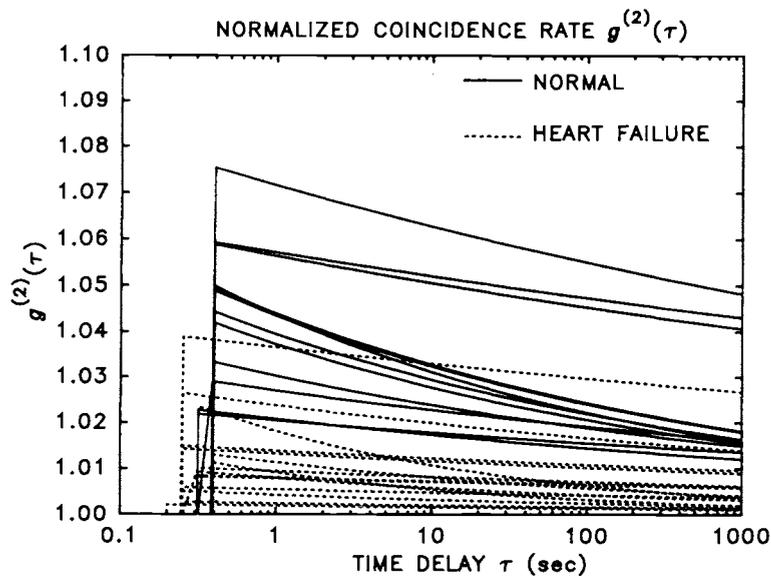


Fig. 14. The normalized coincidence rate $g^{(2)}(\tau)$ for all data sets. Aside from the two heart-failure data sets that fall among the normal data in the FFC (see Fig. 12) the two classes are well separated. The healthy hearts exhibit stronger correlation, which implies greater variability in the heart rate.

The indices used to classify the data exploit the greater correlation present in the underlying point process of normal hearts. This correlation is revealed in greater rate fluctuations. The two heart-failure data sets that were misclassified as normal (9674 and 9706) have relatively large interval variances (see Table 1). The misclassification of these data and their larger interval variances both suggest that these sets exhibit larger fluctuations in rate than the other heart-failure data sets.

The estimated box counting dimension d_B , as a function of the embedding dimension m is presented in Fig. 15. The estimates for all data sets appear to attain their maxima at $m = 5$ or $m = 6$, suggesting that sufficiently high dimensionality have been examined. The normal data exhibit a higher fractal dimension ($2.4 < d_B < 3.1$) than the heart-failure data ($1.6 < d_B < 2.5$). The values of d_B are presented graphically in Fig. 11(b). Because of this separation, d_B , like the FFC, may provide a useful index of the degree of health of the heart.

BOX-COUNTING ANALYSIS OF PHASE-SPACE RECONSTRUCTION

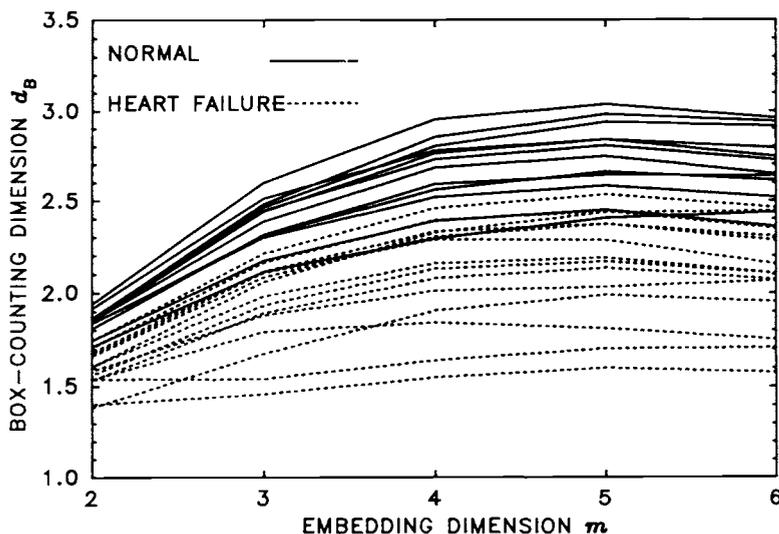


Fig. 15. Box-counting dimension d_B shown as a function of embedding dimension m for all data sets examined. The two classes of data are separated, except for two normal sets that fall among the heart-failure sets for all values of m . The two heart-failure sets that fell among the normal data in the FFC analysis here lie in the middle of the group of heart-failure data, suggesting that different aspects of the underlying point processes are causing the two classes to separate in the FFC and box-counting analyses. Like the FFC, the box-counting dimension d_B provides a possible measure for determining the state of health of the heart.

5. DISCUSSION

Using the statistical properties of point processes and dynamical systems analysis, we have investigated the long-duration correlation in normal and heart-failure patients. The correlation in the underlying point process revealed by count-based analysis is considerably stronger for normal data, and the associated attractors have higher dimension. Indices derived from these analyses may be useful as diagnostic tools for clinically distinguishing between the two groups of patients.

We examined three measures generated from the sequence of intervals obtained from the point process. The ISI histogram of the interbeat (R-R) intervals is generally narrower for heart-failure data. The rescaled range analysis indicates that there is strong correlation in the sequence of intervals, but does not reveal differences in the degree of correlation between the two classes of data. The interval-based PSD also shows strong correlation in the sequence of intervals. Using this measure, the normal data appears to have stronger correlation for intermediate frequencies.

All of the count-based measures appear to reveal differences between the two classes of data. The PND, taken at a sufficiently large counting time, indicates that there are greater rate fluctuations in the normal data than in the heart-failure data. The FFC shows power-law correlation in the sequence of counts for both classes of data, but with substantially stronger correlation for the normal group. The two groups (except for two heart-failure data sets that fell among the normal group) remain separated in the FFC plot for three decades of counting time. It appears that this property may be able to be exploited to conveniently distinguish between the two groups. The fractal dimension of the attractor, estimated via the box counting dimension d_b of the reconstructed phase space, is also generally larger for the normal data. The count-based PSD reveals the same kind of power-law correlation and difference in degree of correlation between the normal and heart-failure groups seen in the FFC, though the region of overlap between the two groups is greater than for the FFC or for d_b .

A number of parameters characterizing the FFC appear to successfully predict the health of the heart. The area under $g^{(2)}(\tau)$ (given by a combination of parameters fundamental to the FFC), T_1 [the counting time at which $F(T) = 1$], and $F(T_0)$ (for T_0 sufficiently large) all properly separate the data (25 of 27 data sets) into healthy and heart-failure classes. Similarly, d_b separates all but two of the data sets into two classes (the two sets that are misclassified by d_b are different from the two misclassified by the FFC).

The FFC is related to the correlation function $g^{(2)}(\tau)$ of the point process through an integral transformation, so that correlation over a range of times contributes to the value of the Fano factor at a particular counting time. Because of this, the FFC is a robust statistical measure of correlation in the underlying process. It therefore also provides a convenient route to estimating $g^{(2)}(\tau)$.

The analysis of physiological data has traditionally been based on the sequence of intervals of an idealized point process obtained from the data. In general it is difficult to associate correlation exhibited in the sequence of intervals with correlation in real time. The association is more readily made, however, for relatively regular point processes where the refractoriness is large, and the average interevent time is small compared to the temporal scale of the fluctuations. These properties appear to apply for the sequence of heartbeats. However, a relationship between correlation in the sequence of intervals and correlation in the sequence of counts cannot be formulated in a general manner since it depends on the details of the underlying point process, which remains to be determined for the heart rate data.

The analyses presented here were based on the entire duration of each of the 27 recordings (ranging from 20 to 24 hours). It will be of interest to determine how the results depend on the duration and local variations in the data, and we expect to carry out such a study in the near future.

6. ACKNOWLEDGMENTS

We thank Ary Goldberger and David Rigney for valuable discussions, and for providing us with the data on which this paper is based. We are grateful to Steven Lowen and Conor Heneghan for insights and assistance. This work was supported by the Office of Naval Research under Grant N00014-92-J-1251 and by the National Science Foundation through the Center for Telecommunications Research at Columbia University.

7. APPENDIX A: THE POISSON POINT PROCESS

The homogeneous Poisson point process (HPP) is ubiquitous in point-process theory, playing the role that the Gaussian process plays in the study of continuous time stochastic processes.³² The HPP is memoryless: the occurrence of an event at time t_0 is independent of the presence or absence of events at any other time $t \neq t_0$. Because of this property both $\{\tau_i\}$ and $\{N_i\}$ form sequences of independent, identically distributed (iid) random variables. These processes are completely characterized by their respective amplitude probability density and mass functions. For the process $\{\tau_i\}$ generated from an HPP, the interevent-time probability density function is the exponential function

$$p_\tau(\tau) = \lambda e^{-\lambda\tau}, \quad (\text{A1})$$

where λ is the average number of events per unit time. The interevent-time mean and variance are $\langle \tau \rangle = 1/\lambda$ and $\text{var}(\tau) = 1/\lambda^2$, respectively. The probability mass function $p_N(N;T)$ of the amplitude of the process of counts $\{N_i\}$ is the Poisson distribution

$$p_N(N;T) = \frac{(\lambda T)^N e^{-\lambda T}}{N!}. \quad (\text{A2})$$

For all counting times T , the mean is equal to the variance of the number of counts, so that $\langle N \rangle = \text{var}(N) = \lambda T$.

The dead-time-modified Poisson point process (DTMP) results from the imposition of absolute refractoriness (dead time) on the HPP. The probability density function of the interevent times $p_\tau(\tau)$ retains its exponential form, but is truncated due to the dead time,

$$p_\tau(\tau) = \begin{cases} 0 & \tau < \tau_d \\ \lambda e^{-\lambda(\tau-\tau_d)} & \tau \geq \tau_d, \end{cases} \quad (\text{A3})$$

where τ_d is the dead time and λ is the rate before dead time is imposed. The process remains renewal; the interevent times are independent and identically distributed. However, the possible overlap of the refractory period across the boundary between adjacent counting intervals results in correlation between the two counts; the counts $\{N_i\}$, therefore, are no longer independent. Furthermore, refractoriness imposes regularity on the point process, thereby reducing $\text{var}(N)$ relative to $\langle N \rangle$. This is easily seen by considering the limit of large initial rate λ ; the imposition of dead time produces a regular, periodic series of impulses with very low count variance. The count variance-to-mean ratio, or Fano factor, provides a convenient measure of the degree of regularity of the underlying point process.

8. APPENDIX B: RELATIONSHIP BETWEEN THE FANO-FACTOR TIME CURVE AND THE NORMALIZED COINCIDENCE RATE

The normalized coincidence rate $g^{(2)}(\tau)$ is defined as

$$g^{(2)}(\tau) \equiv \frac{\Pr\{\mathcal{E}(t, t+dt) \text{ and } \mathcal{E}(t+\tau, t+\tau+dt)\}}{\Pr\{\mathcal{E}(t, t+dt)\}\Pr\{\mathcal{E}(t+\tau, t+\tau+dt)\}}, \quad (\text{B1})$$

where $\mathcal{E}(x, y)$ denotes the occurrence of an event in the interval (x, y) . The function $g^{(2)}(\tau)$ for a point process is analogous to the autocorrelation function for a continuous stochastic process. For an HPP, $g^{(2)}(\tau) = 1$ for all τ . Teich^{36,38} discusses connections among the FFC, the normalized coincidence rate $g^{(2)}(\tau)$, and the power spectral density for a general stationary point process with constant rate.

A normalized coincidence rate of the form

$$g^{(2)}(\tau) = \begin{cases} 0 & |\tau| < \tau_d \\ 1 & \tau_d \leq |\tau| \leq \tau_r \\ 1 + \frac{\delta}{\mu} \left(\frac{|\tau|}{\tau_r} \right)^{\alpha-1} & |\tau| > \tau_r \end{cases} \quad (\text{B2})$$

incorporating both refractoriness and power-law correlation in a simple way,³⁸ corresponds to an FFC of the form

$$F(T) = \begin{cases} 1 - \mu T & T < \tau_d \\ 1 - \mu \tau_d \left(2 - \frac{\tau_d}{T} \right) & \tau_d \leq T \leq \tau_f \\ 1 - \mu \tau_d \left(2 - \frac{\tau_d}{T} \right) + \frac{2\delta\tau_f}{\alpha(\alpha+1)} \left\{ \left(\frac{T}{\tau_f} \right)^\alpha + \alpha \left(\frac{\tau_f}{T} \right) - (\alpha+1) \right\} & T > \tau_f \end{cases} \quad (B3)$$

In these equations μ is the rate of the observed point process (after refractoriness modification), τ_d is the refractoriness (dead time), τ_f is the fractal (power-law) onset time, α is the dominant long-counting-time fractional power-law exponent in the Fano factor, and δ is the excess coincidence rate. Since the FFC appears to be a robust statistical measure, this connection provides a convenient way of estimating $g^{(2)}(\tau)$, which is difficult to obtain directly from the data because of the sparseness of events in very short counting times. In particular, a Fano-factor time curve with a power-law dependence $F(T) \sim T^\alpha$ implies that the underlying point process has a power-law coincidence rate $g^{(2)}(\tau) \sim |\tau|^{\alpha-1}$ and power spectral density that behaves as $S(f) \sim f^{-\alpha}$.

9. REFERENCES

1. A. Babloyantz and A. Destexhe, "Is the normal heart a period oscillator?," *Biol. Cybern.*, Vol. 58, pp. 203-211, 1988.
2. R.D. Berger, S. Akselrod, D. Gordon, and R.J. Cohen, "An efficient algorithm for spectral analysis of heart rate variability," *IEEE Trans. Biomed. Engr.*, Vol. BME-33, pp. 900-904, 1986
3. J.T. Bigger, J.L. Fleiss, R.C. Steinman, L.M. Rolnitzky, R.E. Kleiger and J.N. Rottman, "Frequency domain measures of heart period variability and mortality after myocardial infarction," *Circ.*, Vol. 85, pp. 164-171, 1992.
4. W. Feller, "The asymptotic distribution of the range of sums of independent random variables," *Ann. Math. Stat.*, Vol. 22, pp. 427-432, 1951.
5. R. Freeman, J.P. Saul, M.S. Roberts, R.D. Berger, C. Broadbridge, and R.J. Cohen, "Spectral analysis of heart rate in diabetic autonomic neuropathy," *Arch. Neurol.*, Vol. 48, pp. 185-190, 1991.
6. L. Glass, A. Shrier and J. Belair, "Chaotic cardiac rhythms," in *Chaos*, edited by Arun V. Holden, pp. 237-256, Princeton University Press, Princeton, 1986.
7. A.L. Goldberger, D.R. Rigney, J. Mietus, E.M. Antman, and S. Greenwald, "Nonlinear dynamics in sudden cardiac death syndrome: heartrate oscillations and bifurcations," *Experientia*, Vol. 44, pp. 983-987, 1988.
8. A.L. Goldberger, "Nonlinear dynamics, fractals and chaos: applications to cardiac electrophysiology," *Ann. of Biomed. Engr.*, Vol. 18, pp. 195-198, 1990.
9. A.L. Goldberger, D.R. Rigney, and B.J. West, "Chaos and fractals in Human Physiology," *Sci. Am.*, Vol. 262, pp. 42-49, 1990.
10. P. Grassberger and I. Procaccia, "Measuring the strangeness of strange attractors," *Physica D*, Vol. 9, pp. 189-208, 1983.
11. F. Grüneis, M. Nakao, Y. Mizutani, M. Yamamoto, M. Meesmann, and T. Musha, "Further study on 1/f fluctuations observed in central single neurons during REM sleep," *Biol. Cybern.*, Vol. 68, pp. 193-198, 1993.
12. M.R. Guevara and A. Shrier, "Rhythms produced by high-amplitude periodic stimulation of spontaneously beating aggregates of embryonic chick ventricular myocytes," in *Mathematical approaches to cardiac arrhythmias*, edited by J. Jalife, pp. 11-22, New York Academy of Sciences, New York, 1990.
13. S. Guzzetti, E. Piccaluga, R. Casati, S. Cerutti, F. Lombardi, M. Pagani, and A. Malliani, "Sympathetic predominance in essential hypertension: a study employing spectral analysis of heart rate variability," *J. Hypertens.*, Vol. 6, pp. 711-717, 1988.
14. B. Hoop, L. Liebovitch, and H. Kazemi, "Rescaled range analysis of resting respiration," *Chaos*, Vol. 3, pp. 27-29, 1992.
15. H.E. Hurst, "Long-term storage capacity of reservoirs," *Trans. Amer. Soc. Civil Eng.*, Vol. 116, pp. 770-808, 1951.
16. D.T. Kaplan and R.J. Cohen, "Searching for chaos in fibrillation," in *Mathematical approaches to cardiac arrhythmias*, edited by J. Jalife, pp. 367-374, New York Academy of Sciences, New York, 1990.

17. K.A. Kluge, R.M. Harper, V.L. Schechtman, A.J. Wilson, H.J. Hoffman, and D.P. Southall, "Spectral analysis assessment of respiratory sinus arrhythmia in normal infants and infants who subsequently died of sudden infant death syndrome," *Ped. Res.*, Vol. 24, pp. 677-682, 1988.
18. M. Kobayashi and T. Musha, "1/f fluctuation of heartbeat period," *IEEE Trans. Biomed. Engr.*, Vol. BME-29, pp. 456-457.
19. L.S. Liebovitch, "Introduction to the properties and analysis of fractal objects, processes, and data", in Advanced Methods of Physiological System Modeling, Vol. 2, edited by V.Z. Marmarelis, Plenum, 1989.
20. L.S. Liebovitch and T. Toth, "A fast algorithm to determine fractal dimensions by box counting," *Phys. Lett. A*, Vol. 141, pp. 386-390, 1989.
21. M. Lishner, S. Akselrod, V. Mor Avi, O. Oz, M. Divon, and M. Ravid, "Spectral analysis of heart rate fluctuations. A non-invasive, sensitive method for the early diagnosis of autonomic neuropathy in diabetes mellitus," *J. Auton. Nerv. Sys.*, Vol. 1987, pp.119-125, 1987.
22. S.B. Lowen and M.C. Teich, "Estimating the dimension of a fractal point process," in *SPIE Proc.*, Vol. 2036, to be published, 1993.
23. B.B. Mandelbrot, The Fractal Geometry of Nature, Freeman, New York, 1983.
24. M. Meesmann, F. Grüneis, P. Flachenecker, and K.D. Kniffki, "A new method for analysis of heart rate variability: counting statistics of 1/f fluctuations," *Biol. Cybern.*, Vol. 86, pp. 299-396, 1993.
25. G.B. Moody and R.G. Mark, "Development and evaluation of a 2-lead ECG analysis program," *IEEE Comp. in Cardiol.*, Vol. 9, pp. 39-44, 1982.
26. A. Papoulis, Probability, Random Variables, and Stochastic Process, McGraw-Hill, New York, 1984.
27. C.-K., Peng, J. Mietus, J.M. Hausdorff, S. Havlin, H.E. Stanley, and A.L. Goldberger, "Long-range anti-correlations and non-gaussian behavior of the heartbeat," *Phys. Rev. Lett.*, Vol. 70, pp. 1343-1364, 1992.
28. W.H. Press, B.P. Flannery, S.A. Teukolsky, and W.T. Vetterling, Numerical Recipes in C, Cambridge University Press, Cambridge, 1988.
29. N.L. Powers and B.J. Salvi, "Discharge rate fluctuations in the auditory nerve of the chinchilla," Abstracts of the XIV Midwinter Research Meeting, Assoc. for Res. in Otolaryngology, p. 411, 1991.
30. N.L. Powers and B.J. Salvi, "Comparison of discharge rate fluctuations in the auditory nerve of chickens and chinchillas," Abstracts of the XV Midwinter Research Meeting, Assoc. for Res. in Otolaryngology, p. 101, 1992.
31. F. Ravelli and R. Antolini, "Complex dynamics underlying the human electrocardiogram," *Biol. Cybern.*, Vol. 67, pp. 57-65, 1992.
32. B.E.A. Saleh and M.C. Teich, "Multiplied-Poisson noise in pulse, particle, and photon detection," *Proc. of the IEEE*, Vol. 70, pp. 229-245, 1982.
33. J.P. Saul, P. Albrecht, R.D. Berger, and R.J. Cohen, "Analysis of long term heart rate variability: methods, 1/f scaling and implications," *IEEE Comput. in Cardiol.*, Vol. 14, pp. 419-422, 1988.
34. H. Shono, M. Yamasaki, M. Muro, M. Oga, Y. Ito, K. Shimomura, and H. Sugimori, "Chaos and fractals which 1/f spectrum below 10^2 Hz demonstrates full-term fetal heart rate changes during active phase," *Early Human Dev.*, Vol. 27, pp.111-117, 1991.
35. M.C. Teich and S.M. Khanna, "Pulse-number distribution for the neural spike train in the cat's auditory nerve," *J. Acoust. Soc. Am.*, Vol. 77, pp. 1110-1128, 1985.
36. M.C. Teich, "Fractal character of the auditory neural spike train," *IEEE Trans. Biomed. Eng.*, Vol. 36, pp. 150-160, 1989.
37. M.C. Teich, R.G. Turcott, and S.B. Lowen, "The fractal doubly stochastic Poisson point process as a model for the cochlear neural spike train," in The Mechanics and Biophysics of Hearing (Lecture Notes in Biomathematics, Vol. 87), edited by P. Dallos, C.D. Geisler, J.W. Matthews, M.A. Ruggero, and C.R. Steele, pp. 354-361, Springer-Verlag, New York, 1990.
38. M.C. Teich, "Fractal neuronal firing patterns," in Single Neuron Computation, edited by T. Mckenna, J. Davis, and S. Zornetzer,

pp. 589-625, Academic, Boston, 1992.

39. R.G. Turcott, S.B. Lowen, E. Li, D. Johnson, C. Tsuchitani, and M.C. Teich, "A nonstationary Poisson point process describes the sequence of action potentials over long time scales in Lateral-Superior-Olive auditory neurons," *Biol. Cybern.*, in press, 1993.

40. R.G. Turcott, P.D. Barker, and M.C. Teich, "Long-duration correlation in the sequence of action potentials of the descending contra-lateral movement detector of the locust," submitted for publication, 1993.

41. A. Wolf, J.B. Swift, H.L. Swinney, and J.A. Vastano, "Determining Lyapunov exponents from a time series," *Physica D*, Vol. 16, pp. 285-317, 1985.

42. J.P. Zbilut, G. Mayer-Kress, P.A. Sobotka, M. O'Toole, and J.X. Thomas, Jr., "Bifurcations and intrinsic chaotic and $1/f$ dynamics in an isolated perfused rat heart," *Biol. Cybern.*, Vol. 61, pp. 371-378, 1989.